

Electronics for the Neutral Current Detection Array at the Sudbury Neutrino Observatory

Charles A. Duba

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

Doctor of Philosophy

University of Washington

2006

Program Authorized to Offer Degree:
Department of Physics

University of Washington
Graduate School

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

Charles A. Duba

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:

R.G. Hamish Robertson

Reading Committee:

R.G. Hamish Robertson

John Wilkerson

Thompson Burnett

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 the 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, or to the author.

Signature _____

Date _____

University of Washington

Abstract

Electronics for the Neutral Current Detection Array at the Sudbury Neutrino Observatory

Charles A. Duba

Chairperson of the Supervisory Committee:
Professor R.G. Hamish Robertson
Department of Physics

When the Sudbury Neutrino Observatory (SNO) entered its third stage of operation, the Neutral Current Detectors (NCDs) became the prominent means of measuring the flavor-independent neutrino flux. Neutrons thermalized within SNO's D₂O give an analysis-independent measure of the total neutrino flux, particularly since the NCDs recognize neutrons on an event-by-event basis. We determined and evaluated the electronics requirements for the complete NCD array and demonstrated their ability to measure the neutron flux within SNO. Methods of evaluating detector thresholds, energy linearity, and dead time are presented and applied to the electronics. Predications for array performance were enabled by writing code support for the NCD array in SNOMAN, the collaboration's Monte-Carlo simulation package. The performance of the array has been monitored and verified in the interval before deployment known as the 'cool-down' phase. Methods for NCD contaminant self-analysis are presented, applied and evaluated, providing information on the background contributions. Neutrons are separated from low-energy alphas in a systematic manner to verify that the operation of

the entire NCD data acquisition. The potential of the spare neutron detectors at SNO to be used as a low-background neutron-based supernova array in SNOLAB is discussed.

TABLE OF CONTENTS

List of figures	vi
List of tables	x
Chapter 1: The Solar Neutrino Problem and S.N.O.	1-1
1.1 The Solution to the Solar Neutrino Problem	1-5
1.2 “Second Generation” Solar Neutrino Experiments	1-6
1.2.1 Kamiokande	1-7
1.2.2 SAGE	1-7
1.2.3 Gallex and GNO	1-8
1.3 Super-Kamiokande	1-8
1.4 The Sudbury Neutrino Observatory	1-9
1.4.1 SNO D ₂ O Results	1-13
1.4.2 SNO NaCl Results	1-16
1.5 Current Status of Neutrino Physics	1-18
1.5.1 KamLAND	1-18
1.5.2 SNO NCD Phase	1-20
1.6 3-neutrino family mixing	1-24
1.7 SNO NCDs	1-26
1.7.1 NCD purity	1-27
1.7.2 NCD Cooldown	1-29
1.8 Supernovae and neutrinos	1-30
1.8.1 NCDs and supernovae	1-39
Chapter 2: NCD Array and Electronics	2-41
2.1 NCDs as proportional counters	2-43
2.1.1 Delay line	2-50
2.1.2 NCD Cable	2-52
2.1.3 Chemical Vapor Deposition	2-52

2.1.4 Electropolishing	2-53
2.1.5 Gas Fill	2-57
2.2 Electronics Specifications	2-58
2.3 NCDs as an array	2-60
2.4 NCD signal	2-63
2.5 Preamplifier	2-67
2.6 Shaper/ADC	2-68
2.6.1 Shaper/ADC requirements.....	2-69
2.6.2 Shaping.....	2-70
2.6.3 ADC	2-71
2.6.4 Shaper/ADC VME Communication protocol.....	2-71
2.6.5 VME Controller Issues	2-72
2.6.6 Shaper/ADC Scalers.....	2-73
2.7 Multiplexer/Digitizer	2-75
2.7.1 Multiplexer Design/Function.....	2-75
2.7.2 Signal Multiplexing.....	2-75
2.7.3 Multiplexer Log Amp.....	2-77
2.7.4 Digitization	2-78
2.7.5 MUX Controller Board	2-78
2.7.6 MUX DACs Board.....	2-81
2.7.7 MUX Communication through the IP408	2-86
2.8 Event tagging.....	2-90
2.8.1 SNO Global Trigger	2-91
2.8.2 MTC/D Operation Overview	2-91
2.8.3 NCD Trigger System.....	2-92
2.8.4 NCD Logic Board	2-96
2.8.5 MUX Trigger	2-97

2.9 NCD Trigger Card	2-98
2.9.1 NCD Trigger Card Event Logic.....	2-98
2.9.2 NCD Trigger Card Livetime Counters.....	2-100
2.10 HV Control System.....	2-102
2.10.1 Proportional Counter Gain and Voltage.....	2-103
2.10.2 HV Control Board	2-104
2.10.3 HV DACs Board	2-106
2.10.4 HV interface board.....	2-106
2.10.5 HV Distribution Panel.....	2-107
2.10.6 Spellman MP3P24 HV Power Supply.....	2-109
2.11 Calibration Pulse Generation.....	2-109
2.11.1 Agilent 33120 Pulser	2-109
2.11.2 Analog Characteristics of the Pulser Distribution System	2-110
2.11.3 Digital Characteristics of the Pulser Distribution System.....	2-110
2.12 Low Voltage Power	2-112
2.12.1 VME Power	2-112
2.12.2 Low Voltage Power Distribution and Power Racks	2-113
Chapter 3: NCD Electronics Verification and Control.....	3-114
3.1 NCD Electronics Trigger Thresholds	3-114
3.1.1 NCD ADC/Shaper Trigger Thresholds.....	3-115
3.1.2 NCD Mux Trigger Thresholds.....	3-117
3.2 NCD Electronics Linearity.....	3-119
3.2.1 NCD Gain and Linearity Calibrations.....	3-119
3.2.2 NCD MUX Log Amp Calibrations.....	3-123
3.2.3 Log Amp Non-linearity.....	3-125
3.2.4 NCD ADC/Shapers and Ballistic Deficit	3-127
3.3 NCD Electronics Deadtime Measurements.....	3-129

3.3.1 NCD Shaper/ADC Scaler Deadtime	3-129
3.3.2 NCD Shaper/ADC Deadtime.....	3-129
3.3.3 NCD MUX Trigger Deadtime	3-130
3.3.4 NCD Digitization Deadtime	3-131
3.4 NCD Livetime Counters	3-132
3.5 Random Pulser.....	3-133
3.6 NCD Electronics Model	3-135
3.7 NCD DAQ: ORCA	3-137
3.7.1 DAQ Control.....	3-139
3.7.2 Data Construction	3-140
Chapter 4: Neutrons and NCDs.....	4-142
4.1 Other neutron sources in SNO.....	4-143
4.2 Neutrons from supernovae during cooldown	4-144
Chapter 5: Approaches to NCD Backgrounds	5-153
5.1 Beta-Gamma and Gamma-Gamma Coincidence.....	5-156
5.2 Monte-Carlo	5-169
5.3 NCD contaminant alpha-alpha self-analysis	5-175
5.3.1 Alpha-Alpha Coincidence Cooldown Data.....	5-188
5.3.2 In-Situ Alpha-Alpha Coincidence Measurements	5-192
5.4 In-situ neutron background measurements	5-198
5.5 Neutrons from Solar Neutrino Interactions.....	5-202
Chapter 6: Conclusions.....	6-209
6.1 Verification of NCD Electronics	6-210
6.2 NCD array backgrounds.....	6-211
Appendix A: NCD Communication Tables.....	6-222
Appendix B: Snostream.....	6-238
Appendix C: MUX Schematic	6-248

VITA.....	6-250
BIBLIOGRAPHY	6-213

LIST OF FIGURES

Figure Number	
Figure 1-1: The p-p chain.....	1-2
Figure 1-2: Solar Neutrino Spectrum. Figure is from []	1-3
Figure 1-3: Scale drawing of the SNO cavity with NCDs inside the inner D ₂ O.....	1-10
Figure 1-4: $\nu_e + d \rightarrow p + p + e$ (1.44 MeV) (CC)	1-11
Figure 1-5: $\bar{\nu}_e + d \rightarrow n + n + e^+$ (4.03 MeV) (ACC).....	1-12
Figure 1-6: $\nu_x + e \rightarrow \nu_x + e$ (0 MeV) (ES)	1-12
Figure 1-7: $\nu_x + d \rightarrow p + n + \nu_x$ (2.22 MeV) (NC).....	1-12
Figure 1-8: Solar neutrino fluxes combining best fits from SNO's D ₂ O run with SK... 1-14	1-14
Figure 1-9: Estimated solar neutrino fluxes after the second SNO results..	1-15
Figure 1-10: Allowed mass-squared vs mixing angle after SNO salt phase	1-16
Figure 1-11: Neutrino flux deficit due to oscillations as observed at KamLAND.	1-19
Figure 1-12: Two-neutrino oscillation parameter allowed regions	1-20
Figure 1-13: Effects of narrowing CC/NC and day-night asymmetry at SNO.	1-21
Figure 1-14: Current allowed parameter space for 3-neutrino mixing	1-23
Figure 1-15: Neutrino flavor and mass spectra for the two possible mass hierarchies... 1-26	1-26
Figure 1-16: Fermi Dirac Distribution of escaping thermal neutrinos.....	1-36
Figure 1-17: Potential 1987a neutrino energies as a function of time	1-38
Figure 2-1: Basic NCD Electronics Flowchart.....	2-42
Figure 2-2: NCD Cross-section (not to scale)	2-44
Figure 2-3: Basic Deployed NCD Configuration	2-46
Figure 2-4: The Intercounter Coupler	2-47
Figure 2-5: NCD 'Bell-End' Cable-Counter Connection.....	2-48

Figure 2-6: Flexible Nickel-Teflon Resistive Counter Coupler	2-50
Figure 2-7: Open NCD Preamplifier with labeled connectors	2-68
Figure 2-8: NCD Electropolishing Rig	2-56
Figure 2-9: NCD Electronics Extended Flow Chart	2-62
Figure 2-10: Fully Contained Triton and Proton Drift.....	2-64
Figure 2-11: Partially Contained Triton and Proton Event	2-65
Figure 2-12: Fully Contained NCD Alpha Event	2-65
Figure 2-13: Characteristic NCD microdischarge scope trace	2-66
Figure 2-14: "Seattle" Event on String J3	2-67
Figure 2-15: NCD Shaper/ADC board.....	2-69
Figure 2-16: MUXCB with labeled connections	2-79
Figure 2-17: MUX DACs board with connections	2-82
Figure 2-18: IP408 Block Diagram.....	2-87
Figure 2-19: BSB/CSB Communications Handshake Protocol.....	2-89
Figure 2-20: NCD Trigger Logic.....	2-93
Figure 2-21: Shaper/ADC Trigger Logic.....	2-95
Figure 2-22: NCD MUX Trigger Timing Diagram	2-97
Figure 2-23: NCD Trigger Card Logic	2-99
Figure 2-24: Trigger Card Livetime Counter Logic	2-101
Figure 2-25: HV Distribution Panel with 3 rows tied together	2-108
Figure 2-26: NCD Rack Layouts	2-113
Figure 3-1: ECA NCD Calibration Run across 20 channels.	3-116
Figure 3-2: ECA MUX threshold calibration run across 20 channels	3-118
Figure 3-3: Shaper Channels Occupied During Typical ECA Linearity Test.....	3-120
Figure 3-4: Shaper/ADC Deviations from Linearity During ECA.....	3-121
Figure 3-5: MUX/Scope Deviations from Linearity for an extended ECA run.	3-122
Figure 3-6: Theoretical Standard NCD Multiplexer Calibration Pulse	3-123

Figure 3-7: Digitized Test Pulse after MUX Log Amplifier	3-124
Figure 3-8: Log amp response at different offsets	3-126
Figure 3-9: Log amp linearity fit residuals at 10 mV offset.....	3-127
Figure 3-10: Short Small Pulse Ballistic Deficit as a Function of Pulse Duration....	3-128
Figure 3-11: Separation, in seconds, of consecutive NCD Shaper events	3-130
Figure 3-12: Separation of consecutive MUX events	3-131
Figure 3-13: Separation between consecutive digitized events on separate scopes. .	3-132
Figure 3-14: Simplified NCD Electronics Model for Digitized Pulses.	3-136
Figure 3-15: ORCA during operation.	3-137
Figure 3-16:ORCA.....	3-139
Figure 3-17: Building and merging SNO and NCD data streams	3-140
Figure 4-1: In-rack neutron measurement configuration.	4-146
Figure 5-1: ^{232}Th Decay Chain.	5-154
Figure 5-2: ^{238}U Decay Chain.....	5-155
Figure 5-3: $^{208}\text{Tl} \rightarrow ^{208}\text{Pb}$ Decay Scheme	5-156
Figure 5-4: Prototype NCD Neutron Spectrum fit to a theoretical spectrum.....	5-160
Figure 5-5: Beta-Gamma Coincidence TAC Spectrum Setup.....	5-161
Figure 5-6: TAC spectrum between NaI and NCD events.....	5-162
Figure 5-7: Beta-Gamma coincidence setup	5-163
Figure 5-8: Energy Spectrum of Associated NCD Events	5-164
Figure 5-9: Position of Thorium Tests Along NCD Body	5-165
Figure 5-10: Single-counter gamma detection efficiency	5-167
Figure 5-11: Expected gamma detection efficiency along a string of counters	5-168
Figure 5-12: SNOMAN NCD Geometry Regions and Overlaps	5-171
Figure 5-13: Energy deposited within NCDs from ^{212}Bi decay	5-173
Figure 5-14: Energy deposited in counter during >10 NHit event.	5-174
Figure 5-15: Simplified Counter Cross-cut	5-178

Figure 5-16: Probability of Alpha Detection at Fraction of Max Depth.....	5-179
Figure 5-17: Probability of Multiple Alpha Detections	5-181
Figure 5-18: Alpha Energy Profile for Varying Source Depth	5-186
Figure 5-19: NCD Cooldown Underground Spectrum	5-188
Figure 5-20: Deployed NCD Configuration.....	5-192
Figure 5-21: Digitized Event Scan Sheet	5-194
Figure 5-22: Cleaned ADC Spectra for all channels during open data set	5-196
Figure 5-23: Total open data set Shaper/ADC events below 1MeV	5-200
Figure 5-24: Sum of 29 calibrated channel spectra over the Open Data Set.....	5-202
Figure 5-25: Calibrated ^4He counter spectra summed over the Open Data set	5-203
Figure 5-26: The N2 string during ^{24}Na calibration runs.. ..	5-205
Figure 5-27: Shaper ADC spectrum for alphas between 2.5 MeV and 850 keV	5-206

LIST OF TABLES

Table Number	
Table 2-1: Deployed NCD Counter Summary	2-45
Table 2-2: Deployed NCD String Summary	2-45
Table 5-1: Common Gammas from ^{208}Tl decay	5-157
Table 5-2: Betas Emitted from ^{208}Tl Decay.....	5-158
Table 5-3: SNOMAN NCD Regions	5-170
Table 5-4: Summary of neutron run channels cut.....	5-201

ACKNOWLEDGMENTS

I would like to thank everyone for their understanding, assistance, guidance, and patience during my long bout with the beast. As a reader of this obscure work, you are at least partially responsible for this. Although the world might not thank you for your contributions, I certainly do, even if I fail to mention you by name in the following paragraphs.

Hamish, as my advisor, your suffering made this possible, and I thank you with the utmost humility. I need to thank the academics that guided me around the lab: John Wilkerson, Tom Steiger, Jos VanSchagen, and Joe Germani. I would never have succeeded without the assistance of my senior graduate students Rushdy Ahmed, Jim Beck, Mike Browne, Mike Kelley and Alan Poon. Essential to my continuing sanity were my peer graduate students, who managed to suffer through and with me: Karsten Heeger, Erik Mohrmann, Bruce Henry, Piet Mumm, Adam Cox, Kathryn Miknaitis, Jeff Reid, Rebecca Hoffenberg, and others too numerous to list. There were many others that don't fall into these categories that were required elements, from technical staff to emerging researchers: Mark Howe, Allan Myers, Tim Van Wechel, Duncan Prindle, Greg Harper, Sean McGee and Joe Formaggio. I will never forget the care and dedication that Barb Fulton gave to the cause.

In the end, my life is not a product of science, but of people. The people to whom I am most indebted are my parents Al and Lucy, without whose lifelong dedication and belief, I would never have succeeded. I cannot forget, or at least my brothers will not allow me forget, the fact that they received their doctorates before me. The associated fear of family gatherings without degree completion provided another degree of motivation.

There are those essential friends whose time and patience allowed me to graduate with the remains of my sanity: Eric Kansa, Markus Weckler, Harry Hwang, Gordon Dutrisac, Patrick Culacurcio, Web Pierce, David Pearton, Rafael Vargas, Andy Estes, Amanda Mobernd, Kim Hawley, and Linnea Mobernd. I must also take the time to personally thank and apologize to Annette Plesner, who suffered greatly on my accord.

In the end, the support of Masayo Arakawa pulled me through the darkest times. It was your loving support that allowed me to pass through the final hurdles, and words cannot express this, my deepest gratitude.

DEDICATION

To my grandparents: you are missed.

Chapter 1 The Solar Neutrino Problem and S.N.O.

To expect too much is to have a sentimental view of life and this is a softness that ends in bitterness. Flannery O'Connor

When Pauli first postulated the existence of a neutrino¹, it was to explain nothing. That is to say, to explain the failure of an experiment to meet his expectations. The neutrino was created to account for missing energy and missing angular momentum, as a phantom equalizer of conservation laws. Since the very start of neutrino studies, its weakly interacting nature has hindered every attempt to unearth its properties.

What is a hindrance in detection is a boon in information. Terrestrial observers have for many years noted the photons that come from the Sun. Unfortunately, despite their ubiquity, these particles only provide direct information about the Sun's surface. Combining stellar luminosity measurements with obvious gravitational calculations, astrophysicists were quick to recognize that, although gravitation may counterbalance the outward pressure from the core, it could not explain the longevity of stars².

Stellar fusion is required to explain the power and lifetimes of the stars. The 'p-p chain' was eventually proposed as the method of light-element fusion that would power a

star the size of our Sun. The p-p chain, shown in Figure 1-1, provides the mechanism through which the sun maintains the pressure to counterbalance gravitational collapse.

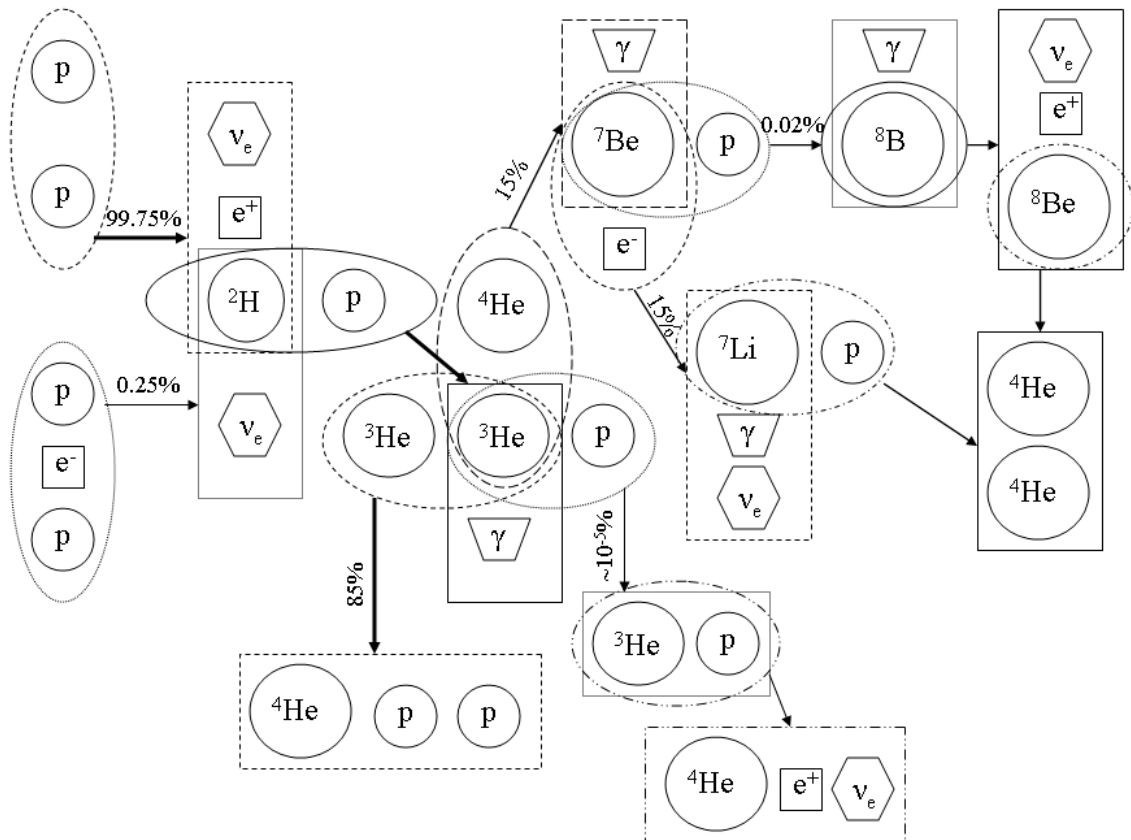


Figure 1-1: The p-p chain

Unfortunately for early astrophysicists, there was no method through which this process could be directly observed—only the core of the Sun is hot and dense enough for significant fusion to occur, and no particle that interacts electromagnetically can escape the core without significant scattering. To a weakly-interacting neutrino, however, the Sun looks nearly transparent—the electron density gradient provides the sole significant effect. Therefore, neutrinos created during stellar fusion processes travel directly to and

through the earth, enabling direct observation. The terrestrial neutrino-energy spectrum from the sun depends on the abundance, density and temperature of stellar core materials. Some of the neutrinos from the sun are sufficiently energetic that they can be measured directly. The expected spectrum of solar neutrino energies in the Standard Solar Model^{3,4} (SSM) is shown in Figure 1-2.

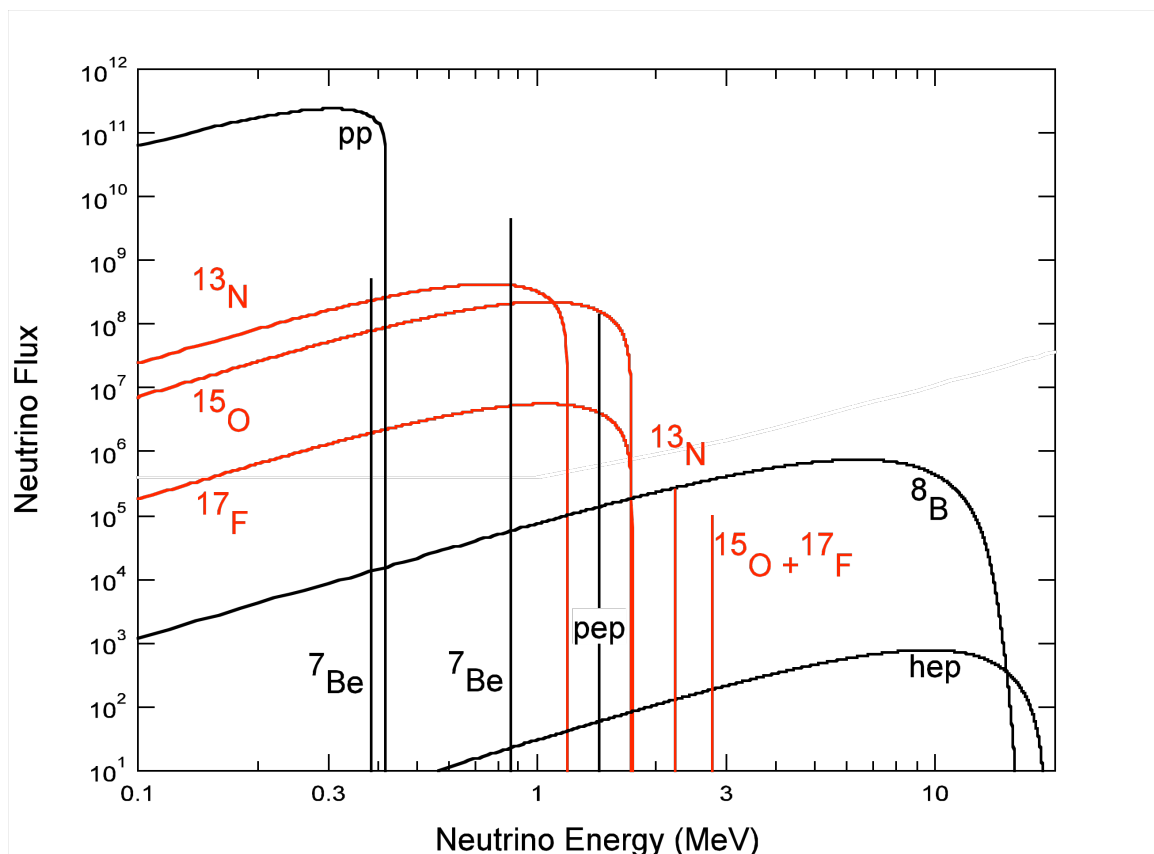



Figure 1-2: Solar Neutrino Spectrum. Figure combines [3] and [4]

The first experimental attempt to measure solar neutrino flux⁵ was done with 1000 gallons of perchlorethylene placed under 2300 feet of limestone. In this prototype experiment, Ray Davis, Jr. and his collaborators estimated the requirements for an accurate measurement. Since this first experiment was only able to place an upper limit

on neutrino flux, John Bahcall's theoretical calculations⁶ were used to predict a useful detector depth and size. The resulting experiment was the first successful stellar neutrino detector⁷ and was built in the Homestake Gold Mine in Lead, South Dakota, with 520 tons of chlorine.

Electron neutrinos were indeed observed through the neutrino-induced inverse beta decay reaction:



Unexpectedly, the total quantity of neutrinos was significantly lower than predicted. The lower bound of the expected rate was 8 Solar Neutrino Units[†], while the observed rate was less than 3 SNU. Uncertainties in both the measurement and prediction were analyzed, but the discrepancy remained. In order to determine where the truth lay, new experiments were built and solar models re-examined.

Each of the subsequent experiments found similar effects as the Homestake results; the solar neutrino flux at the earth was lower than was predicted. While experimentalists verified the old results, theorists attempted to adjust solar models to account for the discrepancies. The array of solar models that were constructed had to obey the limits set forth by stellar lifetimes, luminosities, and even pressure wave propagation. Many unique models were proposed, but none could fill the difference

[†] 1 Solar Neutrino Unit (SNU)= 10^{-36} events per target atom per second

between theory and experiment. This well-defined and unexplained discrepancy was known as the Solar Neutrino Problem.

1.1 The Solution to the Solar Neutrino Problem

The neutrinos detected in the first few solar neutrino experiments were all of the electron “family”, or “flavor”. Electron neutrinos can interact with baryonic matter through one of three weak bosons: The W^+ , W^- and Z^0 . The charged current interaction, or interaction through the exchange of the W^+ and W^- bosons, is only possible with electron-flavor neutrinos. Neutrinos of all active flavors can interact by exchange of a neutral current, or Z^0 boson. Nuclear fusion and fission create electron-flavor neutrinos and anti-neutrinos, so the early neutrino detectors were designed around those flavors and searched for evidence of the charged-current interaction. If some neutrinos arriving from the sun had changed from electron neutrinos to a different flavor, the experiments of the day would be unable to detect them, and would register a deficit.

Bruno Pontecorvo first put forth the concept that neutrinos could oscillate between interaction states in 1957⁸, but the model of how neutrinos were changing flavor as they left the sun was first proposed by Mikheyev and Smirnov in 1985⁹, building on earlier work by Wolfenstein in 1978¹⁰. Neutrino oscillations require a few changes to the standard model of particle physics: One—neutrinos have mass. Two—their mass states and interaction states are not identical. Without mass, the old standard model neutrinos could only travel at the speed of light, and no time would pass in their frame, so no oscillations could occur.

The effective potential energy of a neutrino depends on the local density of leptons and the weak coupling constant, G_F . Since matter in most of the universe is made up of first-generation leptons and quarks, electron-neutrinos experience a matter-dependent effective energy difference from other families of neutrinos because W bosons can be exchanged with electrons. Therefore, an electron neutrino may start with a higher effective energy than a mu or tau neutrino due to electron densities within the sun, and it may “cross over” the effective energies of the mu or tau states on the way out. If the mass/energy state is not perfectly aligned with the interaction state, the neutrino flavor may change as the level crossing occurs. This effect removes many of the expected electron neutrinos, replacing them with either mu or tau neutrinos. The first few generations of neutrino detectors were unable to detect the latter kinds of neutrinos.

1.2 “Second Generation” Solar Neutrino Experiments

After the discovery of the solar neutrino problem, a long series of speculation and evaluation of data followed. The experimental results from Homestake and the theoretical developments of the interim years shaped the design of the next generation of solar neutrino detectors. These “second-generation” solar neutrino experiments utilized previously unavailable technologies and materials along with increased manpower and funding. It was more than 20 years before the “second-generation” of solar neutrino experiments began to report results.

1.2.1 Kamiokande


Kamiokande was originally a water-Cherenkov proton-decay experiment based in the Kamioka lead and zinc mine in Gifu prefecture, Japan. Its purpose evolved to include solar neutrino detection when lower thresholds were realized in 1987¹¹. The Kamiokande-II detector had a mass of 2140 (fiducial 680) metric tons of water, and relied on the electron scattering reaction:



Due to kinematics, the electron is emitted non-isotropically. This allowed the first directional observations of the sun and verification of the neutrino source. However, the detected solar neutrino flux once again fell short of contemporary theoretical predictions, and Kamiokande results were used to design the famous Super-Kamiokande experiment.

1.2.2 SAGE

Starting operation in January, 1990, The Soviet-American Gallium Experiment, SAGE¹², later renamed the Russian-American Gallium Experiment, uses a large mass of metal Ga (~60 metric tons) to detect solar neutrinos. SAGE relies on the reaction,



The presence and concentration of ⁷¹Ge, detected through its beta decay with a half-life of 11 days, indicates the neutrino flux. It is important to note that the threshold for interaction in the Gallium-based detectors is substantially lower than the previous

Chlorine- and water-based detectors. This allows the observation of the low-energy p-p neutrino spectrum, the dominant component of the sun's neutrino flux.


1.2.3 Gallex and GNO

Similar in size to SAGE, but different in method, Gallex¹³ used ^{71}Ga to see inverse beta-decay induced by incident low-energy neutrinos. Unlike SAGE, Gallex's 30.3 tons of gallium was contained in 100 tons of liquid GaCl_3 . Gallex began operation in May, 1991 in the Gran Sasso laboratory, and ended its run in 1997 before being replaced in 1998 by GNO, the Gallium Neutrino Observatory¹⁴. GNO ended operation in 2005 after a scintillating fluid drainage mishap in the Borexino experiment.

1.3 Super-Kamiokande

Kamiokande's search for proton decay, followed by solar and atmospheric neutrino detection is mirrored in the life of the Super-Kamiokande experiment. Super-Kamiokande uses 22,500 metric tons of water as a target, surrounded by 11,000 Photomultiplier Tubes. Just like Kamiokande, Super-K relies on elastic scattering for detection of ^8B neutrinos. Super-K's considerable increase in size and technology from Kamiokande provides the most accurate measurement of neutrino flux due to elastic scattering. Some of the most groundbreaking results from Super-K came from the observation of atmospheric neutrinos, which can interact via the charged current with H and O nuclei.

Atmospheric neutrinos come from the interaction of cosmic rays in the atmosphere. Since muons or antimuons that are created in the first reactions later decay into electrons or positrons, and lepton family number conservation is obeyed, there should always be a muon neutrino and anti-neutrino pair created, while only a single electron neutrino is expected:



Therefore the atmospheric flux of electron neutrinos is only expected to be half of the muon neutrino flux. However, SuperKamiokande was able to observe¹⁵ a statistically significant deviation from that expectation, implying neutrino oscillations. Subsequent measurements showed also a zenith-angle dependence of the deficit, strongly supporting the neutrino oscillation interpretation.

1.4 The Sudbury Neutrino Observatory

The Sudbury Neutrino Observatory¹⁶ is a heavy water Cherenkov neutrino detector housed at the 6800-foot level of the Creighton Mine just outside of Lively, Ontario. A scale drawing of the detector is shown in Figure 1-3.




Figure 1-3: Scale drawing of the SNO cavity with NCDs inside the inner D₂O.

The heart of the detector is 1000 metric tons of ultrapure, 99.92% enriched D₂O. SNO observes neutrinos through Electron Scattering (ES), Charged Current (CC), and Neutral Current (NC) interactions. SNO's array of 9600 photomultiplier tubes record the cone-shaped patterns of emitted Cherenkov light and permit an estimate of the energies based on the number of photoelectrons. From the arrival times of photons, the event position and direction can be reconstructed. SNO also has the unique ability to detect all non-sterile flavors of neutrinos by looking for neutrons dissociated from the D₂O. The dissociation is flavor-independent, since it is mediated by a chargeless virtual Z₀-boson, but requires more than 2.2 MeV of neutrino energy to overcome the binding energy of the neutron with the proton.

The neutron liberated by the neutral-current interaction is quickly thermalized within the D_2O and can be detected one of three ways. The first technique applied at SNO is to look for the neutron to capture on the deuterium and emit a characteristic 6.25-MeV gamma. The second method involves dissolving a chloride salt that will absorb the neutron and emit a cascade of gammas with a total energy of 8.6 MeV. Finally, the neutron can be picked up by 3He proportional counters and tallied independent of the photomultiplier tubes.

Although we neither expect nor have seen any significant antineutrino flux from the sun, SNO does have a significant cross section for anti-neutrinos above 4.03 MeV. An incoming electron anti-neutrino can interact with the proton in a deuteron, initiate an inverse beta decay, releasing a positron and two neutrons in the process.

The major neutrino detection reactions, their interaction thresholds, and their approximate detection thresholds for SNO are as follows:




Figure 1-4: $\nu_e + d \rightarrow p + p + e$ (1.44 MeV) (CC)




Figure 1-5: $\bar{\nu}_e + d \rightarrow n + n + e^+$ (4.03 MeV) (ACC)




Figure 1-6: $\nu_x + e \rightarrow \nu_x + e$ (0 MeV) (ES)




Figure 1-7: $\nu_x + d \rightarrow p + n + \nu_x$ (2.22 MeV) (NC)

1.4.1 SNO D₂O Results

In the period between November 2, 1999 and May 28, 2001, we took calibration, background, and production data with pure D₂O. Results from this first phase of SNO proved to have dramatic implications when combined with the results from SuperKamiokande.

The first results published by the SNO collaboration¹⁷ compared the charged current (CC) rate with the elastic scattering (ES) rate. Without neutrino flavor mixing, the expected electron-neutrino flux from the ⁸B pp-chain decay¹⁸ was $5.05 \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$. SuperKamiokande measured¹⁹ a flux of $2.32 \pm 0.03(\text{stat.})_{-0.07}^{+0.08}(\text{sys.}) \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$, followed by a statistically concordant $2.39 \pm 0.34(\text{stat.})_{-0.14}^{+0.16}(\text{sys.}) \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$ ⁸B electron-neutrino flux measured in SNO. SNO's unique sensitivity the charged-current reaction on deuterium yielded for the ⁸B electron-neutrino flux a value of $1.75 \pm 0.07(\text{stat.})_{-0.11}^{+0.12}(\text{sys.}) \pm 0.05(\text{theor.}) \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$.

If neutrinos of non-electron flavor (ν_μ, ν_τ) are present, then the ES reaction will record their presence via Z-exchange (Figure 1-6). Writing the measured ES rate as though it were due to a pure ν_e flux Φ_{ES} , the relative cross sections for W and Z exchange give:

$$\Phi_{ES} = \phi_e + 0.154\phi_{\mu\tau}.$$

The CC result is a direct measurement of ϕ_e , and these equations can be solved to obtain the total flux $\phi_e + \phi_{\mu\tau} = 5.44 \pm 0.99 \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$, in close agreement with the solar model prediction. This combination of the measurements from SNO and SuperKamiokande is shown graphically in Figure 1-8.




Figure 1-8: Solar neutrino fluxes at the earth combining best fits from SNO's D₂O run with SK. The error ellipses are the 68%, 95%, and 99% joint probability contours for ϕ_e and $\phi_{\mu\tau}$. Plot is from [17].

Ignoring possible sterile neutrino flux, this measurement indicates that two thirds of the neutrinos originating from solar ^8B decay are detected on earth as active, non-electron flavors. The simplest method to account for solar neutrino flavor conversion is to allow neutrinos to have mass states which are not identical to their flavor states, resulting in neutrino oscillations.

The second results from SNO²⁰, still during the D₂O phase, presented a direct measure of the neutral-current interactions. Neutrons capturing on deuterium result in a

6.25 MeV γ ray, and neutral current events release neutrons. Examining the excess of events around this energy resulted in a total flavor-independent active neutrino flux of $\phi_{\text{NC}} = 5.09_{-0.43}^{+0.44}(\text{stat.})_{-0.43}^{+0.46}(\text{sys.}) \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$. Unifying the neutral current flux with electron-neutrino flux gives a combined mu- and tau-neutrino flux of $\phi_{\mu\tau} = 3.41_{-0.45}^{+0.45}(\text{stat.})_{-0.45}^{+0.48}(\text{sys.}) \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$. The reduced error on the neutral current rate sets a tighter range for the allowed ratio of $\frac{\phi_e}{\phi_{\mu\tau}}$, as shown by the ellipses in Figure 1-9.

Note that the measured uncertainty in the NC from ${}^8\text{B}$ flux, represented by the solid blue in the graph is less than the uncertainty in the solar model.




Figure 1-9: Estimated solar neutrino fluxes after the second SNO results. Note the shift in the theoretical predictions due to changes in the solar model from Figure 1-8. Plot is from [22].

1.4.2 SNO NaCl Results

Once the initial pure-D₂O phase was finished, SNO dissolved 2 tons of NaCl salt in the heavy water. This increased the neutron capture efficiency as well as the associated Cherenkov light, thereby enhancing SNO's sensitivity to neutral current interactions. The data presented^{21,22} from the second, or 'Salt' phase of SNO was taken between July 26, 2001 and August 28th, 2003, with a total of 391.4 live days.

This more accurate measurement of the neutral current flux allowed SNO to give equivalent ⁸B neutrino fluxes of $\phi_{\text{NC}} = 4.94^{+0.21}_{-0.21}(\text{stat.})^{+0.38}_{-0.34}(\text{sys.}) \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}$, consistent with first-phase results. The allowed regions of mass-squared difference and mixing between two mass eigenstates are shown in Figure 1-10 for SNO data alone.




Figure 1-10: Allowed mass-squared vs mixing angle after SNO salt phase

The two ‘islands’ of allowed mixing angle and mass-squared difference space have significant differences in their mass-squared properties. Further data and analysis²³ reduced the lower island to a nearly negligible probability, as shown in .




Figure 1-11: Final allowed regions for neutrino mixing prior to KamLAND results. Plot is from [23].

KamLAND, a reactor-based anti-neutrino oscillation experiment, was set at just the right mass scale to confirm the solar results and completely eliminate the lower mass-squared difference region.

1.5 Current Status of Neutrino Physics

1.5.1 KamLAND

The Kamioka Liquid Scintillator Anti-Neutrino Detector looks for the signature oscillation of anti-electron neutrinos emitted from distant power reactors. The ~180 kilometer separation between reactor and detector allows them to test specifically the range of mass squared differences allowed by the Large Mixing Angle (LMA) solution to solar neutrino detector results. An anti-electron neutrino within KamLAND is detectable through the inverse beta decay reaction $\bar{\nu}_e + p \longrightarrow e^+ + n$. The signature of such interactions is the positron-electron annihilation in coincidence with a 2.2 MeV gamma from neutron capture on a proton in the liquid scintillator. The observed positron event rate within KamLAND²⁴ was lower than the rate expected in the absence of neutrino oscillations. A graph of all of the major long-baseline neutrino disappearance experiments is shown as Figure 1-12.




Figure 1-12: Neutrino flux deficit due to oscillations as observed at KamLAND. The gray region was predicted by solar neutrino results. Figure is from [24].

The deficit of electron-flavored antineutrinos at KamLAND not only reaffirms the solar neutrino detector results, it also places further limits²⁵ on the neutrino mass-squared difference, as shown in Figure 1-13.




Figure 1-13: Two-neutrino oscillation parameter allowed region from KamLAND and solar neutrino experiments. Figure is from [25].

1.5.2 SNO NCD Phase

The neutrino experiments up until 2006 have reduced the allowed neutrino mass-squared and mixing angle space in a 2-mass model to what is shown in Figure 1-14²⁶.




Figure 1-14: Effects of narrowing CC/NC and day-night asymmetry at SNO. Figure is from [26]

A day-night asymmetry is a specific prediction of the matter enhancement MSW model, and the best fit mass-squared difference predicts an asymmetry factor

$A = 2 \frac{N - D}{N + D}$ of about two to four percent. A more accurate measurement of neutral current rate can give a much better measurement of the mixing angle between flavor states. The ratio of charged current to neutral current, $\frac{\Phi_{CC}}{\Phi_{NC}} = 0.34^{+0.08}_{-0.06}$, is an approximate measure of $\sin^2(\theta_{12})\cos^4(\theta_{13})$.

Solar neutrino and KamLAND data are often analyzed in the framework of a 2-mass model because the mixing of electron flavor into the 3rd mass eigenstates is known to be small, at the 1 σ level (see section 1.6 for U_{e3}):

$$|U_{e3}|^2 = \sin^2(\theta_{13}) < 0.032$$

The current limitations on 3-neutrino mixing are presented graphically as Figure 1-15.

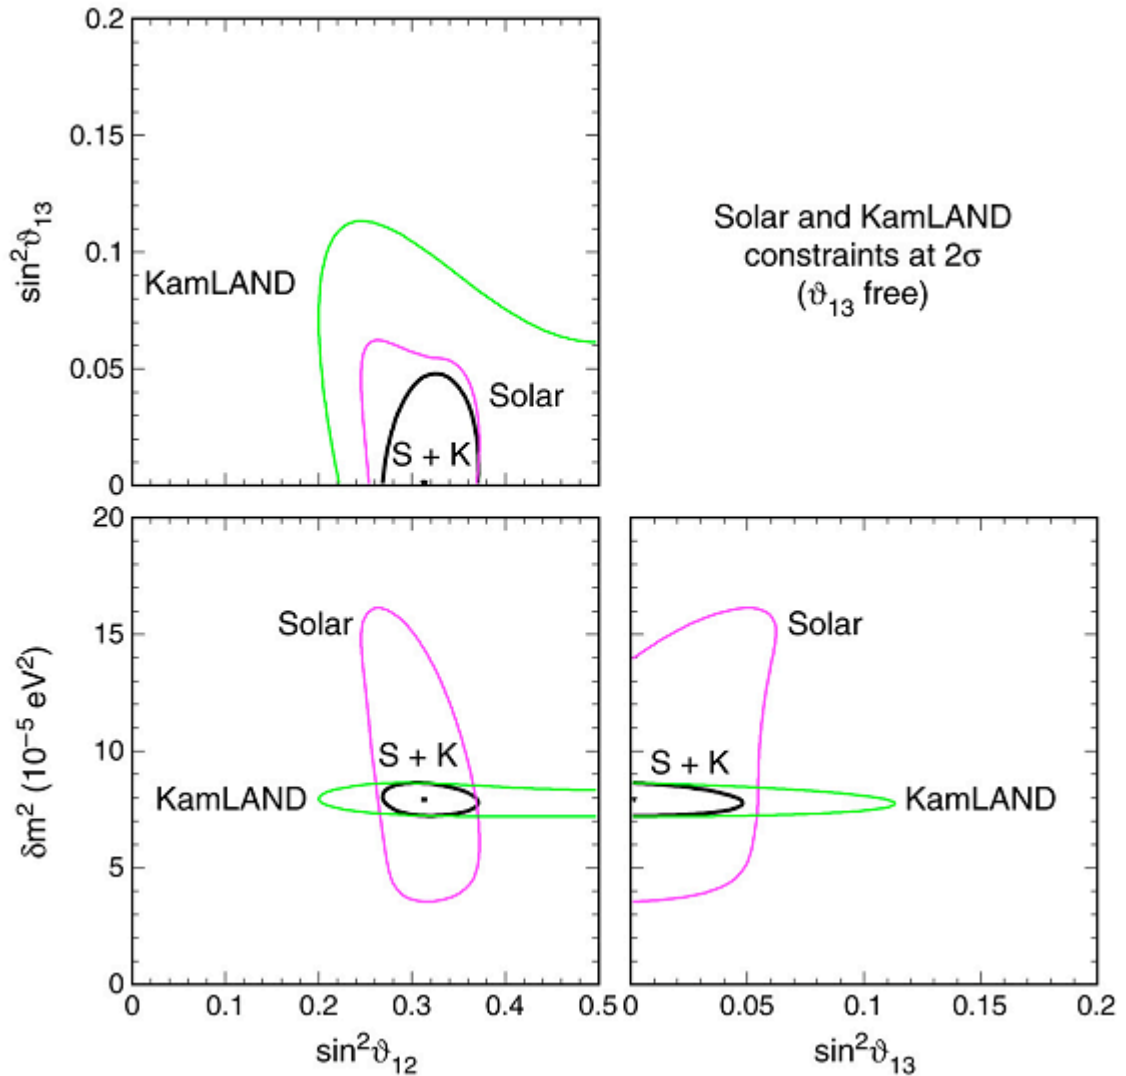


Figure 1-15: Current allowed parameter space for 3-neutrino mixing. Figure is from [26].

At SNO, the uncertainty in the NC/CC ratio has been limited by the neutral current flux measurement. The limits on the neutral current measurement, in turn, have their largest contributions from the energy scale and event reconstruction uncertainties. With the neutral current detectors in place, event-by-event identification is possible, independent of the PMTs. This completely removes the two largest sources of systematic error in the measurement of the NC/CC ratio, in addition to breaking SNO's NC/CC

correlations, and should reduce the neutrino mixing-angle θ_{12} uncertainty to ~ 1.4 degrees²⁷, approximately half of the current value.

1.6 3-neutrino family mixing

Since there are 3 active ‘flavors’ of neutrinos whose mass eigenstates are not identical to their weak eigenstates, the relation between the two bases can be represented by a 3 by 3 matrix. In this structure, neutrino fields of definite flavor ν_α (where α is e , μ , or τ) are linear superpositions of neutrino fields with definite mass ν_i (where i is 1, 2, or 3). The flavor and mass states may be related by a complex unitary matrix $U_{\alpha i}$:

$$\nu_\alpha = \sum_{i=1}^3 U_{\alpha i} \nu_i.$$

Single-particle neutrino states are then related by the complex conjugate, U^* :

$$|\nu_\alpha\rangle = \sum_{i=1}^3 U_{\alpha i}^* |\nu_i\rangle.$$

The most common parameterization of this transformation²⁸ is successive rotations through Euler angles θ_{12} , θ_{23} , and θ_{13} , with a CP-violating phase δ embedded between rotations:

$$U = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\theta_{23} & \sin\theta_{23} \\ 0 & -\sin\theta_{23} & \cos\theta_{23} \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & e^{i\delta} \end{bmatrix} \begin{bmatrix} \cos\theta_{13} & 0 & \sin\theta_{13} \\ 0 & 1 & 0 \\ -\sin\theta_{13} & 0 & \cos\theta_{13} \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & e^{-i\delta} \end{bmatrix} \begin{bmatrix} \cos\theta_{12} & \sin\theta_{12} & 0 \\ -\sin\theta_{12} & \cos\theta_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

This unitary neutrino mixing matrix works out to:

$$\begin{bmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{bmatrix} = \begin{bmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}c_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{bmatrix} \begin{bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{bmatrix},$$

where $s_{ij} = \sin \theta_{ij}$ and $c_{ij} = \cos \theta_{ij}$ for $i, j=1, 2, 3$. In total, this last mixing matrix is a $SO(3)$ rotation transforming a neutrino of mass 1, 2 or 3 into a neutrino of flavor e , μ , or τ . The transform is chosen to match the CKM matrix²⁹ in its modern form³⁰ in which the angles correspond directly to the mixing of two specific generations. The form of this matrix also allows all three real mixing angles to exist in the first quadrant.

This formulation assigns the mass-squared difference measured by solar neutrino oscillations to the mass difference between the first two neutrino mass eigenstates: $\Delta m_{21}^2 = m_2^2 - m_1^2 \approx 8.1 \times 10^{-5} \text{ eV}^2$. Because of the effects of matter enhancement, we know that the mass of the second neutrino is relatively close, and just above the first. The atmospheric neutrino mass-squared difference is a bit more ambiguous, as the oscillation is independent of the sign of the difference, and is measured to be $\Delta m_{23}^2 = \pm |m_3^2 - m_2^2| \approx \pm 3 \times 10^{-3} \text{ eV}^2$. As the sum of the difference between the squares of all three mass states must disappear, we expect the first and third neutrino mass squared difference to be $\Delta m_{13}^2 = \pm |m_1^2 - m_3^2| \approx \pm 3 \times 10^{-3} \text{ eV}^2$. We do not know if the first two mass generations of neutrinos lie far above or below the third. Following the charged lepton mass hierarchy, we might expect the first and second neutrino masses to be below that of the third, but we must allow for an ‘inverted’ hierarchy. The two possibilities, ‘normal’

and ‘inverted’ hierarchies, are shown in Figure 1-16, with the conventional generation assignments.

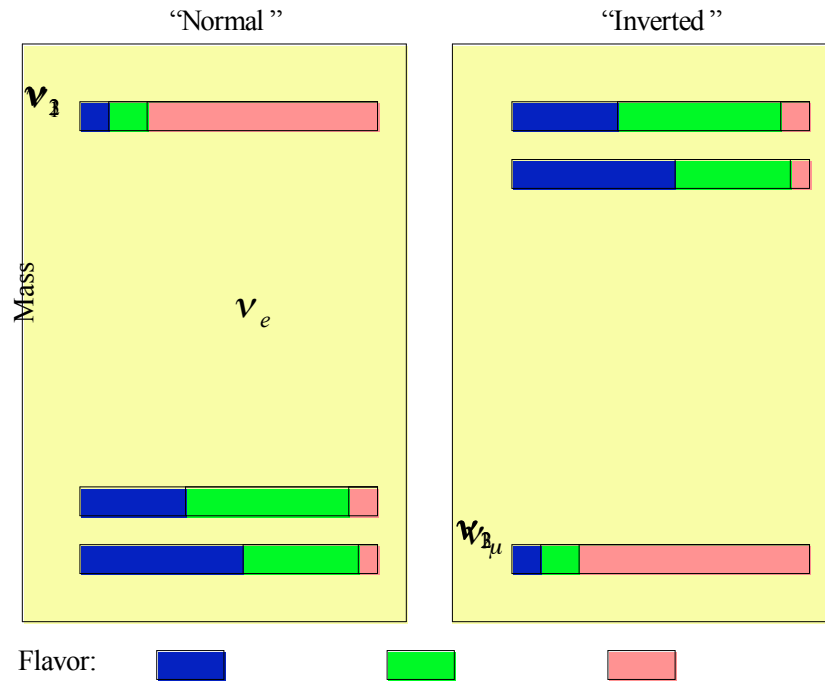



Figure 1-16: Neutrino flavor and mass spectra for the two possible mass hierarchies.

1.7 SNO NCDs

The Sudbury Neutrino Observatory Neutral Current Detection (NCD) array is made up of hundreds of ^3He proportional counters, assembled into 40 strings. The NCDs are designed to detect free neutrons by capture onto the ^3He :



Due to their relatively high energy, the proton and triton leave their electrons behind and convert their kinetic energy into electron-ion pairs as they travel through the

counter gas. The two distinct masses carry identifiable fractions of the total energy (190 keV for the triton, 574 keV for the proton), and deliver unique ionization-current traces if the tracks are perpendicular enough to the center conductor. ^3He proportional counters are a well established and well understood piece of scientific equipment that can be commercially acquired. However, no proportional counters to date have had to conform to the mechanical, electrical and purity requirements placed on them by the SNO detector and its expected signal. The purity and mechanical requirements are discussed here; signal and electrical constraints are discussed in Chapter 2.

1.7.1 NCD purity

The solar neutrino flux results in approximately 5000 neutrons per year inside the D_2O , and the NCD array was designed to contribute less than 1% of the expected signal. As a result of this design, the neutral current detectors were created with ultra-pure materials. The bodies themselves were made from the chemical vapor deposition (CVD) of nickel onto aluminum mandrels, for reasons of purity and thickness uniformity. Although the detectors begin their lives as purified nickel, they are activated by cosmic ray bombardment from surface exposure.

The nickel that makes the body of the tubes is of natural isotopic distribution, that is, roughly 68% ^{58}Ni , 26% ^{60}Ni , 3% ^{63}Ni , 1% ^{62}Ni and 1% ^{64}Ni . No sooner is each ultra-pure detector body formed onto a mandrel than cosmic ray neutrons begin to create radioactive isotopes of cobalt and other elements through spallation. Most of the cosmogenically activated isotopes do not contribute significantly to the

photodisintegration background, as either their half-lives are too short, or their gammas too low in energy to liberate neutrons from the deuteron. The one notable exception is ^{56}Co , an isotope that has a half-life of 77 days. Just over 31% of all ^{56}Co decays result in a gamma with an energy over the 2.224 MeV required to liberate a deuteron-bound neutron.

The only channel for ^{56}Co to enter the system is through cosmogenic neutron activation, and the only exit channel for ^{56}Co to leave the detectors is through decay. Assuming no shielding from bombarding particles, the isotopic decay rate of an originally pure material bombarded at $t=0$, per unit mass, is

$$R(t_E) = R_f \left(1 - e^{-t_E/t_\lambda} \right)$$

where R is the isotopic decay rate per kilogram per day, R_f is the final/surface equilibrium decay rate, (calculated to be about 24 decays/day/kg for ^{56}Co in Ni), t_λ is the half-life of the isotope divided by $\ln(2)$, and t_E is the time of exposure. Naturally, if the tubes are shielded from cosmogenic neutron spallation, then there will be no regeneration of ^{56}Co . The expected rate of ^{56}Co , in decays/day/kg after a cooldown time of t_c is:

$$R(t_c) = R(t_E) e^{-t_c/t_\lambda}$$

COSMO³¹, code specifically designed to estimate spallation radioactivity produced in a pure substance by exposure to cosmic radiation on the Earth, predicts a steady-state ^{56}Co decay rate in natural Ni of 24.22 decays/day/kg. If 400 kilograms of the array were taken from surface-saturation and immediately placed in the D₂O volume, we

would expect 9680 ^{56}Co decays per day. Monte-Carlo simulations through SNOMAN have indicated that .0877% of ^{56}Co decays result in neutron liberation, so that a fully activated array would create just over 8 neutrons per day. In an effort to reduce cosmic ray spallation activation, it was decided that the tubes should be kept protected from cosmic rays by placing them at an underground storage facility until they were ready to be fabricated into counters.

Unfortunately, the tunnel adit at Index, WA, where the tubes were stored had extremely high radon levels (900 pCi/l), whose daughter ^{210}Po plated out onto the nickel counter bodies. All tubes were moved back to UW after the radon was discovered, but not before they had all been contaminated and showed ^{210}Po -alpha decay rates of around 8000 alphas/m²/day. This contamination led to electropolishing, which is described later and had the effect of reducing the Po-alpha activity below 100 alphas/m²/day.

1.7.2 NCD Cooldown

The primary purpose of the cool-down phase was to allow the majority of the cosmogenically-generated ^{56}Co in the nickel of the completed counters to decay, and to prevent further cosmogenic activation. During this cool-down phase, there was ample time to bring the electronics and data acquisition systems online. By the time NCDs were installed, the entire array had ‘cooled down’ for a period between 2 and 6 years. All but 11 of the detectors had been underground for more two years. This means, that aside from the last 11 detectors, the array had less than 0.2% of their steady-state surface levels of ^{56}Co at the time of deployment. Computing with the formula above, we expect to see

less than 1% of the total saturated ^{56}Co decay rate from the entire array. Considering that only 300 meters of detector are in SNO, we expect an initial yield of about 1 neutron in the D_2O volume per 50 days due to ^{56}Co .

During cooldown, the NCD electronics were also tested, verified and prepared for full neutrino runs. The commissioning of the electronics and data acquisition systems is detailed in Chapter 3.

Incidental to the cooldown phase was the monitoring of neutrons emanating from the mine walls that surrounded the cooldown area. The local neutrons provided a constant signal for the NCD array to test its ability on. The extended runs also provided a good understanding of the thermal neutron flux. When the flux is combined with the known neutron generation rate in norite, we were able to accurately verify the amount of ^{238}U and ^{232}Th in the mine walls. Constantly monitoring the neutron flux also allows for detection of time-correlated neutron measurements. A burst of neutrons can be liberated from a target mass by sufficient high-energy neutrino flux, such as the neutrino burst preceding a galactic supernova.

1.8 Supernovae and neutrinos

The closest type-II (core-collapse) supernova in recent history, 1987a, occurred in the Large Magellanic Cloud and was visible to the naked eye. More impressively, it was visible to three neutrino detectors that were online at the time. Irvine-Michigan-Brookhaven³², Kamiokande-II³³, and Baksan³⁴ detected a short burst of neutrinos hours before 1987a became visible.

The energy signature of the neutrinos detected gave information about the thermal energies within the collapsing star on the surface of the “neutrinosphere”. The neutrinosphere is the surface from which a created neutrino is no longer expected to interact on its outward path. In this fashion, a collapsing star is similar to the early universe; we only see the surface of last scattering.

A type-II supernova begins with a progenitor star of substantial mass ($M > 8M_{\odot}$). The star is balanced throughout its life by the pulling of gravity against the pushing of radiation pressure. As the star shrinks from self-gravity, the gravitational energy liberated heats the star, increasing the core temperature and thus the thermodynamic pressure. The increased temperatures and pressures liberate energy in the core by the fusing of light elements, which then feeds back against the pull of gravity, greatly prolonging the life of the star. Fusion begins with lighter elements such as hydrogen and helium, which leave behind the products of their fusion in the core. Eventually these products overwhelm the core, and slight increases in pressure and temperature are no longer able to adequately increase the core’s rate of fusion. The star shrinks until the point at which the core is hot and pressurized enough to begin to fuse the products, be it helium, carbon, oxygen, or silicon. At this point, the fusion of these once-products within the core further increase the temperature and pressure, causing the star to expand until a balance is once again achieved. These steps continue until the silicon in the core is converted into iron, the element with the lowest energy per nucleon. Even though the core has thoroughly used its hydrogen and helium, the majority of the mass of the star is still in these lightest elements that surround the aging star in successive shells.

Once the iron core has reached a mass of $M_c \sim 1.26$ solar masses, the large star can no longer heat its core by fusion, and the gravitational contraction part of the cycle continues unopposed. As the star contracts further, the dense core's gravitational pull overwhelms the resistance put up by thermodynamic pressure, and the star begins to collapse. The electron degeneracy of the stellar core quickly expands, and energetic electrons turn many of the infalling protons into neutrons, driving out electron neutrinos in the process. The gravitational energy liberated by the infall of the iron is snapped up in photodisintegration of that same iron through the endothermic reaction $\gamma + {}^{56}\text{Fe} \rightarrow 13\alpha + 4n$, preventing the core from maintaining hydrodynamic balance. The neutron density of the core then quickly increases, as the core expands its neutron degeneracy. Additionally, nuclei capture electrons within the core through the general reaction $e^- + (Z, A) \rightarrow (Z - 1, A) + \nu_e$, further decreasing the electron degeneracy pressure and releasing neutrinos from the core. These effects combine to allow stellar matter outside the core, as well as some within the core, to collapse with near-free-fall velocities. When the core material's inward kinetic energy finally drops below the massive chemical potential of the neutron degenerate Fermi gas, the core rebounds explosively, sending out a shock wave with particle velocities of tens of thousands of kilometers per second.

Originally, this shock wave was thought to be strong enough to cause the explosion of the star. However, models³⁵ have shown that this rebounding shock front 'stalls' after a few tenths of seconds, falling prey to the still-collapsing outer shells. The pressure waves are unable to progress further than the boundary of this shock front, as the

infall velocity of the outer shells is equivalent to the propagation velocity of pressure waves (speed of sound). At this point, the energy from the pressure waves intensifies and heats the collapsing shells where it has stalled, quickly putting enough energy into the core to dissociate everything at the cost of 8 MeV/nucleon, in a desperate attempt to spread degrees of freedom. The core then continues to heat up as dissociated matter plummets inwards, driving neutrino temperatures up as the shock wave continues to struggle against the incoming tide of outer stellar matter.

Visual supernovae must, at this point, drive the shock front outwards, reviving the explosion, albeit with some delay. However, supernova models are not just inconsistent at how the delayed explosion is reignited, but also uncertain if the supernova becomes visually reignited at all. By observation, we are certain that at least some fraction of supernovae become visible, and we will exclusively deal with those for the remainder of this paper.

Some models predict that high energy neutrinos drive convection just outside the core, mixing enough light elements with hotter, heavier ones to re-ignite the explosion. While a great improvement over previous simulations, these convective models have been shown to be incomplete, in that approximations made in the microphysics and neutrino production areas have had significant effects on increasing the probability of re-ignition. Additionally, the computer resources available for simulations have precluded all attempts to simulate stellar collapse in more than 2 dimensions. Since convection is qualitatively different in 3 dimensions than in 2, this further obfuscates this potential solution.

Despite controversies about the details, there is still a general consensus on certain key neutrino signatures of supernovae. During the neutronization stage of a supernova, the vast majority of neutrinos emitted will be anti-electron flavor. This emission will last between 200 and 600 ms, and will have a capture spectrum that hardens until possible shock breakout. The luminosity will also increase quickly as the core contracts, but the exact profile of this increase is disputed. The entire amount of neutrino energy radiated from the infall of the star will be $\sim 10^{51}$ ergs.

The shock wave forms at around 20 km out from the core, where the first of the pressure waves coalesce, almost immediately after core bounce. The neutrinos from the initial shock wave are not visible, since they lie well within the neutrino-opaque 'neutrinosphere'. However, the initial shock wave velocity of $5 \cdot 10^7$ m/sec quickly carries it out through the neutrinosphere ($5 \cdot 10^4$ - 10^5 km), where the density is low enough that trapped neutrinos can escape. Here, the luminosity in many models is predicted to rise above 10^{54} ergs/second, but only lasts for a brief 3-10 ms. This means the energy released from the shock breakout is on the order of 10^{51} ergs. This breakout is directly dependent on progenitor mass and the density profile of the collapsing Chandrasekhar core. The matter within the shock front has been heated to such a degree that neutrinos will be thermally produced and radiated. Burrows' models claim that this leads to a very abrupt turn-on of the non-electron neutrinos, but others contend that this merely produces a more subtle increase in their luminosity.

The time between when the shockwave stalls into an accretion shock at 100-200 kilometers and the explosion is supposed to reignite, is when most of the neutrinos are

emitted from the dying star. This period of time is at least a few hundred milliseconds long, but could be tens of seconds. During this time, more than 10^{53} ergs of neutrinos are emitted, over 100 times the energy released in the comparatively minor visual burst that shines brighter than galaxies.

In theory, the energy spectra of the neutrinos harden at first, due to increased neutrino mean interaction distance, but the total luminosity drops steadily as the protoneutron star cools. If explosion occurs, the neutrino fluxes decrease by about a factor of two in less than 40 milliseconds, thereafter they decrease on the order of seconds as the embers of the star fade. The distribution of neutrino energies escaping at any time follow a Fermi-Dirac temperature profile, such that the density of states n_r can be calculated as a function of temperature as follows:

$$n_r = \frac{g_r}{e^{\alpha + \frac{\epsilon_r}{kt}} + 1}$$

where g_r is the degeneracy of the level r , $\alpha \equiv -\frac{\mu}{kt}$, μ is the chemical potential, k is Boltzmann's constant, t is the temperature, and ϵ_r is the energy of the level r . Plotting out the distribution for a temperature of 8 MeV yields the solid line of Figure 1-17.




Figure 1-17: Fermi Dirac Distribution of escaping thermal neutrinos

Massive neutrinos could drastically alter the signal from supernovae by oscillating. Since electron neutrinos have a much shorter interaction length, they are emitted only from the outermost surface. This lends them the cooler temperature profile based on that of the outer neutrinosphere. Mu and tau neutrinos, on the other hand, do not have their weak isospin doublets in any appreciable quantities within the neutrinosphere, and thus their interaction length is increased. The temperatures simply are not high enough to freely create muon and tauon pairs. All three active flavors will interact with Z-boson exchange, that effectively set the limit on last scattering radius.

Sterile neutrinos, if they exist, do not interact, even with the Z-boson, and would be able to carry off the thermal energy from the collapsing star. There would be no theoretical neutrinosphere within which the sterile neutrinos would appear opaque. However, sterile neutrinos could only be produced if they have some ($\sim 10^{-4}$) admixture of active flavors. Even then, the star would have to provide the conditions for the active flavors to oscillate into sterile families. Unless the conditions were set up favorably, we would expect matter-enhanced oscillations to remove the sterile neutrino flux as it was being created.

Depending on the particular model being used, neutrino oscillation, or even flavor changing neutral currents, should equilibrate the energy spectra of the different families of neutrinos. For neutrino oscillation, the exact amount of mixing is determined partly by vacuum oscillations, partly by matter-enhanced resonances, and partly by neutrino-neutrino interactions, a highly non-linear mixing term.

Neutrino emission from supernovae has been theorized for many years, but it was not until 1987a that the theories were given observational validation. In 1987, a 20 M_{\odot} blue supergiant star in the Large Magellanic Cloud, over 50 kiloparsecs away, collapsed and exploded, bringing new light to the night sky, but not before leaving a distinct signature in the active neutrino detectors, Kamiokande II³⁶, IMB³⁷, and Baksan³⁸. Both were showered with the sudden stream of higher-energy neutrinos and had a combined signal³⁹ as shown in Figure 1-18:

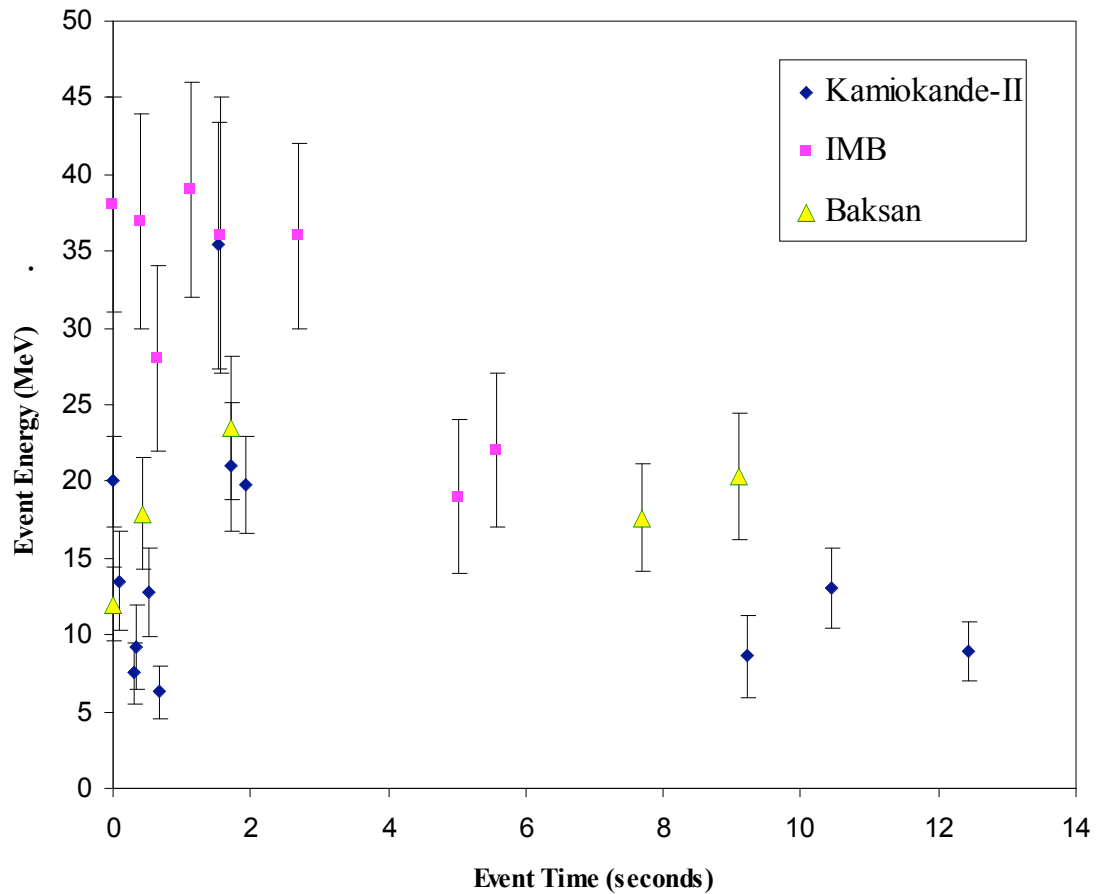


Figure 1-18: Potential 1987a neutrino arrival times and energies. The absolute time of each data set is determined by the first event in the series. Data is compiled from [39].

The arrival times of the events gave the time structure of the core collapse, but more importantly, the distribution of the arrival times of different neutrino types gave a direct measurement of neutrino mass. Dozens of analyses were published, too many to cite, with definite and maximum mass estimates ranging from 50 meV to 26 eV. Figure 1-18 displays the limited data sets from Kamiokande, IMB and Baksan. It is important to note that the relative timing of these two sets of events are only known to the nearest minute; the overlay shown is only one possible event structure.

The frequency of galactic supernovae is also not well known, due to poor statistics and the inability to see much of our galaxy because of obscuring dust. This, combined with the fact that not all supernovae are necessarily visual, leave the frequency of galactic supernovae somewhere between 1 every 10 to 100 years. Additionally, the time structure of collapse is still a point of some contention, since experimental data is not particularly restrictive.

The measurement of the energies of the neutrinos from supernovae will provide a basic understanding of the core temperature profile during collapse, since the different neutrino species experience different effective stellar opacities. However, as mentioned, neutrino oscillation greatly complicates the interpretation. Finally, monitoring the neutrinos from supernovae might provide insight into the theorized R-process creation of heavy elements. Neutrinos from supernovae are theorized to enable the large neutron flux required for the R-process to synthesize the heaviest elements.

1.8.1 NCDs and supernovae

The large flux of neutrinos from a galactic supernova incident on the heavy water inside SNO's acrylic vessel will create a distinctive signal⁴⁰. This work does not investigate the details of that signal, but rather examines the requirement placed on the electronics by the expected signal. However, this work presents a cursory investigation of the expected supernova signal when the NCD array is not in the heavy water—when the NCDs are sensitive to neutrons from the mine wall. A more detailed analysis of the neutron flux and its detection within the SNO control room is given in Chapter 4.

Although the probability of a detectable supernova event during the relatively short-duration cool-down phase was low, this test provided valuable data as to the feasibility of using proportional counters for supernova detection. These data should be used to help assess proposals such as OMNIS⁴¹, LAND⁴², and other⁴³ long-term neutron-based supernova searches. We test supernova event backgrounds, counter geometric efficiencies, neutrino target materials, data acquisition systems, and supernova alert protocol. The results of these tests could then be used to optimize supernova detection strategies for permanent supernova detectors.

Chapter 2 NCD Array and Electronics

The infinite we shall do right away. The finite may take a little longer. S. Ulam

The NCD electronics are meant to simultaneously control and monitor up to 96 detectors in the NCD array. While inside SNO's acrylic vessel, the NCDs need to differentiate between neutrons and other (physics and non-physics) events in order to distinguish potential neutral-current neutrino flux. By design, ^3He proportional counters provide a high neutron cross section as well as the ability to decisively extract a neutron flux.

Reliably distinguishing neutron events is tied up in the shape of the proportional counter's ionization current pulse. The shape of the current pulse carries information that can be used to distinguish neutron events from other types of events, so long as it is digitized with sufficient time resolution. While the characteristic times for neutron identification are limited by lateral straggling and diffusion to 50 ns or longer, extracting position information calls for better time resolution. Fortunately, the rate of neutron detections within the NCDs is small—the expected solar neutrino spectrum²¹, when combined with the neutron liberation cross-section, gives but a few thousand events a

year. However, neutrinos from local supernovae can produce hundreds of events within seconds, presenting us with two completely separate rates of physics events to be quantified.

The NCD electronics were designed with both of these scenarios in mind. Dual data paths provide the low-rate scenario with digitized data while allowing for high-rate data taking. The electronics for each string of counters has a pair of independent thresholds, one for low rate pulse transient digitization and one for high rate signal integration. The first discriminate a current signal, the second an integrated charge. The NCD electronics also provide 8 individually controllable high voltage power supplies that can be manually connected to groups of strings. An overview of the NCD electronics is shown in Figure 2-1.

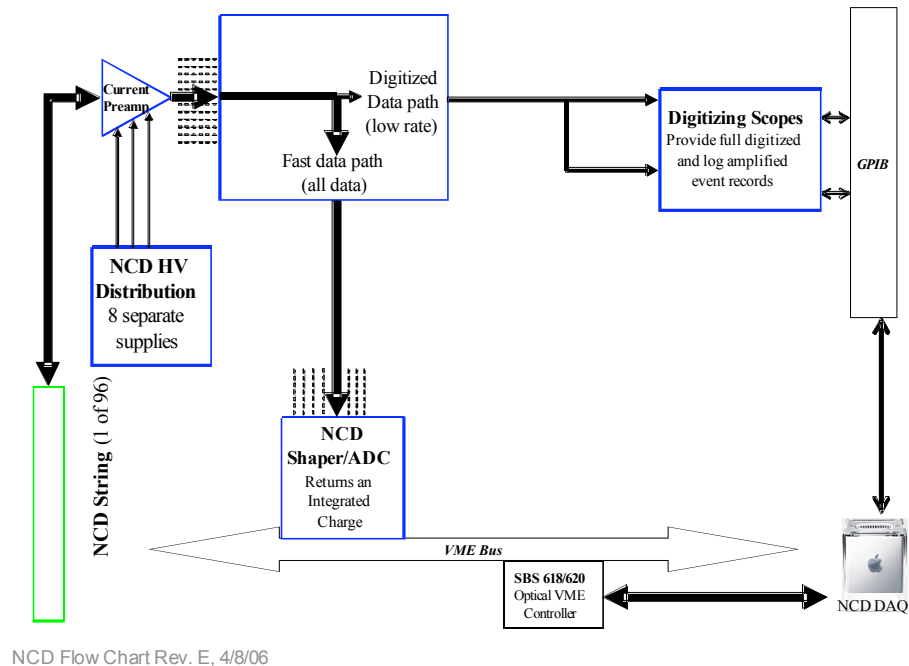


Figure 2-1: Basic NCD Electronics Flowchart

2.1 NCDs as proportional counters

The NCDs main purpose is to catch and identify as many thermal neutrons within SNO as possible. Backgrounds must be kept to a minimum, so contaminants introduced by NCDs are kept in check by minimizing material and assuring cleanliness of each component. The NCDs are unfortunately able to detect more than just neutrons, as Compton-scattered electrons as well as γ 's can ionize and create a cascade directly. Alpha particles, although unable to pass the thickness of the counter wall, can be seen if the contaminant is either on or just within the inner surface of the detectors.

The gas mixture within the NCDs is 85% ^3He and 15% CF_4 by volume, at close to 1928 Torr pressure. The ^3He content provides the counters with high neutron cross-section, and CF_4 provides increased stopping power for ionizing particles as well as improving the performance of ^3He as a counter gas. The outer NCD body is made through chemical vapor deposition, a process which delivers ultra-pure nickel, in an extremely accurate distribution. The ^{238}U and ^{232}Th contaminants are a few pg/g. In order to limit contaminants, the walls of the NCDs are just over 380 microns thick. This thickness is sufficient to resist collapse under vacuum at sea level.

Running through the center of the outer nickel body is the inner anode, a 50- μm -diameter copper wire anode maintained at 1835 V (for a gas gain of 100), used during some of the cool-down phase, or 1944 (for a gain of 200), used for the in-situ phase. This thin center conductor is strung with .3 N of tension.




Figure 2-2: NCD Cross-section (not to scale)

At either end of each counter, a 5 cm long, 5 mm diameter synthetic fused silica “field tube” insulates the nickel body from the copper anode, as well as providing an insensitive region at the end of the counter. The inner wall of the tube feed-through was given a conductive layer of pyrolytic graphite, cracked from methane gas. The layer is electrically connected to the anode. This eliminates the avalanche gain that would otherwise exist in the region inside the detector around the feed through. The 25 millimeter shielded length of anode at each end of a counter, prevents alpha decays from the end caps entering the gas and mimicking neutron signals due to the complicated electric fields and slow charge collection there.

Each assembled NCD counter is a single 2-, 2.5- or 3- meter section of a “string”. In their underground rack configuration as well as in their deployed form, the NCD strings vary in length between 8 and 11 m. While some counters could be welded

together before deployment, the assemblies could not exceed 5.6 m in length due to height constraints above the detector. The remaining welds had to be carried out above the neck of the vessel during deployment⁴⁴. The total length of all underground NCD counters is around 770 meters, but only 398 meters of NCD counter sit within the D₂O volume in 40 strings. Table 2-1 summarizes the deployed NCDs by counter and Table 2-2 summarizes the deployed NCDs by string.

Table 2-1: Deployed NCD Counter Summary

Type	Deployed Counters			Total Length
	2 m	2.5 m	3 m	
⁴ He	12	0	4	36 m
³ He	36	44	60	362 m
Sum:	48	44	64	398 m

Table 2-2: Deployed NCD String Summary

Fill	Deployed Strings			String Length	String Designations
	2 m	2.5 m	3 m		
⁴ He	3	0	1	9 m	I2,I3,I6,I7
³ He	3	0	1	9 m	I1,I4,I5,I8
³ He	2	1	1	9.5 m	J1-J8
³ He	2	0	2	10 m	K1-K8
³ He	0	3	1	10.5 m	L1-L4
³ He	0	3	1	10.5 m	M1-M8
³ He	1	0	3	11 m	N1-N4
Sum:	48	44	64	398 m	

The NCDs were deployed in a 1-meter square grid. Strings placed equidistant from the center of the array are given identical lengths and letter designations. The 40

NCD strings placed in the D₂O, their string names, numbers, and electronic channel identifiers are shown in Figure 2-3.

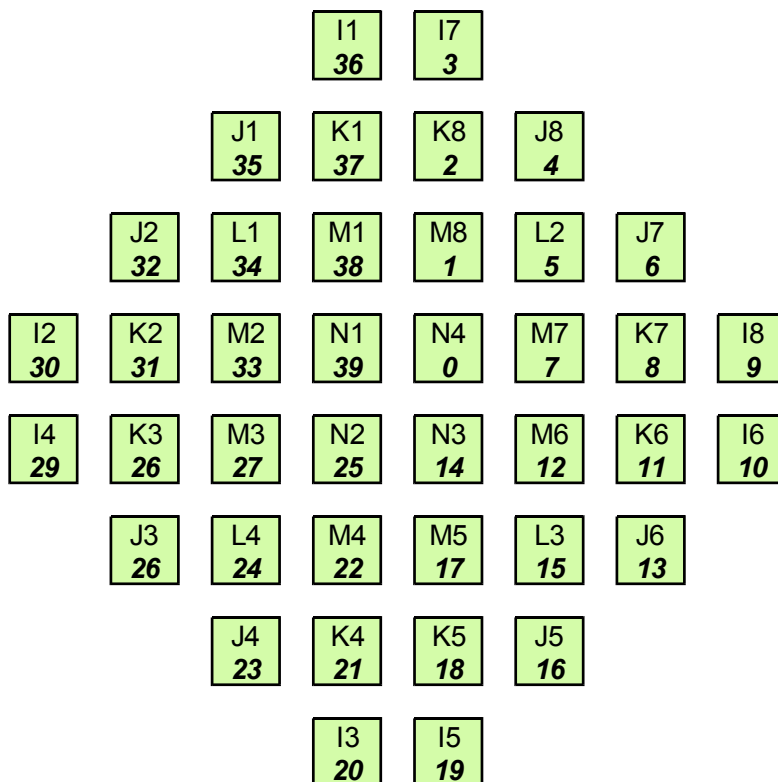


Figure 2-3: Basic Deployed NCD Configuration

Individual NCD counters are linked together by an intercounter coupler, a CVD nickel strip in a helical form. The expected separation between counter anode contacts is 7 mm, but the connectors must be connected before the counter bodies are welded together. The helical form of the intercounter couplers allows for variations of over twice the expected separation. The intercounter couplers attach at the outside nickel end of each counter's feed-through with a 5-mm diameter nickel clamp. A schematic of the intercounter connection is shown in Figure 2-4.




Figure 2-4: The Intercounter Coupler. Drawing is from [45].

Each string is connected to the electronics through a ‘cable bell-end’ at the top of each string. The ‘bell-end’ is a cap of CVD nickel which welds onto an NCD string and has NCD cable built into the other side. Electrically, the ‘bell-end’ connects the grounds

of the cable braid to the counter directly through its outer body, while the anode wire runs through the ‘bell-end’ and a feed-through to the resistive coupler. The bell-end connection with a NCD cable is shown in Figure 2-5.




Figure 2-5: NCD ‘Bell-End’ Cable-Counter Connection. Drawing is from [45].

The resistive coupler matches the 415- Ω impedance of the NCD string to the 93- Ω cable, minimizing unwanted reflections. The resistive coupler, shown as Figure 2-6, is a 1.06-inch diameter Teflon ring that bridges the 1-inch electrical gap by means of an 11 inch, 0.001" diameter "Stablohm" resistance wire which loops back and fourth in a grooved path placed in the Teflon. The path does not encircle the ring; rather, it makes successive half-circles to reduce inductance. Therefore, all of the wire is on one half of the ring. At the ends, the wire conductor on the resistive coupler is attached to a spiked nickel collar which forces electrical contact when placed over nickel feed-throughs. The resistive couplers have a resistance of 325 Ω .




Figure 2-6: Flexible Nickel-Teflon Resistive Counter Coupler. Drawing is from [45].

2.1.1 Delay line

As the charge from particles cascading towards the center of the NCD arrives at the anode, it is free to redistribute its charge in both directions inside the counter. The counters behave as transmission lines with a characteristic impedance of 415Ω . Rather than terminate the far end of each string to prevent reflections, the anode of the bottom counter in every string runs into an unterminated delay line also of 415Ω nominal impedance. The anode of the top counter attaches to an NCD cable through the 325Ω

resistive coupler. This allows the upward traveling half of the signal to propagate cleanly into the preamplifier, followed in time by the downward-going half of the signal reflected from the open end.

The delay lines have round trip delays of 80-90 ns. At the end of the delay line, total reflection is guaranteed by its infinite impedance termination. The reflected signal has the same polarity as the original signal, albeit dispersed, with a delay to rise time of approximately 10. Given that current travels along the center anode at approximately 0.9 c, and that the deployed strings vary in length between 9 and 11 meters (section 2.2), the arrival times at the top of a string can vary between 1 ns for the direct pulse from an event at the top of a counter and 166 ns for the reflected pulse from an event at the top of an 11-meter string. The smallest difference in arrival times is from an event right next to the delay line, or bottom of the counter. This signal should take the same time to propagate up the anode, but the delayed reflection would only have an 85 ns delay. As a function of distance from the delay line (or distance from the bottom of the active counter), z , the delayed pulse will take an extra delta-t:

$$\Delta t = t_d + \frac{2z}{v_c},$$

Where t_d is the round-trip delay of the delay line, v_c is the propagation speed in the NCD string. Since the reflection arrives not more than 170 ns after the direct pulse leading edge, extracting position information depends on pulse amplitude, pulse width, and signal-to-noise ratio. A reliable Z-position can be found for approximately half of the neutron events, with a bias for accurately finding events having tracks parallel to the

anode, because they make the shortest pulses. An improved method⁴⁶ is under investigation.

2.1.2 NCD Cable

The NCD cable is designed to remain fully immersed within the D₂O for years of deployment. Each is cut to the appropriate length to run the distance from its counter end to the top of SNO's neck, where it connects to the preamplifiers. In order to reduce the stress at the cable-counter connector as well as maintain the vertical alignment of NCD tubes, the cables are slightly buoyant within the D₂O, with a specific gravity of 0.955. As the specific density of D₂O is 1.106, each liter of cable should produce a net 1.51N of buoyant force. The cable is a specially fabricated 93□ coaxial cable with a 0.5 mm-diameter copper central conductor surrounded by a woven-ribbon copper ground shield. Additionally, the cable is designed with materials that minimize possible radioactive background contribution: Pure copper, low-density polyethylene dielectric, and high-density polyethylene jacket.

2.1.3 Chemical Vapor Deposition

The stringent cleanliness limits of SNO require that all NCD material be as free of contaminants as possible. Consequently, Chemical Vapor Deposition (CVD) has been used to manufacture not just the counter bodies, but also many of the miscellaneous parts. The CVD process is limited, clean, and perfectly fit for making the NCDs. Chemical Vapor Deposition can be used to condense many different elemental metals, but the

chemical properties and elemental strength of nickel provide the best fit for the NCD requirements.

The CVD process starts with the reaction under pressure of powdered nickel (Ni) with heated carbon monoxide (CO) in a mixing chamber where the gaseous carbonyl $\text{Ni}(\text{CO})_4$ is formed. A liquid at room temperature, the carbonyl can be revaporized and passed over a heated aluminum or stainless steel mandrel, where the carbonyl ‘cracks’ and the metal deposits slowly upon a surface. The CO is recycled and bound to more metal at the chamber.

Nickel binds with CO into nickel carbonyl at 100° C, and nickel carbonyl cracks at 200° C. Although some other elements, such as Cr, Fe, Co, W, and Mo are also carried through the CVD process, most of the U and Th are not. Although it was believed at the time of construction that nearly all of the U and Th would be removed through CVD, chemists later discovered⁴⁷ that Th forms a carbonyl under some conditions, which may account for the final determination of 5.2 pg/g Th in the NCDs (see section 5.3.1). The CVD processing was carried out by Mirotech and its successor, CVD systems in Toronto, Canada.

2.1.4 Electropolishing

The elevated levels of ^{222}Rn daughters, specifically ^{210}Po , that had concentrated on the surface of the NCDs was the main target of the electropolishing. Po displaces Ni into acidic solution, but can be kept in solution if encouraged by an electrical potential that was supplied with ample current (600 A). About 25 μm was etched from the inside

of the tubes to assure removal of any Al particles from the mandrel. A smaller amount was removed from the exterior to reduce the ^{17}O , ^{18}O (α, n) reactions in the heavy water. Since no mandrel aluminum was present on the outside of the NCDs, only a short run on the outside of the counter bodies was required.

Essentially, electropolishing is electroplating done in reverse. The object that is being polished is immersed in an electrolytic solution (1M H_2SO_4) and held at a positive voltage with respect to the Ni cathode that is placed along the tube axis. Neutral nickel atoms enter the solution as Ni^{++} and are reneutralized at the cathode. If the counters had been held at a lower potential relative to a metal-ion rich solution, then metal would have attached itself to the body of the counter, resulting in electroplating.

As the NCD counters were designed to minimize contamination, they were already thin. For structural reasons, we could not afford to lose much additional metal. Since the rate of metal lost to solution in electropolishing is a function of current density, solution, geometry and temperature, all factors had to be controlled throughout the process. If too little polonium is removed, the surface polonium contamination will drown out critical data. If too much nickel is lost in the process, the tube might collapse under the load of heavy water to which it will be exposed. Electropolishing removes material in the direction of field lines, so that conductive surfaces lose material in a unidirectional fashion, and it does so without stressing or occluding. The uniformity of electropolished surfaces allowed us to remove mass without affecting structural integrity.

In the UW NCD electropolishing rig shown in Figure 2-7, we ran two separate sections, one each for the inside and outside surfaces. In the first section, each NCD was electropolished twice. Each of the three successive electropolishes used acid of increasing cleanliness, culminating in a fresh batch of acid for the final internal rinse. The acid was removed from the system after it was used for the external electropolish of a tube. Since electropolishing nickel only becomes effective at temperatures above 65° C, the acid used for electropolishing had to be heated.

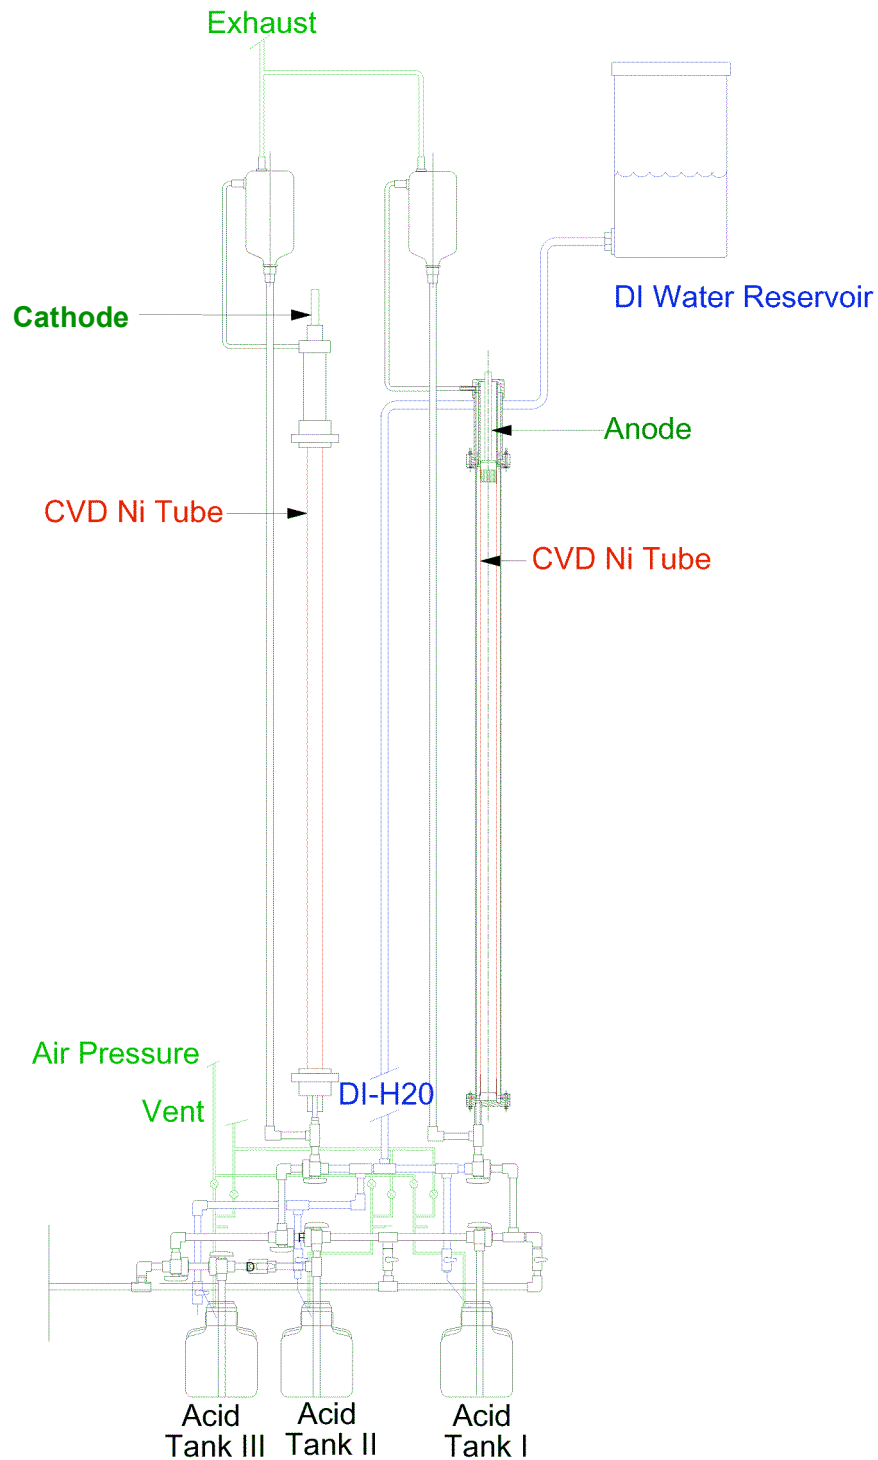


Figure 2-7: NCD Electropolishing Rig. The outside of the NCD tube is cleaned in the first electropolish using Tank I. The inside is then taken through two electropolishes, using acid Tank II followed by Tank III. The used acid is cycled to the right (to earlier electropolishes) so that the first acid (Tank I) has been used twice for previous NCD tubes, while the last acid used (Tank III) is new.

The temperature on the external electropolishing rig was increased by heating tape wrapped around the NCD containment tube. The anode voltage was monitored in comparison to the NCD containment tube, and the current recorded every minute during electropolishes. After electropolishing the outside of the NCD, the acid would be drained from the system and the NCD rinsed 3 times with purified DI water.

The internal electropolishing rig did not allow for heating tape, as the nickel tube served as the acid containment device. In order to remove a uniform thickness of nickel, the temperature of the acid was monitored by an infrared sensor as the polishing progressed. Temperature, in addition to current, was recorded for all internal electropolishes. After each electropolish, the used acid, noticeably more green, was moved further up the system and the tube was rinsed three times with purified DI water. The green color, due to Ni^{++} , was monitored optically to determine when the desired amount had been removed.

2.1.5 Gas Fill

The NCDs, like most proportional counters, rely on a thin conductive center wire held at high positive voltage relative to the outer walls. The counter is filled with a mixture of low-electronegativity gases whose admixture and pressure were carefully selected for highest capture efficiency while attempting to avoid space-charge saturation. Additionally, the pressure within the NCDs, in combination with the structural strength of the counter body, had to be sufficient to withstand the 17.65 meter column of D_2O on top of the elevated underground air pressure of 1000 Torr. In total, the NCDs have to be of

sufficient structural strength and internal pressure to sustain 2440 Torr at their base. In testing, the integrity of each NCD was determined under vacuum test. Further information about the counter gas and gas gain is provided in the section covering the high voltage requirements of the NCDs (2.10.1).

2.2 Electronics Specifications

The electronics must identify as many neutron events in the counters as possible, without losing events due to lack of time resolution or coincident events. Neutron event position resolution requires that each event be separable from its own reflection, arriving between 85 ns and 150 ns later. Such resolution is readily available on digital oscilloscopes. Digitizing scopes have dead-time between events as the data is read out, so coincident event overlap is possible. To avoid loss of interesting data in which two neutrons or two alphas are detected within a short interval, a pair of digitizers is used.

If the NCD array were only to measure neutrons born from solar neutrino interactions we would expect, with a 0.285 capture efficiency⁴⁸, to detect 4 neutrons per day. The chance of two neutrons arriving during the same second is 0.012% per day, or an average of 1 time every 17 years. With background alphas included, the dual event rate is still only a few events per month. Uncorrelated triple coincidences should be less frequent than the lifetime of the detector. Thus, a minimum of two digitizers with busy times on the order of one second per event are necessary to avoid loss of both Poisson-distributed signal data and time-correlated data such as arise from the uranium and thorium alpha decay chains.

In order for NCD and PMT data to be associated, the electronics for both systems must communicate. As SNO already had working trigger handling and data acquisition protocols, the NCD electronics and DAQ were made to conform to SNO standards. During full operation of the SNO detector with NCDs installed, all events within NCDs must fire SNO-wide triggers, with the specification that they originate from the NCD subsystem.

If the main SNO electronics are offline, the NCD electronics must be able to independently keep track of runs as well as events, primarily to remain “supernova-live” as much as possible. Therefore, the NCD electronics must operate with an independent data acquisition system that is able to correlate events with the main SNO DAQ. The NCD electronics must also be able to generate event numbers, so-called Global Trigger Identifiers, or GTID, to correlate coincident events, and to attach GPS times.

The underground environment of SNO and the NCDs places a number of requirements on the electronics. The HV system for the NCDs must be remotely monitored and operated, with power cycling resulting in a safe condition. Additionally, the electronics must be able to self-calibrate and diagnose system failures in real time, so that single-counter issues do not increase system-wide downtime. The infrequent access afforded by the mine environment demands robust remote operation as well as system reset and restart. Finally, the unpredictable nature of mine power and access require extra attention to data backup and integrity.

2.3 NCDs as an array

In this and subsequent sections, we discuss the array-wide signal-processing electronics. Schematic diagrams are too voluminous to be included here, but are available on a SNO collaboration website¹.

All signals that come from the charge deposited on the NCD center anode are run through a signal preamplifier, which separates the current pulse from the high DC voltage. The amplified pulse is carried into one of four multiplexing (MUX) boxes, where it is split into the ‘fast’ and ‘slow’ data paths. The ‘fast’ data path buffers the input signal directly from the MUX box into a shaper. Here, the peak of the integrated signal is measured by an ADC and converted into a digital charge value if the event exceeds the Shaper/ADC threshold. If the event has sufficient peak current to exceed the digitizing threshold in the MUX box, a delayed signal will be routed via a video switch to a logarithmic amplifier and into both digitizers.

The ‘fast’ data path is designed to give basic event information at high rates. In the event of a nearby galactic supernova, neutron rates in the D₂O volume can be expected to reach up to hundreds of hertz. Although the digitizers and the ‘slow’ data path have no ability to maintain this sort of throughput, galactic supernovae should produce neutrons far above the expected background rates in the detectors. Thus, the Shaper/ADCs in the ‘fast’ data path can be used to separate out neutron events by

¹ http://ewiserver.npl.washington.edu/sno/ncd/ncd_internal/ncd_elecdaq.html

measuring integrated charge alone. The 'fast' data path is designed to handle thousands of hertz of data, to allow not only for galactic supernovae, but also for neutron sources that may be introduced into the D₂O volume for calibration.

The 'fast' path ADC information passes through a VME bus to a controller that connects to the controlling computer over an isolated fiber-optic line. In parallel, an event that triggers digitization on the 'slow' path results in the MUX controller placing the hit information onto the line of the 32-bit differential I/O module while triggering the appropriate digitizer. The computer grabs the 'slow' path hit information through VME, and begins to query the digitizer for the digitized event record. Figure 2-8 shows the electronics and data path. It is important to note the independence of each data read-out thread-- the MUX and Shaper/ADC systems run independently.

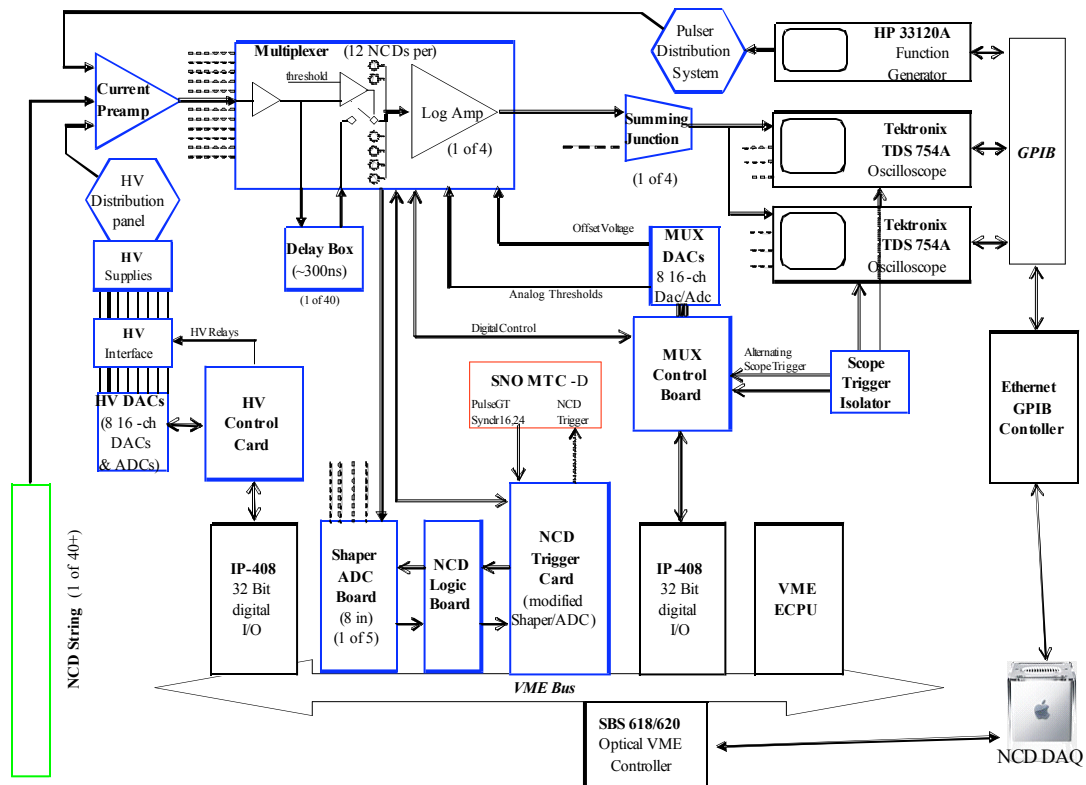


Figure 2-8: NCD Electronics Extended Flow Chart

The NCD system relies on the NCD Trigger Card (NCDTC) for its counting and triggering. The trigger pulses sent from both MUX and shaper systems fire the NCDTC. SNO identifies time-correlated events by requiring each detector element to increment a counter controlled by the Master Trigger Card Digital (MTCD). When any individual detector element observes a recordable event, the NCDTC will send a trigger request to the MTCD. The MTCD's response to a trigger request depends on the current state of the system. The details of the NCD triggering system are discussed in detail in section 2.8.

An independent Embedded CPU (ECPU), Motorola's MV167, prevents short-term computer function from interfering with electronics performance by controlling the electronics from within the VME crate. The ECPU control increases the data acquisition speed of the system by sharing the event organization and data packing duties with the control computer, and maintaining a shared memory buffer between the computer and hardware.

The high voltage distribution system is controlled by electronics very similar in design to the threshold control hardware in the MUX control. There are 8 Spellman 3kV power supplies that connect to a high voltage distribution panel. The power supplies can be manually cabled to different numbers of strings, and each individual power supply has sufficient current and stability to power the entire NCD array under normal conditions. The high voltage control board regulates the HV DACs board and switches the power to the Spellman supplies through an array of relays. The digital to analog converter board (DACs) is used in the high voltage system to set the NCD anode voltages to within a few volts. The ORCA⁴⁹ control of the NCD high voltage system is very similar to the SNO Hardware Acquisition and Real time Control⁵⁰ (SHaRC) interface used to control the high-voltage on the PMT array.

2.4 NCD signal

The signal from the preamplifier naturally depends on the event inside the NCD. Neutron detection is of primary interest, and the counter is designed to bring out their most characteristic interaction features. When a neutron is captured on ³He inside the

counter as discussed in 1.7, a proton and a triton emerge in opposite directions. Thermalized neutrons will result in 764 keV of kinetic energy being shared unequally as dictated by simple momentum conservation between the triton and proton. The position, azimuthal angle and radial angle of the reaction determine what the signal will look like. The triton will not travel as far as the proton, unless the proton ranges out into the counter wall. A fully contained event in which the ionization trail travels non-parallel to the anode will result in a characteristic drift signature with 2 peaks as shown in Figure 2-9.




Figure 2-9: Fully Contained Triton and Proton Drift, delogged, in nanoseconds vs. volts.

If one of the decay particles ranges out in the Ni wall, some total energy is lost from the gas. The signal will be qualitatively different as well, for example as shown here in Figure 2-10.




Figure 2-10: Partially Contained Triton and Proton Event, delogged, in nanoseconds vs. volts.

Alpha particles tracking through the counter gas will deposit much more energy in the same distance. A characteristic event is shown here in Figure 2-11.




Figure 2-11: Fully Contained NCD Alpha Event, delogged, in nanoseconds vs. volts.

The electronics also play host to electronic events not necessarily occurring as a result of ionizing radiation. Microdischarge events, where small amounts of charge

suddenly track across to ground without opening up a channel for more current, also occur. They are typically short-time events, lasting less than 100 ns, with very little net charge, but with significant current. A characteristic microdischarge event can be seen in Figure 2-12.




Figure 2-12: Characteristic NCD microdischarge scope trace, delogged, in nanoseconds vs. volts.

There are also events which are harder to distinguish from interesting physics events; one such is known as a ‘Seattle’ event. A typical ‘Seattle’ event is shown in Figure 2-13. The name originates from a workshop held in December 2004 where these events were first identified. They originate in the J3 (#26) string and currently have an unknown cause.




Figure 2-13: "Seattle" Event on String J3, delogged, in nanoseconds vs. volts.

2.5 Preamplifier

The preamplifiers used for the NCD array⁵¹ are current-mode transimpedance amplifiers with an input resistance adjusted to terminate the 93- Ω cable properly. The NCD Preamps have two main sections: a HV section and a surface-mount integrated circuit board. The HV section is lead-mounted to reduce microdischarge due to surface leakage. The circuit has potentiometers for adjustment to both input impedance and gain. The preamps have a pulser input terminated in 100 k Ω to the input, to provide current injection. The pulse injection point is directly at a FET gate and separated from the cable by a ferrite bead. As a result, only low-frequency characteristics, up to \sim 1MHz, can be reliably calibrated through the pulser input. High-frequency response must be measured

by direct insertion at the preamplifier input. A picture of a preamplifier is shown as Figure 2-14.




Figure 2-14: Open NCD Preamplifier with labeled connectors

2.6 Shaper/ADC

The primary function of the Shaper/ADC cards⁵², shown in Figure 2-15, is to integrate the signal current arriving from the preamplifiers into a total charge per event. The relatively short dead time for each event provides a reliable method of quantifying events at rates that would be unattainable with full signal digitization. Each input channel's trigger threshold level and ADC gains are individually programmable, and can be adjusted during data acquisition. This 'fast' path allows for neutron counting when pulse shape analysis-based event identification is unnecessary, and low energy event analysis when neutron rates are low.



Figure 2-15: NCD Shaper/ADC board

2.6.1 Shaper/ADC requirements

During a supernova burst near the galactic core at around 8 Megaparsecs, we can expect hundreds of neutrons within a 10-second interval. Although the digitizing scopes will not be able to keep up with such a rate, the Shaper/ADCs must be able to get an accurate neutron event count. As very few background events are expected in any 10-second window, an accurate count of neutrons can be acquired by making an energy cut alone. In fact, the expected background in the neutron-capture energy region during any supernova burst should be less than 0.03 events with the exception of certain strings with

high rates of non-physics events. This will allow for very accurate neutron capture counts in both high- and low-rate situations. The few digitized events recorded will only provide assurance that the signals are neutrons.

The charge-integration and high-rate capabilities of the Shaper/ADC also open up the window into lower energy events from the NCDs. Lower energy physics events, mostly expected from electrons resulting from Compton scattering gammas and beta decays, can be observed and tallied for later time correlation analysis. Triggering on lower-energy events within the NCDs might also allow for NCD-SNO event correlation analysis. If, for instance, the NCDs can trigger off of a lower-energy gamma or beta emitted in the ^{208}Tl decay, then it might be possible to look at SNO's PMTs for the corresponding 2.6 MeV gamma (section 5.1).

2.6.2 Shaping

The shaping network allows for a more precise evaluation of the total energy of a particle passing through the NCD. Each channel of the NCD Shaper/ADC board has a four-stage integrating and shaping network, along with an independent peak-detecting circuit. Because the input is a current rather than a charge signal, no differentiator is required. After the pulse is integrated, it is split. One path heads directly into the ADC while the other leads through a dual differentiator and then a discriminator which triggers the sampling of the ADC.

The NCD signal passes through the shaping network of 4 successive Maxim MAX437 op-amp based 4.0- μs integrators and a 50- μs preamp-tail cancellation

network. This imperfect ballistic signal splits in two, the first path of which heads directly into the signal input of the MAX120, a 12-Bit BCMOS sampling ADC with a 350 ns acquisition time and a 1.6 μ s conversion time. The second path triggers the ADC by providing a trigger during the peak of the pulse. The triggering time and level are discussed in the triggering section 2.8.3 below.

2.6.3 ADC

Previous batches of the MAX120 chips used in the main SNO electronics proved unstable under a manufacturer-approved 15 V DC. The MAX120s used in the NCD electronics are from a later batch, which has had the previous flaw corrected. The NCD MAX120s are run at the default 12 V DC, a voltage at which both versions have been problem-free.

Although for the most part, these ADCs work sufficiently well, there were some known issues with the pulse height digitization in the NCD Shaper/ADC cards. The biggest issue was a <1% 2-bit preference. This preference could be seen when running a swept-amplitude pulse train into the shaper inputs. Every other channel has slightly more/less than the previous, and every fourth is higher (lower) still. This effect is documented in 2.6.5.

2.6.4 Shaper/ADC VME Communication protocol

The NCD Shaper/ADC board is designed for operation on a standard 16 MHz VME Bus. The NCD Shaper boards follow standard VME protocol with a 16-bit address

space and a 16-bit data register. Operations requiring more than one VME clock cycle are handled by ‘information-ready’/passive registers.

Each Shaper/ADC board should be run in ‘event-lock’ mode, in which the first signal to go over threshold on a counter will create a board-level lockout pulse. This signal is carried out through a direct cable connection to the NCD Logic Board and fanned out to all active Shaper/ADCs. This lockout pulse inhibits all channels on every board from acquiring additional data while the current event is digitized. This lockout pulse is also responsible for requesting a GTID assignment, and is often referred to as the ‘Board Triggered’ pulse as well as the ‘Multiboard Enable/Disable.’ The event lockout lasts for approximately 230 μ s, during which time, only the scalers (see section 2.6.6) will continue firing.

By convention, each NCD board is addressed starting with a base address that is an even multiple of 100 between 8000 and 8f00. Table A-1 lists all of the NCD Shaper/ADC registers and their functions. The ‘address’ listed is given as an addition to the base address of a particular board.

2.6.5 VME Controller Issues

The VME crate is controlled by a SBS Technologies 618/810 VME Controller Card. The 618/810 connects directly to a SBS PCI card on the DAQ computer through an optical fiber. This PCI card has mapped memory, storing the contents of the VME registers in memory available to be read by the DAQ computer. Unfortunately, there was an undocumented feature of this PCI card that occurred as the bus and processor speeds

increased: the card would claim data was completely written into its memory when it was not adequately ready.

Introducing a 50-ns wait allowed only 1 error in 1 million reads and a 150-ns wait completely removed errors from the readout. However, that time delay, given the required accuracy and data rates, was unacceptable. Fortunately, reading the mapping register on the PCI controller card after writing to it removed all errors. This was then implemented in all systems in February of 2005.

2.6.6 Shaper/ADC Scalers

Every channel on the NCD Shaper Cards has a 16-bit scaler to keep track of the total potential trigger signals it receives. Additionally, on each card there is one 28-bit full-board scaler to keep track of the total event count on each board. These scalers count independently of the analog-to-digital conversions, indicating how many events may have been missed. Each scaler fires on the leading edge of its channel's discriminator, and will not retrigger unless the shaped signal returns to below threshold again. Therefore, the larger dead-time experienced by a channel due to a busy ADC or board-level lockout is not experienced by the scaler.

Due to the effects of shaping, the effective dead time of the scaler is dependent on both the energy and the shape of the pulse. Square waves take approximately $6 \mu\text{s}$ plus half the pulse width to reach peak value through the shaper. It takes slightly more than twice that time for the peak to return to below threshold, leading to an effective dead-time on each scaler of approximately 3 times $6 \mu\text{s}$ plus 1.5 times the pulse width. Most

neutron events are similar to a 1 μ s-width square wave, so they result in approximately 20 μ s of dead-time.

In November of 2004, the triggers for the scaler system were modified to allow a more accurate count of total events that could fire the discriminators. Any pulse significant enough to fire the discriminator starts a series of 6 flip-flops. The first flip-flop resets itself after the pulse goes low again. The next all rely on a 4-MHz (16 MHz VME clock divided by 4) scaler shift clock to enable each of their Q-outputs. Each Q-output is hooked into the D-latch input of the next scaler, effectively delaying the signal. After 3 flip-flops, the Q-output is hooked up to be the clock of the next flip-flop. The D-latch input of the scaler logic clocks on the logical AND of the Shaper/ADC line and the ADC converter busy line. The Q-output of this scaler sends a short pulse out, only between 250 and 500 ns.

The NCD Shaper Card Scalers are reset at the start of every run and are enabled immediately before the card's ADC functions are enabled. On DAQ request, the value stored in each scaler is latched into an individual register between BASE+20 and BASE+2E for readout. Given the reasonably low scaler deadtime, the channel scalers can provide an accurate count of actual missed events on a channel-by-channel basis.

2.7 Multiplexer/Digitizer

2.7.1 Multiplexer Design/Function

The maximum drift time for primary ionization in the 85:15 $^4\text{He}:\text{CF}_4$ gas mix at 2.5 atm is 2.5 μs . The fastest structure in the ion pulse is limited by lateral straggling to a few tens of nanoseconds. In order to have a chance at accurately determining Z-position, a digitization sampling time no more than 2 ns is desired. The ion tails caused by motion of positive ions persist for many microseconds after the primary ionization hits the anode. Therefore, for simple and complete event-by-event signal identification, each of the 40 strings of NCDs requires nanosecond resolution for periods of microseconds. Through multiplexing, we are able to take 8 channels of digitizers and convert all 40 channels into digitized data, with the ability to record simultaneous events. Naturally, we must have individually adjustable thresholds, for which we use the DACs boards. In the event of near-simultaneous firings, we attempt to digitize and read out as many events as possible, leaving the remainder to the Shaper/ADCs.

2.7.2 Signal Multiplexing

Signals arriving from the current preamplifiers are split into two parallel buffer amplifiers. The first 2x amplifier copies the input line to a shaper/ADC output line. This will copy the input line without discrimination of events, so that a separate threshold can be maintained on the shaper data. The second amplifier drives both a delay line and a discriminator. The 320-ns delay line is noninductively-wound RG-58 on a PVC spool.

The frequency-dependent attenuation in the line is compensated for by a peaking network on the input side of the delay line.

The threshold for firing a MUX gate is set through an AD9696, which compares the negative going signal with a voltage set by the DACs. Although the reference level on the discriminator can be set between +9.4 mV and -41 mV, values below -37.6 mV cause an LT1011 comparator to set the 74LS109 MUX FF D-input to zero, disabling the channel entirely. When not disabled, the FF input of the MUX is set to 1 at a positive going transition from the discriminator and reset to zero on initiation of digitizer readout from the MUX Controller Card.

Each MUX box has 12 AD8180 high-speed video switches and a 13th dummy line to reduce signal shifts by providing the common output with a default connection when no gates are open. The analog signal, taken through the AD8180 to the common output of all 12 channels, is then split into a linear and a logarithmic amplifier. It is important to note that there is only one analog line on the output side of the AD8180. Multiple channels firing will add together on this line. An AD8055 handles the linear amplification, which then connects to a test output line.

The analog signal from the MUX is also taken through an AD8307 logarithmic amplifier. Logarithmic amplification enables the system to digitize a large dynamic range with fixed relative precision.

2.7.3 Multiplexer Log Amp

A signal of sufficient peak voltage V_{in} causes the AD9696 to fire a 74LS109 FF which triggers the appropriate AD8180 switch, allowing the signal to pass through to the log amplifier. Each multiplexer board has one AD8307 log amplifier, which is individually controlled. More detail on the shaping and triggering of this system can be found in appendix C. The output voltage of the AD8307, V_o , is represented in standard form by the equation:

$$V_o = a \log\left(1 + \frac{V_{in}}{b}\right) + c,$$

where a , b , and c are constants to be determined by calibration.

The offset of the log amplifier for each multiplexer is controlled by setting the 13th DACs output on the MUX board. This biases the AD8307's offset line, thereby setting the zero-value of the log amplifier. As the voltage on this line can vary from 0 to 2.5 volts, standard settings fall somewhere in the range of $(80/255 \text{ or } 0x50/FF) * 2.5$ volts.

The offsets are individually determined for each log amplifier by injecting a 500-kHz, 10-mV peak-to-peak sine wave directly into the MUX input and examining the amplified pulse. The offset was set to the voltage at which the positive peak of the sine wave just began to rectify. This allows small reverse-polarity pulses (up to +10 mV) as is expected from thermal noise and electromagnetic interference to pass through the MUX without chance of misidentification.

2.7.4 Digitization

Digitization is handled through the use of Tektronix 754A digitizing oscilloscopes. The 754A is a 500MHz, 4-Channel digitizing oscilloscope with a sample rate of 2 Gigasamples/Second, divisible by either 2 or 4 channels. Each digitizer has 8 bits of vertical resolution with record lengths up to 500,000 points. Nominally, we have set the time base to 2 ns per point and 100 mV per division. The scopes are set to trigger on a TTL pulse applied to the auxiliary trigger. The scope trigger output line is monitored by the electronics to determine the availability of each oscilloscope.

2.7.5 MUX Controller Board

The multiplexers interface with the system through the MUX Controller Board. The MUXCB monitors the firing of the multiplexers and triggers scopes accordingly. The MUXCB can then reset the multiplexer channels as well as monitor the status of the digitizers directly. Setting and readback of the MUX threshold values is done through the MUX DAC/ADC board. A picture of the MUXCB with connectors labeled, is shown as Figure 2-16.




Figure 2-16: MUXCB with labeled connections

The MUXCB is a 6U-sized (233.5 mm by 160 mm) circuit board, designed for horizontal mounting in the NCD electronics rack. The MUXCB is a purely digital board; all analog work is handled by the MUX DACs Board, described in the following section. At the heart of the MUXCB's digital logic is the Altera 9480, a 208-pin Field Programmable Gate Array (FPGA). Each MUX sends digital communication directly to the MUXCB through one of eight 40-pin MUX Control Cables. The command set for the MUXCB is described below.

When an event on a MUX channel goes above the set threshold, the MUX asserts a 'MUX_Box_Triggered' signal down the control line directly to the MUXCB. Multiple

channels firing on the same multiplexer do not change the 'MUX Box Triggered' state. The MUXCB cannot know how many channels in a triggered MUX box are triggered until it actively queries the box.

On the firing of the first MUX, a scope trigger and an NCD_GTID pulse are asserted for a time between 2 and 3 MUXCB 10-MHz clock cycles. If the system is in "Toggle Scopes" mode, the MUXCB will attempt to trigger the scope which did not fire on the last digitized event. If the MUXCB is in single scope mode, the system will merely attempt to retrigger the same scope every time. Note that system performance can be suboptimal if there is an uneven scope processing or communication delay that software does not account for². The 'Scope Busy' state of the digitizer can be read through the 'Read Event Register' command. In the event that both scopes are busy and the system is forced to reset, the MUXCB will attempt to trigger the next scope in the sequence.

As soon as the control computer has completed the acquisition of digitized data from an event, the MUXCB must be told to reset the triggered MUX. "Mux Arm/Rearm" is a one-command function that forces the MUX to send a reset pulse to all of the flip-flops that control the AD8180 switches. Without a specific request to the MUXCB, the MUX switches will continuously allow events to pass through, but the

² A possible future upgrade to the logic and data acquisition system is the addition of a 'missed-event' flag which would be added to the programming of the MUXCB Altera. This flag will monitor the firing of the various MUX boxes, and determine if any 'MUX Box Triggered' lines go high while in a state that would prevent the firing of a digitizer. This state can occur when both scopes are busy or during the short amount of time between when the oscilloscope completes digitization and the 'scope busy' line flips the trigger to the other scope.

MUXCB will never trigger on a new event from that MUX. In this case, the ‘Mux Triggered’ line, which is set by a logical OR of all 8 MUX-Box-Triggered lines, will be continuously asserted.

The MUXCB is purely digital, and holds neither the digitization threshold values nor the log amp offsets. These settings, which are direct line inputs on the MUX boards, come from the MUX DACs board. The MUX DACs relays all of its communication through 2 ribbon cables that connect to the MUXCB. The first cable, the ‘Address Line,’ carries the address of the channel that the MUXCB is adjusting or reading. The second cable, the ‘Data Line,’ carries the digital value of the threshold setting request as well as the digitized value of the last channel read. Note that both the Value and Address lines have separate conductors for input and output signals, and there are no bi-directional lines.

2.7.6 MUX DACs Board

The MUX DACs board is the analog extension of the MUXCB. It relies on constant communication with the MUXCB for setting adjustment and readout, but will hold the line values in event of a MUXCB disconnect or power off. The heart of the MUX DACs board lies in its 16 8-channel DACs and ADCs. With 10 26-pin ribbon-cable connections, the MUX DACs holds the analog MUX threshold and log offset values down on differential line pairs. Each MUX board has its own cable to the MUX DACs board, currently leaving 2 unused and unimplemented spare lines. A picture of the MUX DACs board is shown as Figure 2-17.




Figure 2-17: MUX DACs board with connections

The MUXCB can only command the MUX DACs to either set or read back a threshold. In either event, the MUX DACs expects the MUXCB to first assert an address which specifies which DACs or ADC channel will be used. The MUXCB then needs to follow with a read or write request before the information is removed from the address lines. Write and read operations to identical single-channel addresses access the same pin on the board, but the data paths diverge from there.

A threshold ‘set’ request requires that the MUXCB place the channel of the setting on the address line, then the setting itself on the data line, followed by a single high-going pulse on the ‘set threshold’ part of the address line. Once the MUX DACs

has placed the threshold onto the differential MUX setting line, it will assert a ‘Threshold set complete’ on the data line. This is the handshake the MUXCB looks for before it returns a ‘DAC Write Complete’ signal back to the control computer through the IP408. The IP408 is an I/O board described in more detail in section 2.7.7.

The DACs board receives and relays all digital communication in the system through two parallel cables, the DAC_Address_Bus and the DAC_Data_Bus. The DAC_Address_Bus is a 20-pin dual differential bus carrying commands from the MUXCB to the DACs as well as the busy status of the DAC’s ADCs. The first differential line carries the ADC_Busy line to the MUXCB, which the DACs board will assert if any of the MAX158 ADCs are putting out an INT* (ADC Busy) pulse. This prevents the MUXCB from grabbing incomplete data from the DAC_Data_Bus. The next 2 differential lines carry the read or write request of the MUXCB, with the MUX Box and Channel being selected by the next 3 and final 4 lines, respectively.

The DAC_Data_Bus is a 34-pin dual differential bus which carries information to or from the DACs or ADCs. The first 8 differential lines (16 total) are dedicated to the MUXCB reading the ADCs, whereas the second 8 differential lines are reserved for the MUXCB’s digital DAC value. In this manner, the first half of the pins are reserved for input to the MUXCB and the second half are reserved for output to the DACs. The last differential line is grounded on both sides.

Before the MUXCB writes a MUX threshold value, it must be first instructed which MUX box and channel will have its threshold set. This is done with a ‘Write

Channel Address' command which proceeds the 'Set MUX Threshold' command. Once this has been done and a 'DAC Write' command issued, the MUXCB places the contents of the 'Computer Data Register' onto the 'DAC Data' bus, and waits for two 10-MHz clock cycles to pass. This places a minimum time between the 'Write DAC' request and the 'DAC Write Strobe' being passed unto the MUX DACs board of 100 ns, to prevent the write request preceding the actual write data. The copying and waiting is done inside the Altera on the MUXCB.

When the differential DAC "Write Strobe" is received at the DACs board on the Address Bus, it is put directly into the "WR" input of the AD7228A DAC. Only the previous "Write Channel Address"-specified 8-channel AD7228A digital-to-analog converter is sent the "DAC Write Strobe". The DAC digital values, however, are carried through the Data Bus into all of the MUX DACs board AD7228A's D0 through D7 data inputs. The output channel on the AD7228A's is specified through the Address Bus and copied to all of the DAC input A1-A3 lines. Therefore, all channel and data information is sent to all DACs on the board, but only the board that receives the "DAC Write Strobe" uses any of that information.

The MUX DACs board sends the analog MUX trigger threshold directly into each MUX box through a 26-pin ribbon connector. Every 26-pin connector carries 12 analog thresholds, one log offset, and 13 ground lines. The odd pins are all grounded, and the 12 analog thresholds are in order of channel number, starting at pin 2 and going to pin 24. Pin 26 holds the log offset. There are a total of 10 connectors, 8 for a full 96-string array

and 2 spares, for a potential 120 thresholds and 8 log offsets, since the spares do not carry log offsets.

Since there are 13 analog lines for each MUX and the DAC and ADC channels come in groups of 8, each MUX threshold cable has 2 sets of 8-channel DAC/ADCs assigned to it. The highest 3 channels in each group of 16 channels are routed to the 'spare DACs' lines. Therefore, counting from 0, a request from the MUXCB to set channel 12 of MUX 0 would result in setting the 'log offset' for MUX 0, whereas a request to set channel 13 of MUX 0 would set Spare 0, channel 0. As there are 8 total standard DAC/ADC channels, we have 24 of these 'spare' lines.

Due to the relatively long response time of an ADC, the MUXCB must separate the reading, latching, and returning of data with significant amounts of time. As soon as the MUXCB receives the 'Threshold Read' request, all values currently held on the threshold section of the 'DAC_Data_Bus' are reset and a 'Read_ADC' is sent out to the DACs Board. The MUXCB will then wait until the DACs board returns an 'ADC_Busy' before continuing. The MUX DACs receives the 'Read_ADC' through the address line, and decodes the address associated with this request with a LS154 decoder. The read request is sent to the individual MAX158 ADC in question, and the channel address is carried to all ADCs. Since only the ADC which has the desired channel receives an 'ADC_Read' request into its 'RD' input, it is the only ADC to use the channel information. As soon as the MAX158 gets the request, its 'INT' line goes high, and this 'Busy_ADC' is carried back to the MUXCB across the Address bus.

On receiving a 'Busy_ADC', the MUXCB waits 2 10-MHz clock cycles before sending off an internal 'ADC_Data_Ready' signal, which attaches the values on the DACs Data Bus onto the MUXCB's data return register. Simultaneous with the start of this, the MUXCB begins to wait for another 2 10-MHz clock cycles before locking in the values on the data return register.

2.7.7 MUX Communication through the IP408

The IP408 is a commercial 32-channel combination input/output I/O card⁵³ that sits in an AVME9660 VMEbus Carrier Board⁵⁴. Three IP408s are used in the full system, one for each MUX system, high-voltage control system, and pulser distribution system. Each IP408 mounts as a daughter card in the AVME9660, and all VME commands, for input and output, are taken through the carrier board.

The IP408 is controlled by a Field Programmable Gate Array, with 8 (channels 0-7) of the 32 channels able, but not used in our instance, to create interrupts on the VME Bus. The input buffers are tied to the field input signals through a 100k Ω series-connected resistor, limiting the input current. The IP408 allows for selectable input or output on each channel by connecting 32 open-drain outputs in tandem with 32 input buffers. Each of the outputs is an open drain of an individual MOSFET which is driven by the FPGA. Each of the transistors' sources are connected in common, which allows for digital control of up to 32 low-side switches. Note that sending a '1' to the output on a channel will close the circuit and turn the switch 'ON', while a '0' will leave the circuit open, turning the switch 'OFF'. Reading the input register of an output channel will

yield the current output state of the channel, as the input buffers are connected in tandem with the output MOSFETs. It is important to recognize that the IP408 does not drive the communications line on output, it merely connects the line to drain or leaves it open. On input, the IP408 must have the line open (OFF on output), and then it can read the high or low TTL state. On power-up, the IP408 always starts in an all-OFF state. A block diagram of the IP408 is shown in Figure 2-18.

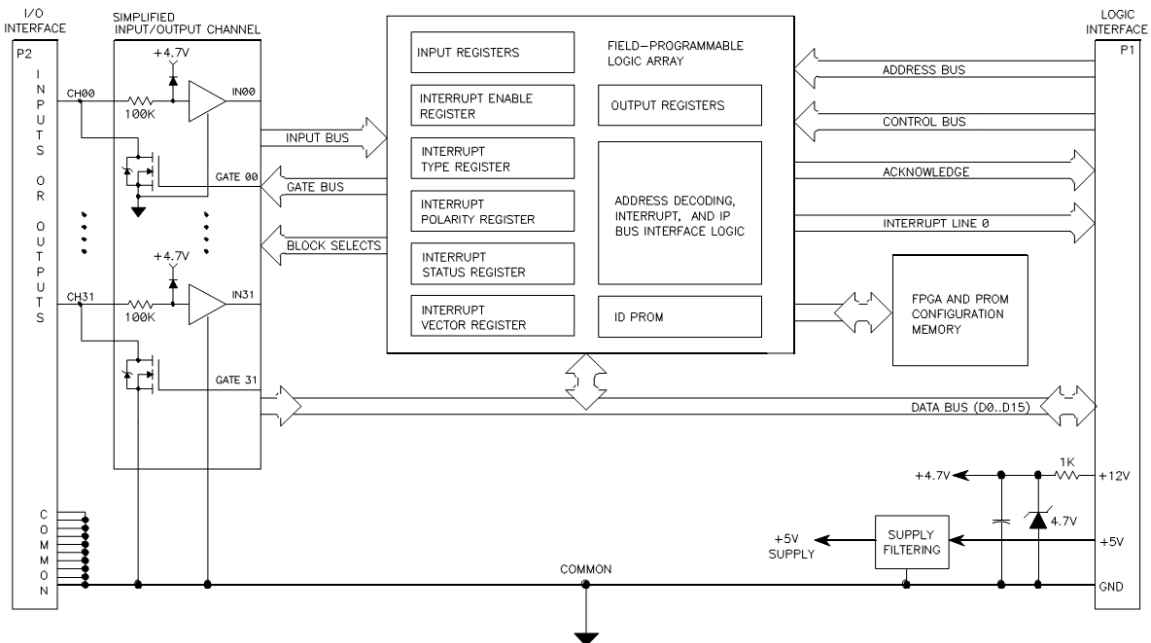


Figure 2-18: IP408 Block Diagram. Figure is from [53].

The AVME9660 interfaces the IP408 with the world, the VME crate on one side, and the 50-pin ribbon cables on the other. Note that pin 1 of the ribbon cables should connect to the bottom of the IP408 module's input section. The communications which run from the IP408 through the AVME9660 to the MUXCB are detailed here:

The MUXCB has sole write authority to the lowest 19 bits on the line. These include: a 3 bit Controller Return Register (CRR), a 15 bit Data Register (DR), and a Controller Status Bit (CSB). The computer/IP408 has sole write authority to the other 13 bits, composed of an 8 bit Computer Data Register (CDR), a Bus Status Bit (BSB), a Reset bit (RB) and a 3-bit Write Register (WR).

The 32 bit bus from the IP408, IO<31...0>, is mapped to the controller as shown in

Table A-3.

The MUX controller has 7 commands plus a MUX Reset. The sequence for the issuance of a command is that the desired command code is loaded into the (Write Register) WR<2...0>, any data required for the command is loaded into the (Command Data Register) CDR<7...0> and then the BSB is set high. It is the low to high transition of the BSB that announces to the controller that there is a command for it. After the controller has finished executing the command, it sets the CSB and puts the return code for the command on the CRR and if that command has data to be sent to the computer,

the data is loaded into DR<14...0>. The computer then lowers the BSB bit and the controller lowers its CSB signifying to the computer that it is ready to receive another command. Figure 2-19 graphically demonstrates this communication protocol.

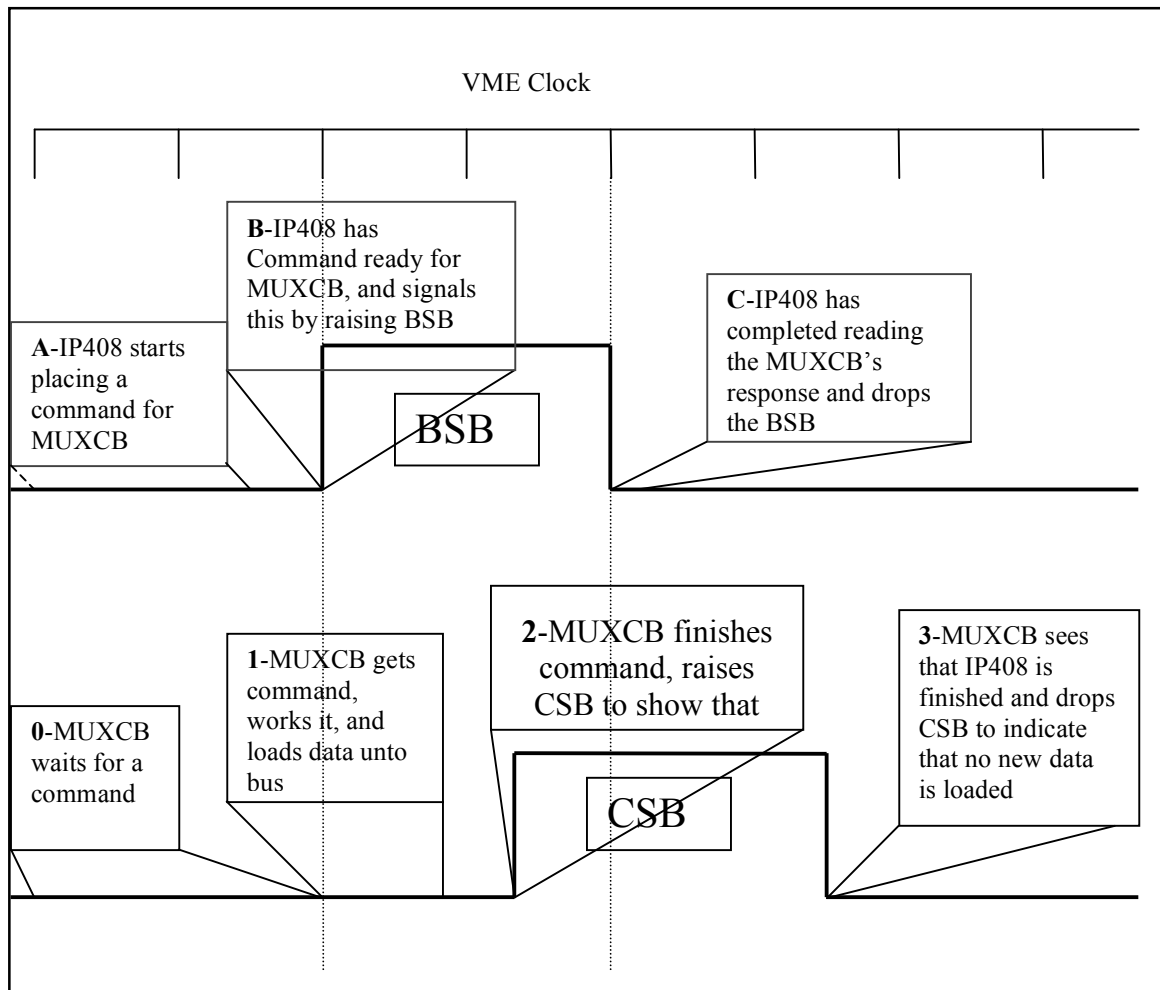


Figure 2-19: BSB/CSB Communications Handshake Protocol. It is important to note that the IP408 uses the VME clock for timing, while the NCD control boards do not.

The MUXCB write register, bits 2 to 0, carries the command descriptor to the MUXCB. The command, whose descriptors are described in detail on Table A-5 in

appendix A, can be: Reset MUXCB, Write Channel Address, Write DAC Value, Read ADC Value, Query MUX Status, Rearm MUX, Read Event Register, and Read Selected MUX. Additionally, a reset can be issued from the controlling computer by sending a mark across the reset line.

2.8 Event tagging

The NCD trigger system⁵⁵ is designed for operating in both SNO-integrated as well as stand-alone modes. SNO-integrated mode requires that the NCD trigger assignment conform to the Global Trigger IDentification (GTID) protocol. In the stand-alone mode, the NCD trigger system must be able to provide unique event identifiers as well as accurate timing values in a format compatible with SNO data.

If the SNO trigger system is not currently acquiring an event, the MTCD will send out a Global Trigger pulse (GTRIG) to each of the elements. The element that triggered the MTCD is recorded and put in the SNO event data header. Any other triggers that arrive at the MTCD within 10 ns will be latched into the event as being associated triggers. Any element triggers that arrive at the MTCD between 10 ns and the end of the LOCKOUT phase (~400 ns) will be ignored and its late GTID request used to raise the ‘missed trigger’ bit.

The NCD Trigger counter is incremented in step with all GTID counters in other SNO DAQ electronics by the GTRIG that is sent out by the (MTCD). The controlling NCD Electronics computer grabs the GTID information from the NCD Trigger counter and associates this value with any events coming in at this time.

2.8.1 SNO Global Trigger

SNO's Global Trigger is a system-wide event that results in the recording of data. All hardware connected to the trigger system must be either allowed or disallowed to set off the Global Trigger at the MTC/D. Hardware enabled to fire the global trigger is 'masked in' while disabled hardware is referred to as being 'masked out'.

Each global trigger is assigned a number on a 24-bit counter providing a unique identifying number within a given run. Operations with this Global Trigger ID (GTID), such as loading and reading, are done with two separate requests, the lower-16 bits and the upper-8 bits acting independently. Synchronizing all Front End Cards and peripheral devices is done through a synchronous clear command (SYNCLR) carried on the trigger cable. The NCD electronics conform to the counting and command structure that the SNO electronics expect from every attached experimental element.

2.8.2 MTC/D Operation Overview

SNO's Master Trigger Card/Digital (MTC/D) receives trigger pulses from the PMT array, calibration sources, monitoring hardware, NCD hardware, and short-term test equipment. In turn, the MTC/D must decide whether the pulses it is receiving warrant a global trigger being sent to the Front End Cards and NCD Trigger system. If the Global Trigger is sent out, all GTID counters must latch any pertinent data with the appropriate ID. The MTC/D is tasked with maintaining synchronized GTIDs across all triggered devices, as well as storing all relevant event-wide information in on-board memory for inclusion in the event header.

The MTC/D expects a 20-ns-long candidate trigger pulse from the electronics component that wants to signal a trigger. The candidate triggers come from every potential trigger source in SNO, from the analog sum of fired PMTs in the MTC/Analog to the NCD electronics. The MTC/D can ‘mask in’ or ‘mask out’ any combination of candidate triggers (inputs), so that the MTC/D will only fire the global trigger (output) from certain sources. Ten nanoseconds after a candidate trigger pulse that is masked in, or enabled to fire the global trigger, the MTC/D will latch in all of its source trigger states. This gives the state of all trigger sources within 10 ns of the triggering event, even if they are ‘masked out’. During the 60 ns after the trigger word latch, the MTC enters the LOCKOUT phase, during which any additional trigger source will add a ‘Missed Trigger’ to the trigger word.

Only the MTC/D is permitted to assign a GTID, but the MTC/D never sends out the value of the GTID. Therefore, any device attached to SNO that has relevant, event-specific information must keep track of the 24-bit GTID internally. If such a device observes an event satisfying its own trigger criterion, it may request that the current GTID be assigned to the event, but it may not index the GTID autonomously; only the global trigger from the MTC/D causes the GTID to be assigned and advanced.

2.8.3 NCD Trigger System

The Shaper/ADCs and multiplexers begin the NCD triggers by sending out a trigger pulse. The trigger pulse is fired anytime an active channel in either system has an event over threshold, the specifics of which are defined separately for the shaped and

digitized data. The trigger pulse from each system runs through a digital OR in either the Shaper Logic Board or the MUX Controller. The summed triggers are sent separately to the NCD Trigger Card, where the event is assigned identifiers.

A diagram of the basic NCD trigger logic is shown in Figure 2-20.

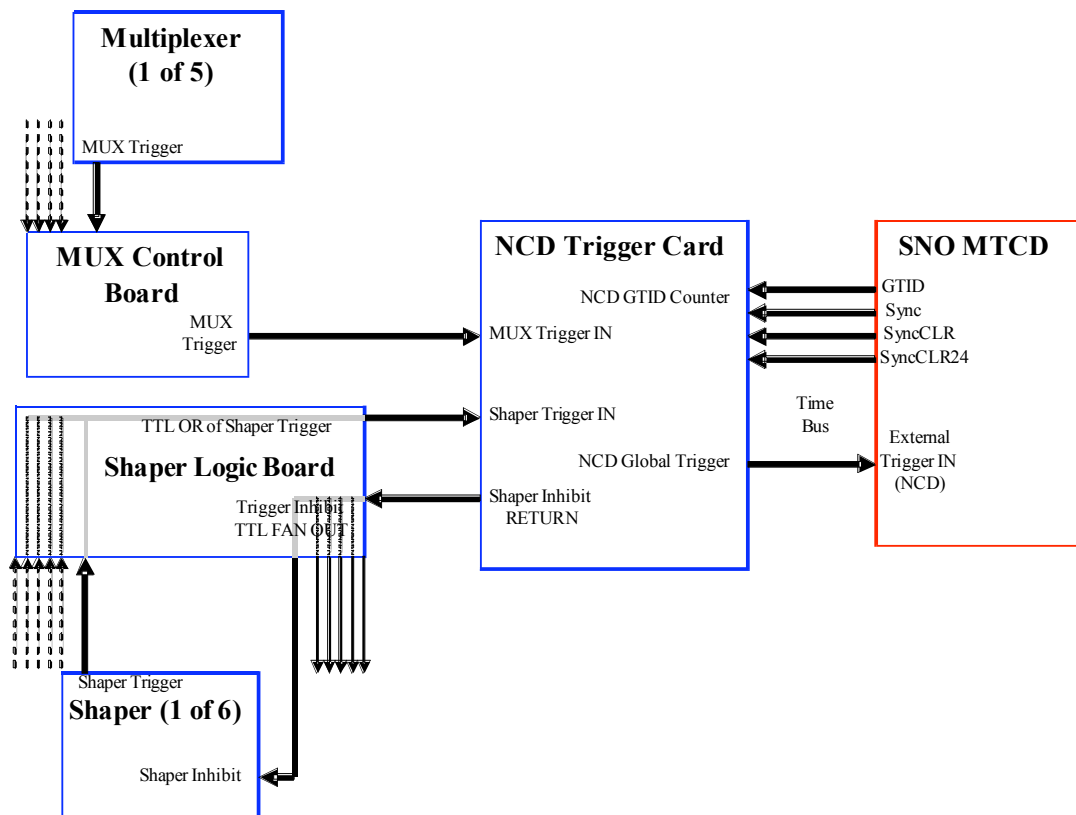


Figure 2-20: NCD Trigger Logic

As described in section 2.5, the signal going into the Shaper-ADC is shaped through four integrating circuits before being split off into a signal and a trigger generating section. The trigger generating circuit differentiates the pulse to make a steep bi-polar pulse, the zero crossing naturally centered on the peak of the shaped signal. This

differentiated signal then runs through an LT1191 op-amp in simple multiplicative gain mode, whose output is sent directly into a comparator. The zero-crossing is necessary but not sufficient to fire the comparator, as sufficiency would create triggers during the smallest of analog pulses.

Meanwhile, the integrated signal is constantly compared to the DC level held by the programmable buffered AD7726 DAC. If the signal pulse is greater than the DC level, the comparator will assert a sample and hold (S/H) when the differentiated pulse has its next zero-crossing. The timing of the differentiated pulse means that the pulse height is compared typically at the peak of the analog pulse, ~6 microseconds after its onset.

The sample and hold generated is a logic pulse that starts and stops the analog-to-digital conversion. The 'Sample-ADC' is latched high until the read of the ADC is complete. Multiple triggering during a single Sample-ADC firing is impossible.

A copy of the comparator output is also sent to a scaler counter which records the number of downward-going zero-crossings while the signal is above the charge threshold. This scaler allows the ADC to report the total number of times the channel has had events over threshold, even if the ADC was busy and unable to record the event. If the scaler advances during conversion, we must assume the data of the latter events lost.

During every A-D conversion, within 180 ns of the onset of the Sample ADC being asserted, a Channel Triggered is fired. Each Shaper/ADC board runs a logical OR on all of the Channel Triggered lines and places the result on the Board Triggered line.

The Board Triggered, a TTL state, is connected to a LEMO output connector at the top of the front panel of the card. This signal, which can be used for event lockout as discussed in section 2.5, is sent through the Shaper Logic Board where it is turned into a Global Shaper Event Signal. A timing diagram is shown as Figure 2-21.

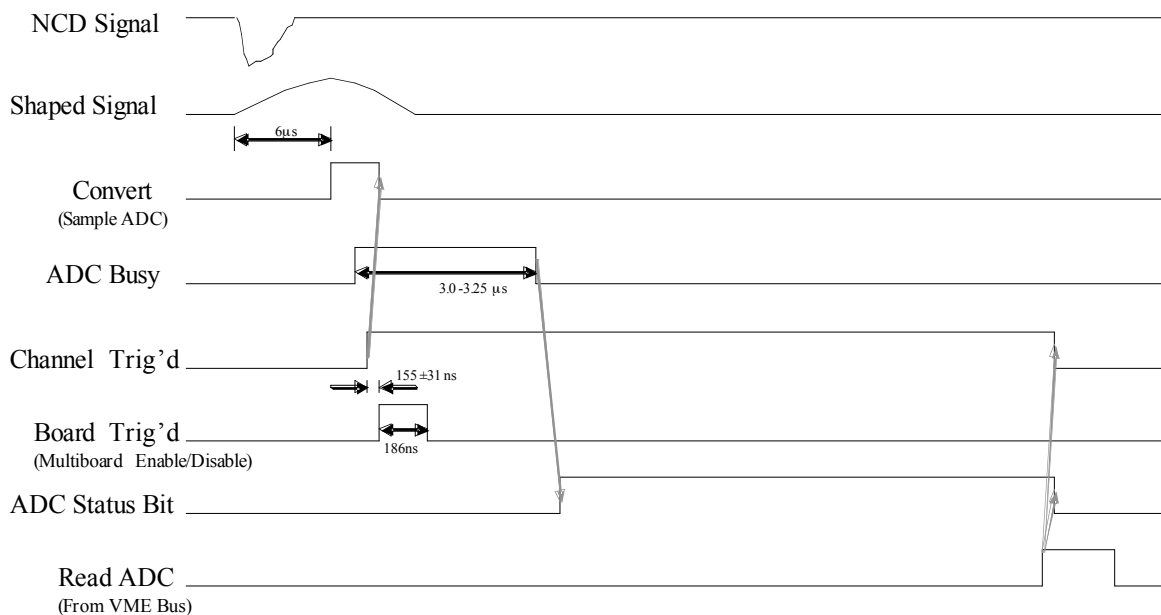


Figure 2-21: Shaper/ADC Trigger Logic

There are two triggering modes that the NCD Shaper/ADCs can be set to: Multiboard Mode and Continuous Mode. Multiboard is the standard mode for operation of the NCDs, while Continuous Mode should only be set during non-physics (eg. testing or calibration) runs.

In Multiboard mode, the Board Triggered line turns on when any of the Channel Triggered lines fire, as shown above. While the Board Triggered line is held high, there is no internal data acquisition inhibit. During Multiboard mode, the Inhibit line must be

set externally by holding the Inhibit In to TTL high through a LEMO connector. The Board Triggered line must be externally monitored through TTL levels on a separate LEMO connector. The NCD Shaper Logic Board monitors the Board Triggered line of each Shaper Board in the array, and puts out an Inhibit In pulse to every board.

Continuous mode turns off the board-level triggering, so that no Board Triggered pulse will be sent away from the board. Continuous mode also ignores board-wide Inhibit pulses, allowing synchronous data acquisition at the cost of positive event identification. Single channels still lock themselves out from further conversions while in the process of converting a signal.

2.8.4 NCD Logic Board

The NCD Logic Board is the interface between the Shaper/ADC and the NCD Trigger Card. On the start of a signal, the NCD Logic Board makes a logical OR of up to 8 Board Triggered signals and fans them out through two identical Global Shaper Event Signal lines. One of these lines is sent into the Shaper Event In LEMO connector on the front of the NCD Trigger Card.

On the return side, the NCD Logic Board accepts a Global Shaper Inhibit Out returning from the MTC Trigger Card. The Global Shaper Inhibit Out is fanned out to each Shaper/ADC's Board Inhibit IN LEMO connector. The NCD Logic Board is a 6-U sized VME card that puts out and reads TTL logic signals through LEMO connectors..

2.8.5 MUX Trigger

As mentioned in section 2.7.5, events that go above the digitization threshold will cause the MUX board to send a 'MUX Box Triggered' pulse to the MUX Controller. Note that this trigger pulse will fire as soon as the NCD signal crosses threshold, unlike the Shaper/ADC pulse which must go through a shaping network. The MUX Controller, in turn, will send a trigger to the appropriate scope. If a scope trigger is sent out to either scope, a Global MUX Event signal is sent out through a BNC connector on the MUX Controller. The Global MUX Event signal is a NIM-logic pulse that is carried directly to the NCD Trigger Card. A MUX timing trigger diagram is shown as Figure 2-22.




Figure 2-22: NCD MUX Trigger Timing Diagram

2.9 NCD Trigger Card

The NCD Trigger Card is designed for operation on a standard 16MHz VME Bus. The NCD Trigger card is made with the digital portion of the NCD Shaper boards, and therefore uses a similar communication protocol. In contrast to the NCD Shaper boards, the NCD Trigger card is addressed starting with a base address of 7000. Additionally, the Altera 7192 FPGA that controls the NCD Trigger card has been reprogrammed to match NCD events with times and SNO event numbers rather than amplitudes.

Primarily, the NCD Trigger Card functions as the hardware interface between the NCD and SNO electronics. The NCD Trigger Card communicates with SNO electronics through the Master Trigger Card/Digital while counting event triggers on the NCD Shaper/ADC boards and the MUXCB. Integration with the SNO Trigger System⁵⁶ requires that every subcomponent synchronize with a 24-bit GTID counting protocol and interface through a 4-signal 'time bus' (see Figure 2-20). The NCD trigger card also maintains four separate clocks that track the NCD electronics livetimes on a run-to-run basis.

2.9.1 NCD Trigger Card Event Logic

The NCD Trigger Card (NCD-TC) watches for triggers in the NCD Shaper/ADC boards through the LEMO input at the top of the card. The signal comes from the NCD Shaper Logic Board's sum of the NCD Shaper/ADC board enable/disable outputs, sending one pulse per ADC firing. When such a Global Shaper Event Signal is received, the NCD-TC sets the Shaper Event Bit on. Coincident with this bit is the NCD Global

Trigger Out pulse, which is sent through ECL logic on a BNC connector bound to SNO's MTC/D. This signal is connected to the MTC/D at the External NCD Trigger input. The NCD-TC then prepares to latch the GTID register on receipt of a GTID pulse from the MTC/D, so long as the Shaper Event Bit is high when the Global Trigger returns from the MTC/D. In this condition, the Valid Shaper GTID Clock Bit is set high upon receipt of the MTC/D's GTID, and will remain set high until it is read or reset by software. The NCD Trigger Card timing logic is shown below as Figure 2-23.

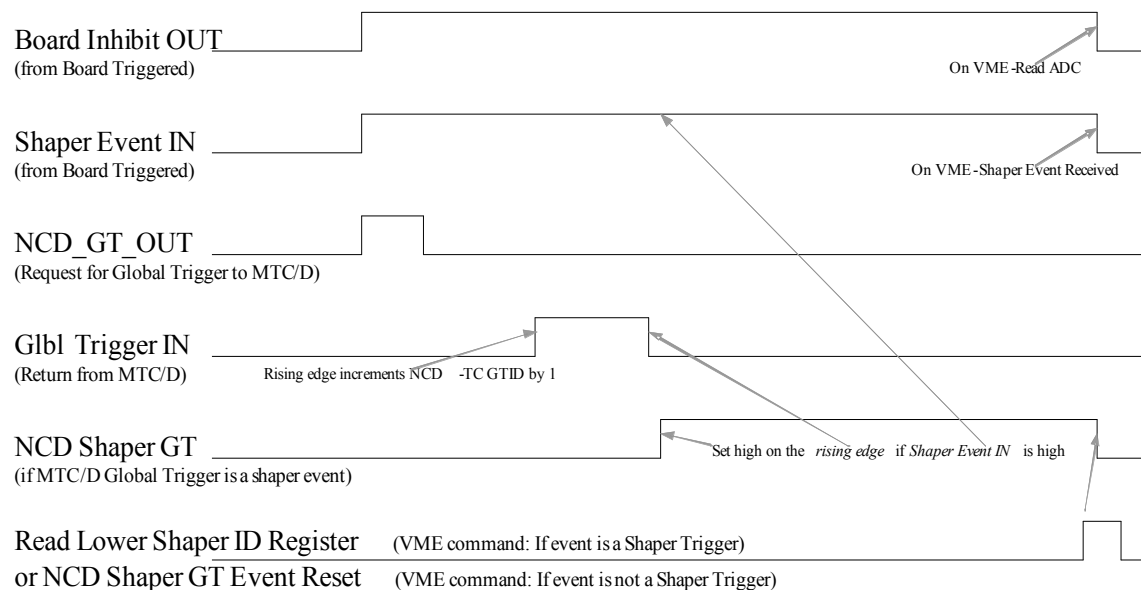


Figure 2-23: NCD Trigger Card Logic

The NCD Trigger Card also accepts NCD MUX events from the NCD MUXCB, corresponding to one (up to four-channel) digitization per pulse.

After each event, the NCD Trigger board assigns a 24-bit GTID to the event. Should the SNO MTC/D not be available, the NCD Trigger board can also assign a 48-bit clock counter to the GTID. This optionally-active local NCD Trigger clock counts cycles of an internal 10 MHz precision clock. The GTID associated with the event is handed off on a read request on the VME Bus. All of the read and write requests are shown in Table A-8 and Table A-9.

2.9.2 NCD Trigger Card Livetime Counters

The NCD Trigger Cards have a 10 MHz clock that is divided down to a 1-MHz timing signal inside the NCD Trigger Card FPGA. This timing signal splits four ways and eventually runs into one of four separate 40-bit counters. The counters may be all be globally reset or individually enabled or disabled. At the beginning and end of every run, each livetime counter must be read and latched in. A diagram of the Trigger Card Livetime Counters can be seen in Figure 2-24.

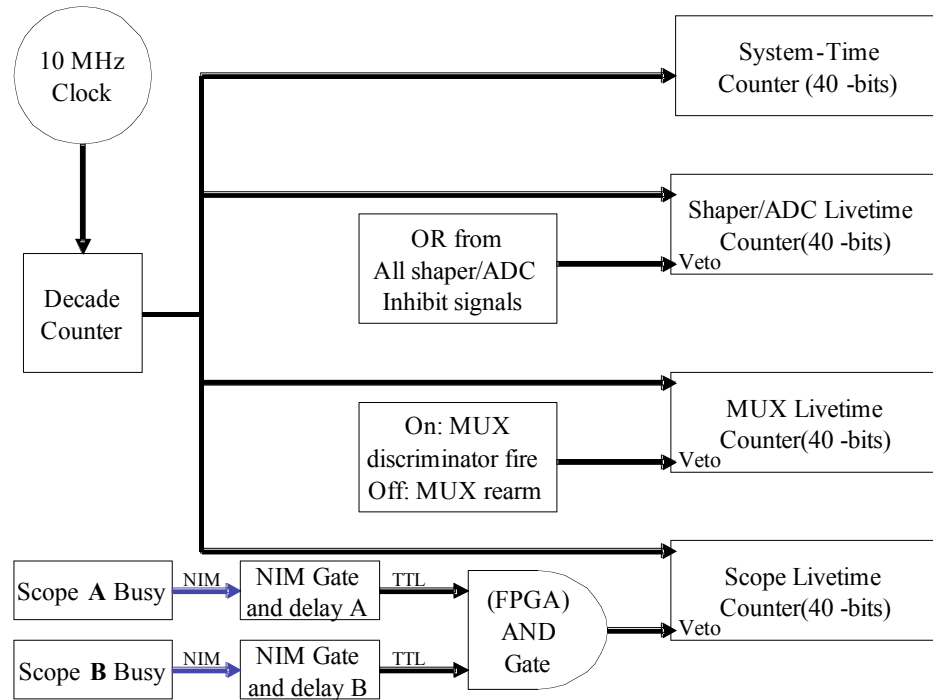


Figure 2-24: Trigger Card Livetime Counter Logic

The first timing signal goes straight into its 40-bit counter, effectively measuring the total number of 1-MHz timing signals that are seen. This is the clock that measures the amount of time the whole system has been functional. As 40-bit rollover occurs every 12.73 days, there will be occasional rollover. Although the counters in the Altera FPGA can see overflow in internal counters, this function is not used. Instead, we check for overflow manually, as run length should never approach 12 days.

The second timing signal counts the time that the system has been operational while none of the shaper inhibit signals are on. Effectively, an inhibit signal from any Shaper/ADC will veto this counter for about 230 μ s (see section 3.3.2) and the resultant

total count at the end of a run will be the amount of time the Shaper/ADC system was ready to take new data.

The third counter reads the livetime of the MUX system. Whenever a MUX discriminator fires, this counter is turned off. When the MUX is rearmed the third counter once again counts. This would be an accurate measure of dead-time except for the fact that the MUX discriminator stays live while the event is being digitized. This 20- μ s window will not be counted as 'live-time' automatically, but can be used in within any given analysis routine, for example to look for short bursts above MUX threshold.

The fourth counter measures the effective livetime of the digitizing scopes. As there are two scopes in the current system, this counter is vetoed on the logical AND of both Scope Busy signals. Unfortunately, the Scope Busy signals are in NIM standard, whereas the NCD-TC FPGA expects a TTL signal. Currently, a NIM gate and delay module is being used to generate a fixed-length 'busy' signal for each of the scopes, rather than reading them directly.

All of the counters increment on the rising edge of the 1 MHz signal. The values from the counters are latched into the read registers on the falling edge following a latch request. No counter will function unless its enable line has been set high, and all of the counters may be simultaneously reset by software.

2.10 HV Control System

Given the size and gas fill of the NCDs, the high voltage system must be able hold the center conductor at 1944 V for a nominal gas gain of 200. The gain of the

counter depends heavily on this voltage, with a 0.05% shift in voltage (1 V shift at nominal Voltage) resulting in 0.7% change in gain. Although designed for eight, the final deployment system has six separately controlled channels. This independence allows for defective or questionable electronics channels or NCD strings to be isolated and removed from the HV system while disabling as little of the array as possible.

2.10.1 Proportional Counter Gain and Voltage

For the purposes of calculating the voltage-dependence of gain M , we may treat the NCDs like any standard proportional counter, using the Diethorn parameters⁵⁷:

$$\ln M = \frac{V}{\ln(b/a)} \frac{\ln 2}{\Delta V} \left[\ln \frac{V}{a \ln(b/a) P} - \ln K \right],$$

where V is the voltage difference between anode and counter wall, ‘ a ’ describes the radius of the anode, ‘ b ’ describes the radius of the cathode, ‘ P ’ is the counter gas pressure and ‘ K ’ is a constant of the gas mixture. The gas gain is, therefore, a function of both the physical characteristics of the counter and the gas itself. There is a delicate balance between the gas gain on the counter and the noise on the electronics played against the self-screening effect of charge at high gas gains and the loss of energy resolution. As the voltage across the counter is increased, the signal-to-noise ratio increases, but the resolution decreases at the point where particles inside the gas ionize the majority of charge carriers. Essentially, at sufficient counter voltage, proportional counters become Geiger-Muller counters that report only the presence of an ionizing particle, with little information about its energy.

2.10.2 HV Control Board

The HV Control Board allows for control of the HV DACs to set the voltage and the low-voltage relays that allow supplies to be switched on and off. All communication to the HVCB is handled through a 50-pin ribbon cable from an IP-408 I/O module sitting on a carrier board in a VME crate. The HVCB sets and monitors the voltage for the HV Power Supplies through the HV DACs Board, whereas it reads and sets the power status of the HV PS's directly. The HVCB also reads back the current draw of each HV supply independently.

The HVCB is identical in design to the MUXCB, but the Altera has been reprogrammed to deal with a different instruction set, and a few of the IC components have been modified or removed. The main control and monitoring of the HV components is done through the HV DACs board, detailed later. In order to increase voltage setting precision, the HV Control system adds two DACs channels together, weighting one 256 times the other one.

Each of the 8 Spellman HV power supplies receives its power through a relay controlled by the HVCB. A single command from the IP408 can set all 8 relays to HV power on or off. The power state of these relays to HV Power Supplies is monitored directly through a BNC mount, identical in position on the MUXCB with the 'Scope_H_Busy' line. The HVCB's HV_Power_State input expects a constant 5V to be delivered from the power supply which powers the Spellman supplies. If at any time, the 5V should drop below TTL high, all of the relays are immediately switched off. This

protects the HV system from the hazards of rapid turn-on in the event power to the HV supplies is lost, while HV DAC settings remain high. In the event the HV DACs board loses power, all of the DACs are forced to return in a low state. If the HVCB loses power, all the relays switch off, and will not switch back on until the HVCB is brought back on and explicitly instructed to reconnect relays.

The computer and the HVCB communicate over a 32 bit bus through an IP408 identical in configuration to the MUXCB. By convention, the IP408 for the HVCB is placed at the address 6100 (slot 2/4), as opposed to 6000 for the MUXCB, and 6200 for the Signal Distribution System IP408. Similar to the MUXCB, the HVCB has sole write authority to the lowest 19 bits on the line. These include: a 3 bit Controller Return Register (CRR), a 15 bit Data Register (DR), and a Controller Status Bit (CSB). The computer has sole write authority to the other 13 bits, composed of an 8 bit Computer Data Register (CDR), a Bus Status Bit (BSB), a Reset bit (RB) and a 3-bit Write Register (WR).

The 32 bit bus from the IP408, IO<31..0> is mapped to the HVCB as shown in Table A-6. The commands are similar to the control and return commands given by and to the MUX CB system.

The controller has 7 commands plus a HV Controller Reset. The sequence for the issuance of a command is that the desired command code is loaded into WR<2...0>, any data required for the command is loaded into CDR<7...0> and then the BSB is set high. It is the low to high transition of the BSB that announces to the controller that there is a

command for it. After the controller has finished executing the command, it sets the CSB and puts the return code for the command on the CRR and if that command has data to be sent to the computer, the data is loaded into DR<14...0>. The computer then lowers the BSB bit and the controller lowers its CSB signifying to the computer that it is ready to receive another command.

2.10.3 HV DACs Board

The HV DACs Board is the analog arm of the HV Control Board. It has 2 digital control ribbon cables which connect to the HV Control Board, and 4 ribbon cables which connect it to the HV Interface Board. The HVIB-bound ribbon cables carry coarse and fine voltage settings, as well as read back HV Power Supply current and voltages.

The analog voltages are set in two separate DACs, a 'coarse' and a 'fine' line. All of the 'coarse' analog outputs come from the zeroth DAC line, and the 'fine' values come out from the first DAC line. The values carried off vary from 0 to 3 Volts DC. The voltage read-back is done through the sixth DAC line, and the current draw is read back in the seventh DAC line.

2.10.4 HV interface board

The HV interface board controls the HV power supplies at the direction of the HV DACs board. Additionally, the HVIB controls the power going to each of the eight Spellman HV supplies with a relay. The HVIB also carries the current draw, voltage setting, and voltage reading between the Spellman supplies and the HV DAC. Finally,

the HVIB directly provides the HV Controller board with the power state of the HV distribution system, so that any HV power failure can be seen by the control electronics.

Every Spellman MP3P24 power supply has an output range of 0 to 3kV, controllable by adjusting the voltage difference across the analog V_{prog} inputs between 0 and 10V. The HVIB combines the HV-coarse and HV-fine outputs from the HV DACs board by passing the DAC currents through either a 100Ω or a $25.5k\Omega$ resistor before passing the signal on to the MP3P24. This provides the analog V_{prog} with a weighted sum of the coarse to fine DACs of 255 to 1, respectively. The maximum voltage from the HV DACs board has been set to 7.5 Volts, effectively limiting the maximum HV to 2250V.

The HVIB has 8 D0061A relays that switch the power to the MP3P24 power supplies. The relays remain closed so long as the corresponding input BNC carries a positive TTL signal. The relay state is set by the HV Controller card relay outputs. Additionally, the HV Controller card monitors the power state of the board itself through a separate two pin connector. If the HVIB power status line ever turns off, the controller board immediately opens all of the relays (see section 2.10.2).

2.10.5 HV Distribution Panel

The high voltage supply lines run directly from the Spellman supplies to the back of the HV Distribution Panel (see Figure 2-25), on which the HV interface board is mounted. There are twelve female SHV ‘in’ connectors which are mounted on the back of the distribution panel (not shown), eight of which are expected to be in operation

during standard running. Each of the eight connectors attaches to an entire row of fourteen female SHV ‘out’ connectors on the front side of the HV distribution panel. The last four SHV ‘in’ connectors each hook up to one of the four half-rows of seven ‘out’ connectors at the bottom of the HV Distribution Panel.

A basic diagram of the HV Distribution Panel is shown as Figure 2-25, with SHV groups indicated by like colors and proximity.




Figure 2-25: HV Distribution Panel with 3 rows tied together

The four half-rows are to be used for test purposes. Either an external power supply or one of the primary power supplies can be routed in through one of the four SHV inputs on the backside of the HV distribution panel. If primary power supplies are diverted or fail, the extra two SHV outputs on the front side of the HV distribution panel

can be used to ‘jump’ the voltage across the primary rows. Note that Figure 2-25 shows the panel with the first 3 lines tied together.

2.10.6 Spellman MP3P24 HV Power Supply

The full deployment system has six Spellman MP3P24 HV Power Supplies⁵⁸. The Spellman MP3P24 power supply has a 10-pin interface in the configuration as shown in Table A-10.

2.11 Calibration Pulse Generation

2.11.1 Agilent 33120 Pulser

The electronics test pulses are created by an Agilent 33120A 15 MHz Arbitrary Waveform Generator. The Agilent 33120 connects into the first of the daisy-chained Pulser Distribution Cards (PDCs) detailed below. The control of the 33120 is done through isolated GPIB directly from ORCA. Although the standard test pulse (see section 2.3.3) is expected to be the standard output of this device, arbitrary waveforms can be stored and sent out through ORCA.

The output of the Agilent 33120 is adjustable to 12 bits of precision, with an accuracy of 1%. The output amplitude can vary from 50 mV p-p to 10 V p-p, and all signals pass through a 30-dB attenuator (see below). The output impedance is fixed at 50 Ω . The pulse rise-time of the Agilent 33120 is 18 ns.

2.11.2 Analog Characteristics of the Pulser Distribution System

In standard operation, the signal output from the Agilent 33120 is brought through a 30 dB attenuator before reaching the PDS. Inside the PDS, each signal is amplified by a factor of 7.22 through an AD8055 and distributed to individual channels. Just before exiting the PDS system, each signal is then brought back down by a 0.275X attenuator. The rise-time for the entire PDS is approximately 15 ns, so that in series with the Agilent 33120, the rise-time is around 24 ns. A diagram and more detailed description can be found in section 3.6, where analog measurements of the system are presented.

2.11.3 Digital Characteristics of the Pulser Distribution System

The NCD electronics pulser distribution system allows for an arbitrary waveform to be passed through to any subset of electronics channels. The NCD PDS consists of one Pulser Distribution Optocoupler Board and up to 8 Pulser Distribution Cards. Control of the PDS is exerted through an IP408 I/O module that connects to the PDOB through a 50-pin ribbon cable. The channels can be individually set in groups of 8, and up to 128 total channels connected on 8 separate Pulser Distribution Cards can be set to arbitrary on/off states. Each individual command can only set the on state of 8 channels, but all channels can have the pulser signal turned off with a 'reset' command.

The PDOB optically isolates all signals coming from the IP408 through 14 optical isolators. The PDOC reroutes all input lines from the IP408 through 20-pin ribbon cables which carry identical signals. The only return signal that the PDOB returns is on the 'Instruction Complete' line, which is first reverse-isolated, and then run across a parallel

series of jumpers back to the IP408 return lines. Effectively, this allows for the 'Instruction Complete' signal to return on any combination of the first 8 IP408 lines.

The PDCs have 1 input line, 1 direct output line, and 16 selectable lines. In standard operation, the pulser will connect into the input of one of the PDCs. The direct output line will carry the pulser signal to the next PDC's input line, and the final PDC will have a 50-ohm terminator placed on its direct output. The 16 selectable lines on each of the up to 8 PDCs are split into a lower and upper group of 8.

The PDC's only function is to connect or disconnect lines from receiving the pulser signal. On receiving an event strobe (2nd channel of PDOB-PDC cable), the PDC looks for an address on the 7th, 8th and 9th channel of the 20-pin PDOB-PDC cable. If the three bits match the PDC's address, then it checks the 6th address line to see which group of 8, high or low, the channel setting is intended for. The PDC then sets the 8 selected channels of the ALS575 8-channel switch to match the on/off state of the 8 data bits that are sitting on the 10th-17th bit of the PDOB-PDC cable. After the PDC sets the switches, it asserts an 'instruction completed' on the 19th channel of the PDOB-PDC cable. The PDC will continue to hold the 'instruction completed' line high until the PDOB lowers the 'Strobe' line, the expected response when the DAQ has acknowledged completion of the switch settings. On a 'Clear' command, that is, when the PDOB-PDC cable carries a high 4th bit, all PDCs will turn off all switches, regardless of any other information carried on the line. The successful completion of a 'clear' command will also return an 'Instruction completed' return when done. A list of the commands that can be sent and lines that return values is shown as Table A-11.

2.12 Low Voltage Power

The NCD electronics require +5, +12, +15, -6, and -12 Volts of DC power at different points. Each voltage is provided by a separate power supply with a redundant in-rack backup in case of individual failure. There is an additional +5 Volt isolated power supply to prevent ground loops within the multiplexer system. A full list of voltage requirements for each NCD electronics subsystem, along with the maximum current draw, is given in Table A-12.

2.12.1 VME Power

The standard VME crate switching power supplies were replaced with linear power supplies to enhance stability and decrease noise. Although slight, the switching supply noise generated during operation was unacceptable during operation of the Shaper/ADC boards. Although the addition of more power supplies allows for increased power draw, it also has much higher capacitance. When the VME crate is empty, the non-switching VME power supplies will hold voltage much longer, and must be power-cycled with care.

2.12.2 Low Voltage Power Distribution and Power Racks

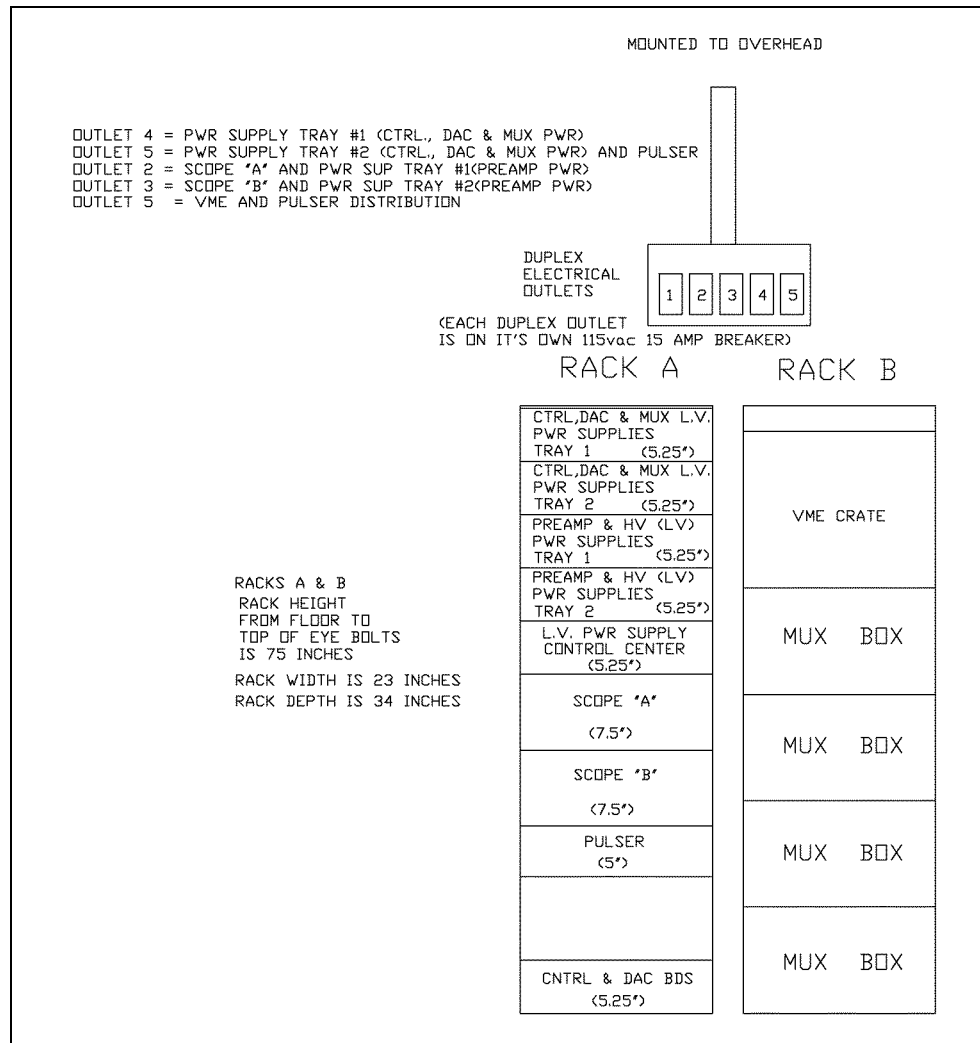


Figure 2-26: NCD Rack Layouts

During operation, each of the DC power supplies must draw AC power. Given the high power load of the different components, 5 separately-fused 15-amp 115 V-AC outlets were required. The layout for the two electronics racks used during cooldown is shown in Figure 2-26. Note the redundant trays of power supplies.

Chapter 3 NCD Electronics Verification and Control

The complexity of the Neutral Current Detector Electronics requires regular analog calibrations in order to assure data accuracy and consistency. These calibrations fall under three major categories: MUX log amp calibrations, MUX - Shaper threshold calibrations, and Gain calibrations. Electronics state and connectivity are verified by ORCA, the DAQ system detailed below.

Each of the NCD Electronics Calibrations (NCD ECA) are run at weekly intervals to validate data from the previous week and insure consistent data in the week following. The weekly interval is an optimization choice⁵⁹, balancing ECA dead-time with the risk of significant data loss in the event of a failure.

3.1 NCD Electronics Trigger Thresholds

Signal processing and digitization only takes place if discriminators are triggered by sufficient signal amplitudes. Each channel has 2 discriminators, one on the current and one on the current integrated over a 6- μ s shaping time. The setting, stability and efficiency characteristic of these thresholds must be measured.

3.1.1 NCD ADC/Shaper Trigger Thresholds

In order to monitor the actual threshold levels, NCD ECA runs include a varying-amplitude pulser run. The thresholds are mapped out as a function of pulser amplitude and shaper channel. The thresholds are determined by the pulser amplitude at which the Shaper/ADC acquires 50% of the test pulses.

The threshold test pulses are a series of 1.1- μ s one-cycle sine waves with amplitude varying from 125 mV to 25 mV in 1 mV increments. All electronics channels are sent the threshold test pulses simultaneously, with a frequency of 20 Hz.

The algorithm used to determine the threshold setting on each channel is based on that used for SNO-PMT electronics. In this algorithm, for each channel, the pulser is set to the lowest end of the calibration range. The algorithm, implemented automatically by the data acquisition system, records the number of triggers received by the electronics at each pulser amplitude, starting at the low end. Once the ratio of triggers generated in the electronics system to pulses sent into it goes above 50%, the bin number is saved as b_1 . The algorithm sets a second bin b_2 by starting at the largest pulse and decreasing the amplitude until the ratio of pulses to triggers drops below 50%. If the bins agree, that is, $b_1 = b_2 + 1$, the threshold is fixed.

If bins b_1 and b_2 differ by more than one, the precision is decreased by merging every 2 bins and rerunning the check. If this next check fails, further mergers are made until 10 bins are required. If the algorithm succeeds within 10 bins, the uncertainty in the

threshold is recorded to be the bin width where the condition $b_1 = b_2 + 1$ was first met.

After 10 bins, the algorithm determines that there is no solid threshold.

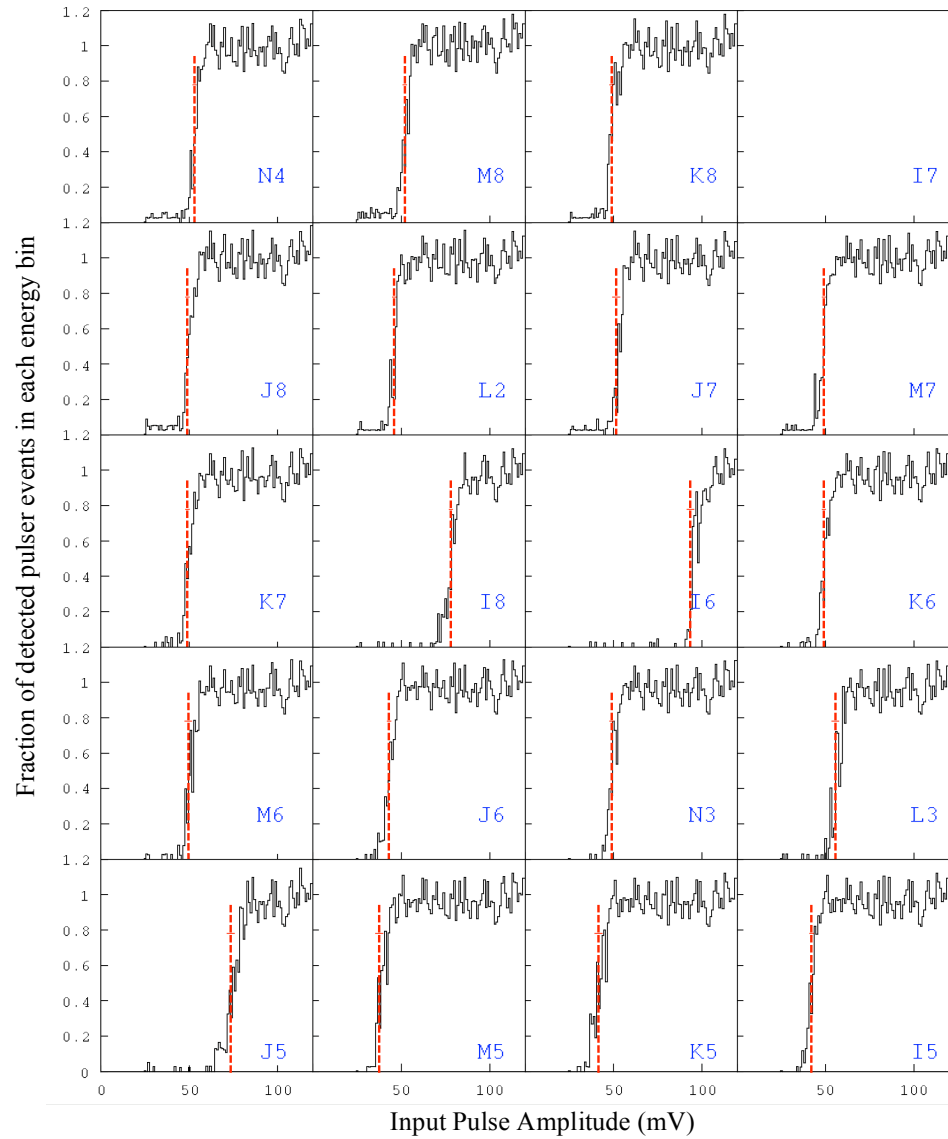


Figure 3-1: ECA NCD Calibration Run results. There are events below threshold which trigger the system due to background level variations. Uncertainties on thresholds determined by the rebinning method are 2 to 4 channels wide, and are indicated for each channel by a small horizontal dash.

An example of 20 channels in an NCD Shaper/ADC calibration run is shown in Figure 3-1. The data are from run 54323 and the plot is from [59]. Note that the

threshold on I7 is set above the highest (125 mV) level. Each plot shows the ADC counts as a function of pulser amplitude, and the dotted line shown is the threshold determined by the algorithm presented above. The counts that appear below the determined threshold arise when test pulses combine with noise to trigger the system. Essentially, the width of the threshold is a determination of the noise on that channel.

3.1.2 NCD Mux Trigger Thresholds

The trigger threshold for the NCD Multiplexer system is determined with a method similar to the Shaper/ADCs, above. Pulses sent in from the Pulser Distribution System scan the amplitude range from 25 mV to 125 mV.

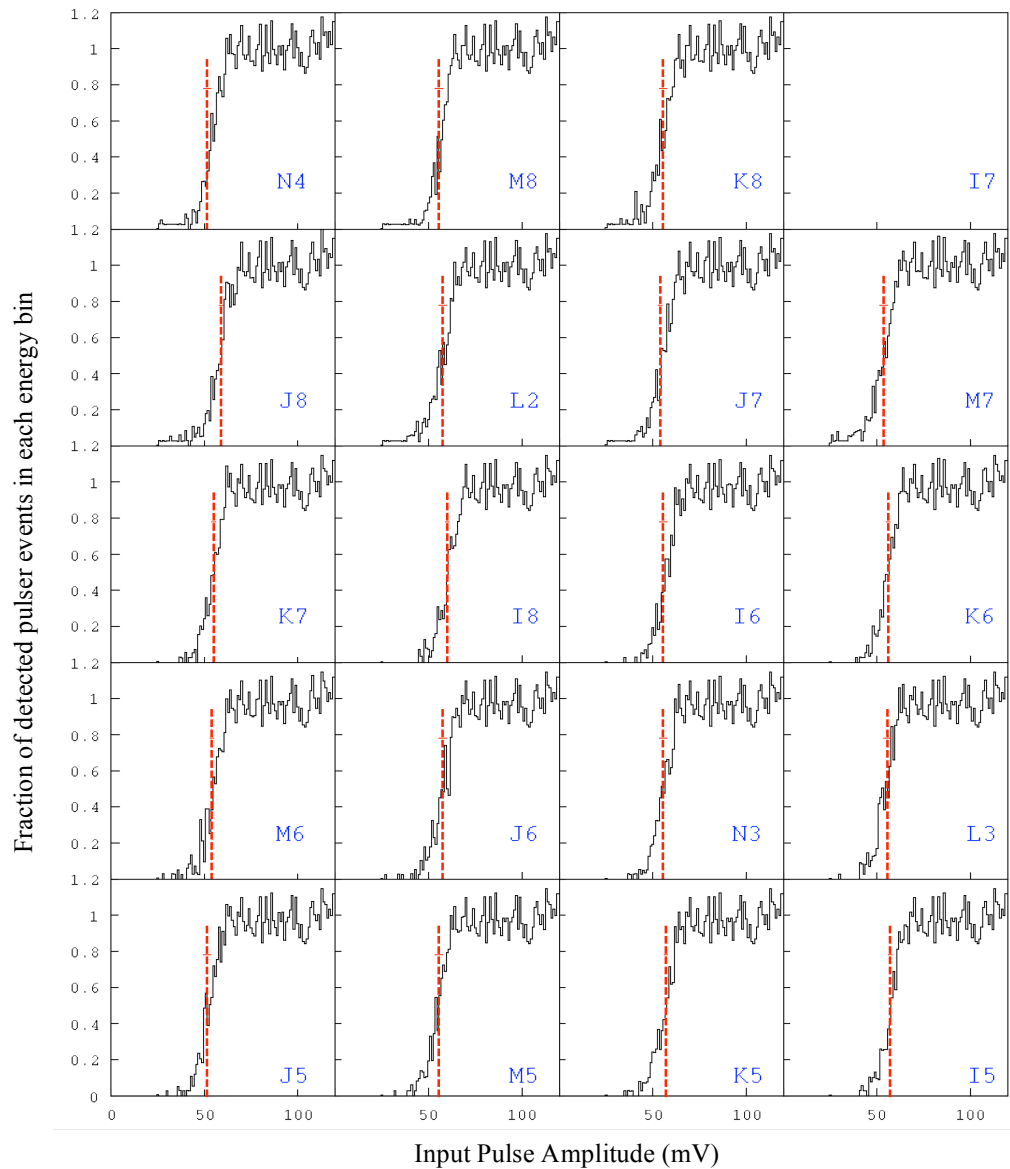


Figure 3-2: ECA MUX threshold calibration run across 20 channels. Uncertainties on thresholds determined by the rebinning method are 2 to 4 channels wide, and are represented as small horizontal lines near the 0.8 fractional level.

An example of threshold calibration for the MUX thresholds is given as Figure 3-2. Note the reduced accuracy on triggers when compared to the Shaper/ADC system, an expected difference due to the lack of current integration, which increases the bandwidth and noise at threshold.

3.2 NCD Electronics Linearity

3.2.1 NCD Gain and Linearity Calibrations

Each NCD Electronics Calibration includes a specific linearity run. In the NCD ECA linearity run, the Pulse Distribution System routes a square-wave pulse into each preamplifier one at a time. The standard linearity pulses are 10 1.1- μ s square waves with varying amplitude from 625 mV down to 145 mV in 5 120 mV steps. A once-a-month extended NCD ECA includes 20 steps from 1001 mV to 51 mV of 50 mV each. Pulses for Shaper/ADC linearity tests during every NCD ECA run are sent out at 50 Hz, while for log amp calibrations, pulses are sent out at 0.8 Hz.

Due to the extended duration of the linearity tests, real physics events are expected to arrive during such an ECA run. At this time, these events must be manually extracted, but a DAMN mask (SNO data tag) can be set up to remove non-ECA events.

A typical Shaper/ADC linearity test is shown in Figure 3-3.

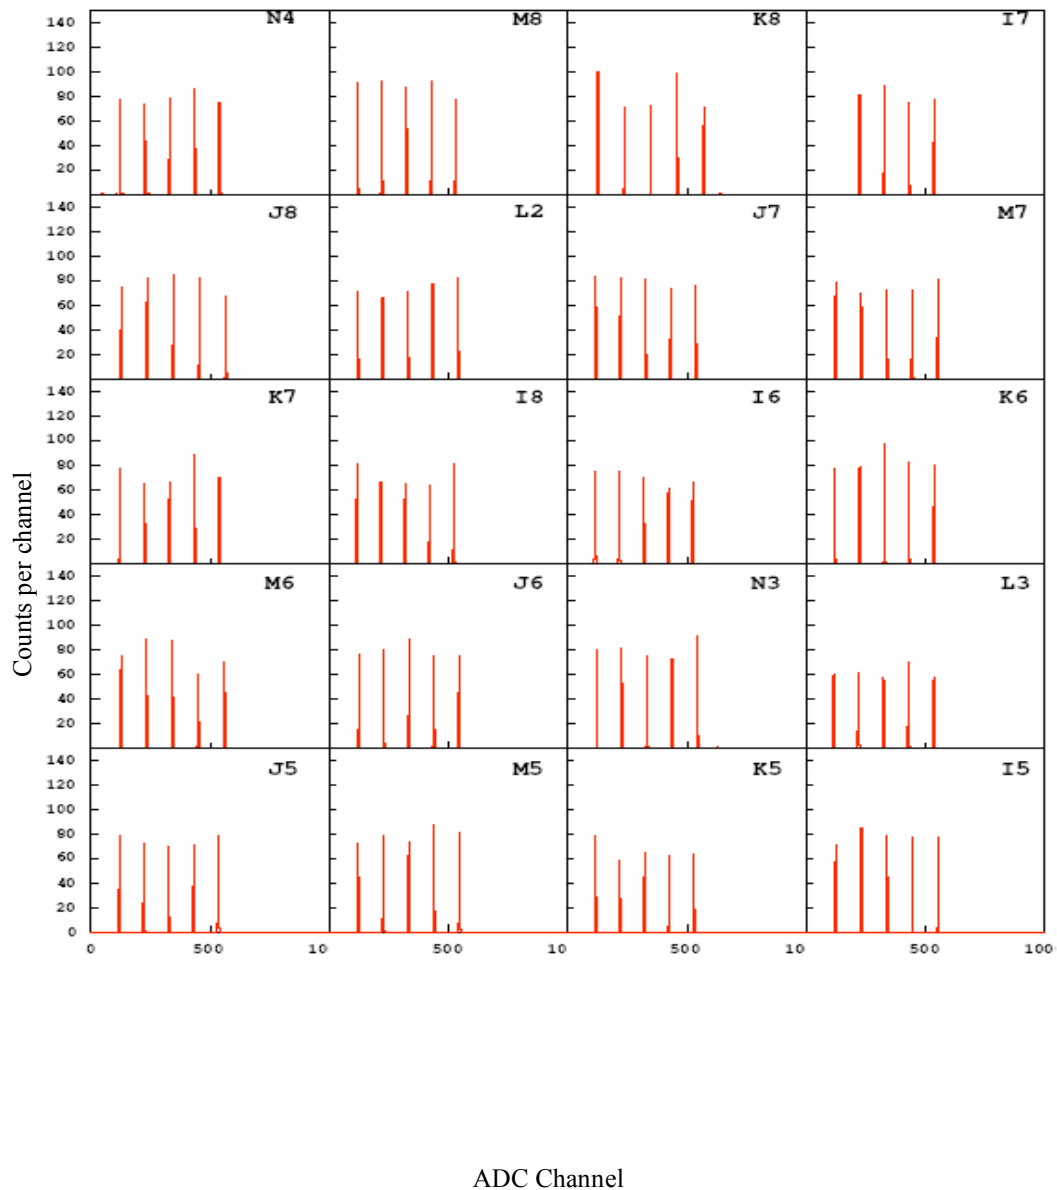


Figure 3-3: Shaper Channels Occupied During Typical ECA Linearity Test (Run 54745)

These plots also show that the electronic pulser resolution is very narrow compared to the resolutions expected for the NCD physics events. The linearity of the Shaper/ADCs can be more easily seen by examining the deviation from linearity as a

function of test pulse amplitude. The results for these 16 channels are shown in Figure 3-4.




Figure 3-4: Shaper/ADC Deviations from Linearity During ECA-- Note that the maximum deviation is less than 0.5%

The ECA Linearity tests for the MUX/Scopes require a valid log amplifier calibration to have been obtained beforehand. The test pulses are “delogged” (see section

3.2.2) and the peak amplitude is compared to a linear function. The resultant deviations for pulses input on the same 16 strings are shown as Figure 3-5.

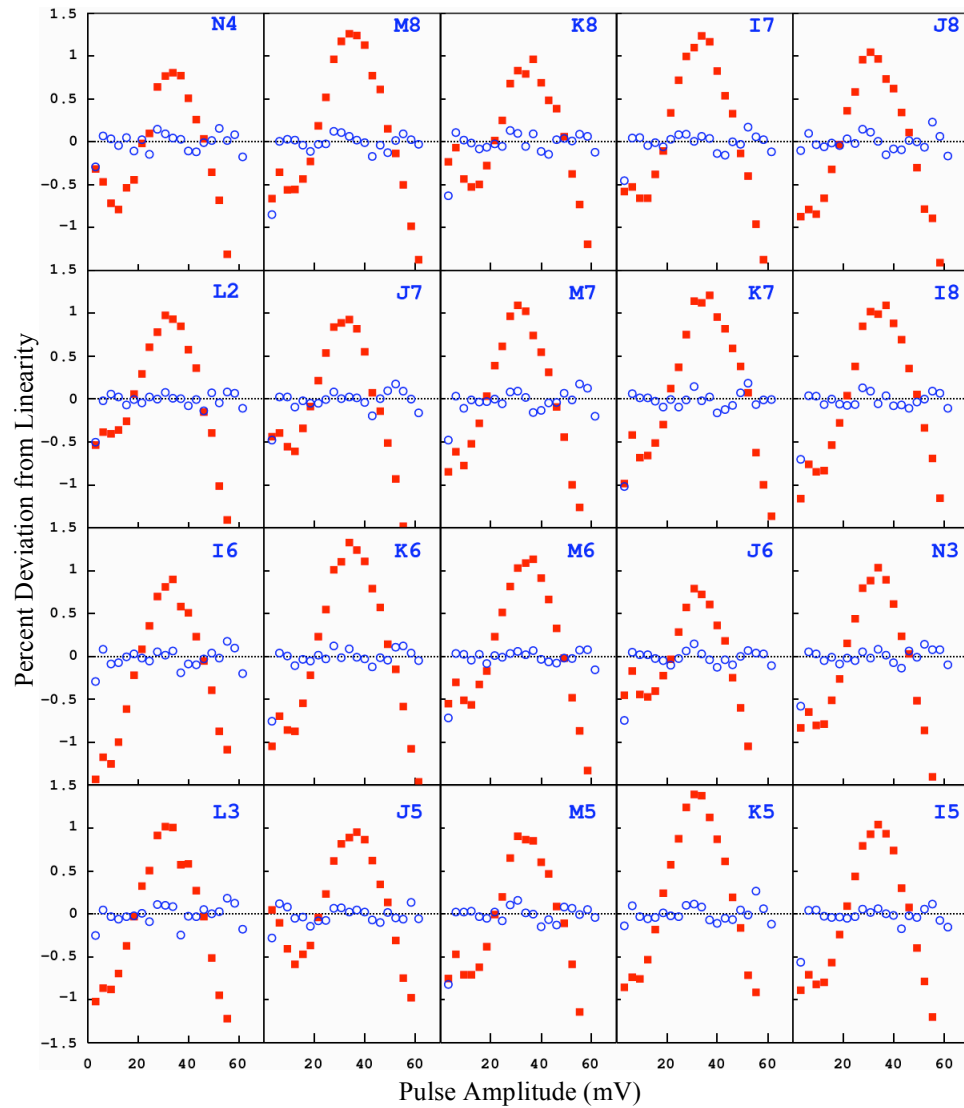


Figure 3-5: MUX/Scope Deviations from Linearity for an extended ECA run. These deviations, while they could be corrected numerically, are less than the uncertainties from other sources.

3.2.2 NCD MUX Log Amp Calibrations

The dynamic response of each log amp in the NCD Multiplexers is determined by injecting a well-quantified pulse into the system. The oscilloscope record is then analyzed and a response function generated.

The standard calibration pulse for the log amplifier is 1 offset complete cycle of a sine-wave pulse, beginning and ending where the derivative is zero, with a square wave trigger pulse a few microseconds ahead of the pulse. The logarithmic transfer function is gleaned by analyzing the log output of the known sinusoidal waveform. The square wave pulse provides consistent triggering of the system, while the sine wave provides a continuous function with minimal high-frequency content for logamp calibration. The ideal standard calibration pulse, on input to the preamplifier, is shown as Figure 3-6.

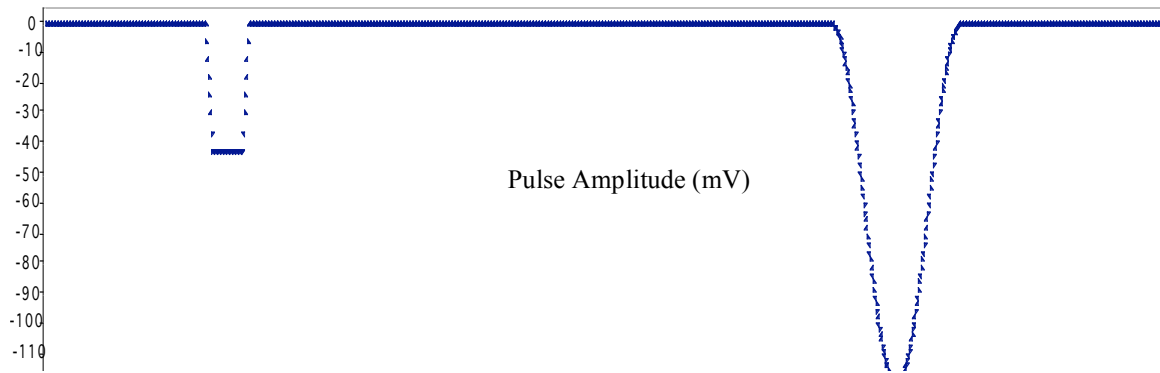


Figure 3-6: Theoretical Standard NCD Multiplexer Calibration Pulse, mV vs ns.

An example of this calibration pulse sequence after passing through the MUX and log amp is shown in Figure 3-7.




Figure 3-7: Digitized Test Pulse after MUX Log Amplifier

This also shows some of the effects of pulse reflection at the input to the preamplifier, most clearly visible on the square part of the trigger, which shows multiple peaks. Mainly this structure is caused by the pulse reflecting from the 325- Ω resistive coupler between the 93- Ω cable and the 415- Ω NCD. The coupler provides matching impedance for pulses traveling from the NCDs to the preamp, but not in the other direction. The influence of these reflections is included in the ECA analysis of log amp calibrations. Further examination of this pulse reflection can be found in [59]. It is important to note that this analysis has also uncovered non-logarithmic behavior from the MUX Logamps.

3.2.3 Log Amp Non-linearity

The original specification for the logarithmic amplifier called for the AD606, and the preparatory work for the MUX system included its measured transfer function. However, the AD8307, a faster logarithmic amplifier with a wider dynamic range, became available before the MUX hardware was built. Consequently, the production MUX system includes the AD8307 amplifier.

In section 2.7 it was noted that the signal headed into the MUX is amplified times 2, source terminated, and sent through a 320-ns delay line. There is approximately unity gain from the preamplifier to the log amp.

The original specifications for the system called for offsets up to 20 mV. Given that the AD8307's offset setting range is only 0 to 1 mV, a 26 dB attenuator (1:20) was placed in front of the log amp. Since the signals could be expected to vary from 20 mV to -2.5 V, the dynamic range at the input to the log amp is 1 mV to -125 mV. In retrospect, this was not an optimal design choice because input noise at the log amp contributes to the noise budget, albeit at a minor level.

The current output of the AD8307 varies by 40 μ A per 20 dB of voltage at the input. In typical operation, an internal 12.5 k Ω resistor converts the current to a voltage. Unfortunately, the capacitance associated with this output channel limits the minimum rise time to over 400 ns. To reduce this time, the log amp output is connected to the summing junction of an inverting op-amp. This drastically reduces the impedance, and after passing through a second op-amp, provides a nominal voltage gain of 140 mV per

decade change at the input. The actual gain measured is approximately 163 mV per decade. The expected signal's dynamic range is a factor of 150, or 2.176 decades, which therefore corresponds to 355 mV range at the scope input.

Fitting the response function of the output from the log amp is done⁶⁰ using the equation $V_{out} = K \ln D(V_{offset} + V_{in})$, where K and D are constraints. Figure 3-8 shows the fitted function at known offsets and input voltages plotted on top of the measured data. The negative offsets and input voltages are shown in Figure 3-8 as positive voltages.




Figure 3-8: Log amp response at different offsets, fitted by expected response function. Plot is from [60]

The 2-3% deviations from linearity that the log amplifier exhibits are actually visibly non-random. A residual plot, given for an offset of 10 mV with many input test values,

clearly demonstrates this in Figure 3-9. This behavior is a manifestation of the particular architecture of the AD8307 chip, a cascade of saturating amplifiers.




Figure 3-9: Log amp linearity fit residuals at 10 mV offset. Residuals are shown in mV deviation from log-function plus offset best fit. Plot is from [60]

Although this non-linearity can be fit to a combined sine and natural log function⁶⁰, the resultant transform is not easily reversible. Should the introduced non-linearity prove an analysis problem, this technique may be investigated further, but at present the non-logarithmic behavior does not appear to compromise the analysis.

3.2.4 NCD ADC/Shapers and Ballistic Deficit

Due to the variable event durations within the NCDs, the ballistic signal is not ideal—a fraction of the integrated pulse decays during the acquisition. The amount of signal that is lost relative to an ideal integrator is known as ballistic deficit. This total

deficit depends on the amplitude and duration of each pulse, so we attempt to correlate the fractional ballistic deficit directly to the pulse duration for each total integrated charge. For small integrated charges for small durations, as would be expected for neutron signatures, the ballistic deficit may not be seen on individual channels. However, after averaging data and correcting for pulser round-off error, we can see the $\sim 2\%$ dropoff over $2.5 \mu\text{s}$, as can be seen in Figure 3-10. The dimensions of the counter and the drift speed in the gas limit the maximum pulse duration to about $2.5 \mu\text{s}$.




Figure 3-10: Short small pulse ballistic deficit as a function of pulse duration, for rectangular pulses with a fixed total charge. This expected deviation will not be corrected for as it is not the dominant error term and cannot be removed from data without digitization records.

3.3 NCD Electronics Deadtime Measurements

Since the physical objectives of SNO are the determination of rates (solar neutrinos, supernova neutrinos, atmospheric neutrinos, etc), a detailed understanding and measurement is essential. For each of the data acquisition paths, we investigate the experimentally determined deadtimes to compare with the predicted deadtimes in Chapter 2.

3.3.1 NCD Shaper/ADC Scaler Deadtime

The scalers in the NCD Shaper/ADC system have the shortest deadtime in the NCD electronics. The actual trigger deadtimes are difficult to show in graphical format, as the scalers do not have times associated with their triggers. We will instead present the result of measurements⁶¹ that place the scaler deadtime at 20 μ s, about 3.5x the shaping time of 6 μ s.

3.3.2 NCD Shaper/ADC Deadtime

The Shaper/ADC main trigger path has a lockout time during which no board may contribute a signal, and thus, by design, a fixed deadtime independent of string number. This deadtime was predicted to be in the hundreds of microseconds. Looking at the time separation between each subsequent event in the open data set in Figure 3-11, we can see that the expected rise in trigger rate is cut short by the measured NCD system deadtime of

$225 \pm 20 \mu\text{s}$. The uncertainty was determined using the same rebinning algorithm described for the threshold analysis in section 3.1.1.




Figure 3-11: Separation, in seconds, of consecutive NCD Shaper events

3.3.3 NCD MUX Trigger Deadtime

The MUX trigger system does not have an enforced lockout time like that of the Shaper/ADC system. Multiple hits are permitted during single events, but the GTID and clock information are assigned by the first hit. Thus, a deadtime exists that is primarily limited by computer response speed through the VME and IP408. This independent system shows the MUX system response is variable, with a deadtime of $780 \pm 80 \mu\text{s}$. This extended deadtime is undoubtedly due to the DAQ computer response time combined

with the IP408 communication speed. A plot of event time separation is shown as Figure 3-12.




Figure 3-12: Separation of consecutive MUX events

3.3.4 NCD Digitization Deadtime

The NCD Digitizers present a little more difficulty in quantifying deadtime. Each scope can acquire $20\ \mu\text{s}$ of data once triggered, but the slow event download time through GPIB limits the retrigger rate of each scope to less than 2 Hz. The difficulty in deadtime analysis comes with the multiplicity of oscilloscopes; while one scope is triggered, the other is still live, waiting for an event. Thus, the lower limit on deadtime between digitized events is set by the NCD MUX minimum retrigger time. The upper limit is

nearly the single-digitizer event retrigger rate, for if the two scopes fire in rapid succession, one has to wait for the full event download time before the third event can be covered. The plot of digitized event separation times is shown as Figure 3-13. Note the lowest non-zero time separation is in the 700-800 μ s range, consistent with the MUX trigger response, but the mean scope data rates must be held to 1 Hz or less to avoid serious losses.




Figure 3-13: Separation between consecutive digitized events on separate scopes. Single scope retrigger times are much larger due to the GPIB data transfer.

3.4 NCD Livetime Counters

The 4 Livetime Counters, described in detail in section 2.9.2, count the total number of 1 MHz clock pulses for which each electronics system is able to acquire data.

The Shaper/ADC, MUX, and Digitizer livetime counters are compared with the total system livetime counter, giving a fractional livetime for each system during the run. The livetime counters are not reset at the beginning of each run, and overflow is expected every 12.73 days.

The livetime read by the data is related to the expected deadtime on each system, so long as the events are distributed in a predictable fashion.

3.5 Random Pulser

We expect solar neutrinos, and thus neutrons liberated by solar neutrinos, to arrive as a randomly distributed pulse train. In principle, the livetime counters provide an accurate livetime for event rate measurement. However, when cuts are made to remove bursts, the additional deadtime bears an unknown relationship to that registered by the livetime counters, since they are read out only occasionally and not event-by-event. A better measure of deadtime can be obtained by running a pulser whose random time distribution mimics that expected from the signal⁶¹.

The sampling of deadtime provided by the random pulser is only valid for random pulse trains. Signals with time-correlated events would not have their sampling fraction accurately determined by the sampling fraction of random pulser events. Random pulsers can be monitored and their events recorded in the scopes, the Shaper/ADCs, and the Shaper/ADC scalars. The differences can provide an accurate measure of the data loss for random signal event trains.

Let the total deadtime predicted by the sampling fraction of the random pulser be τ_d . Given that we will know (from the scaler) the true number of pulser events N_p , we will be able to measure the difference, or lost pulses, L_p , between the pulses sent and acquired during any interval. The most probable total deadtime as a fraction of the total run time T is the number of pulses missed divided by the number generated, with a Poisson-distributed uncertainty on the number of events lost:

$$\frac{\tau_d}{T} = \frac{L_p}{N_p} \left(1 \pm \frac{1}{\sqrt{L_p}} \right)$$

In any given run, the number of detected versus true events will be reduced by a factor of the livetime divided by the total run time. If the true event count is C_{true} , and the number of observed events is C_{obs} , then:

$$\frac{T - \tau_d}{T} C_{true} = C_{obs}$$

The true signal count should also be a simple factor of the true signal rate, R_t and the run time T :

$$C_{true} = R_t T$$

The uncertainty on the true signal count, however, depends on the uncertainty in the poisson-distributed signal itself as well as the uncertainty in the number of events lost to deadtime. As the former depends only on the poisson distribution of the true signal,

and the latter depends only on the uncertainty in deadtime that results from the random pulser, these two errors are uncorrelated. Indeed, we may add them in quadrature:

$$C_{true} = C_{obs} \left[\frac{T}{T \left(1 - \frac{L_p}{N_p} \right)} \right] \left(1 \pm \frac{1}{\sqrt{C_{obs}}} \right) \left(1 \pm \frac{L_p}{(N_p - L_p) L_p} \right)$$

It is evident from examination of these relationships that there are competing terms for the optimum random pulser rate. On one hand, too many random pulser pulses interfere with the signal sampling. On the other hand, too few pulses and the system deadtime is not well understood. The choice of optimum random pulser rate, set as a Poisson distribution about 0.005 Hz, is detailed in [61].

3.6 NCD Electronics Model

The net analog effects on an NCD-generated signal must be understood for the signal source to be recognized. In the effort to represent the global effect on both real and artificial pulses, a unified NCD electronics model⁶² has been developed and is shown as Figure 3-14.

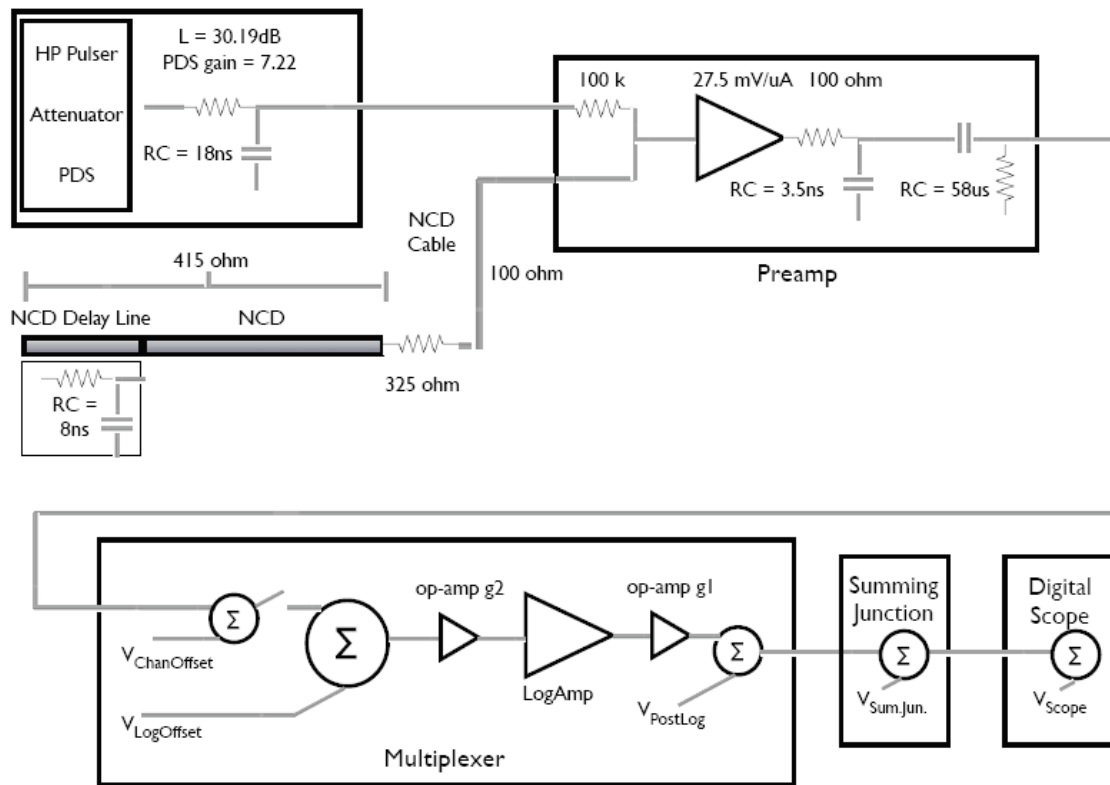


Figure 3-14: Simplified NCD Electronics Model for Digitized Pulses. Note that a low-pass filter exists between the Multiplexer LogAmp and the op-amp g1, but is not placed on this diagram.

3.7 NCD DAQ: ORCA

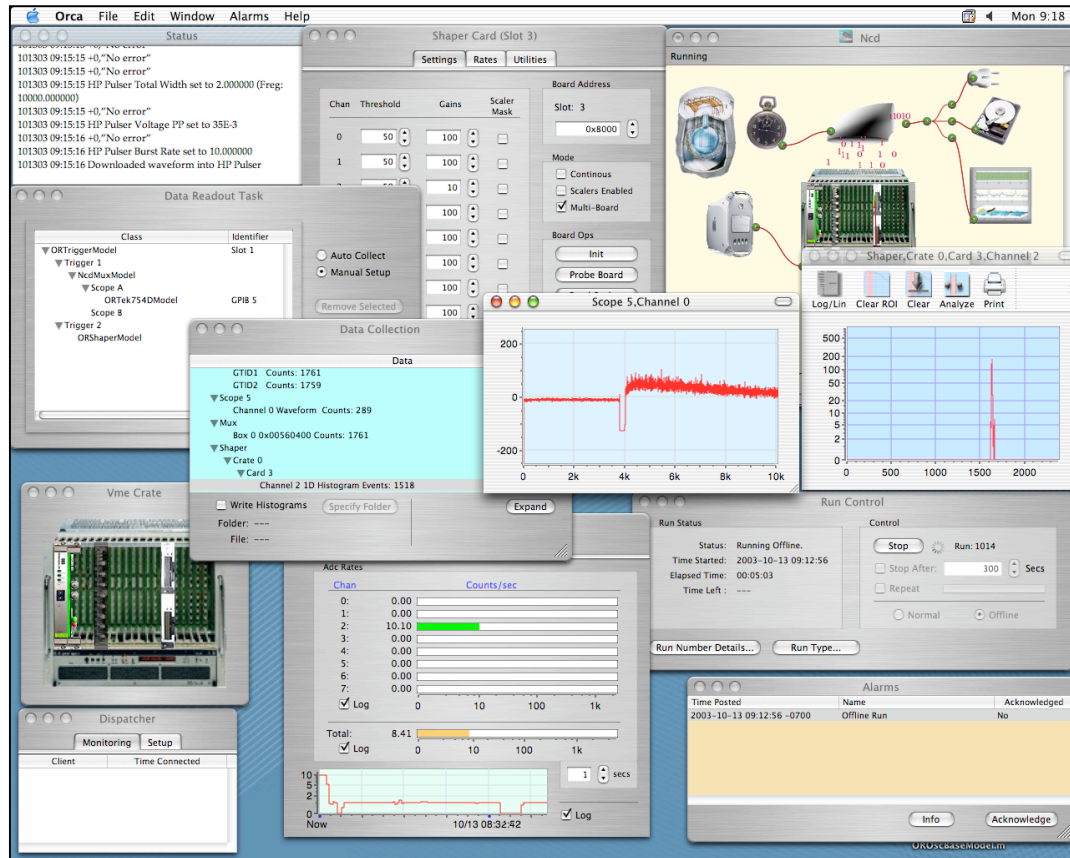


Figure 3-15: ORCA during operation. Window in top left is status, which has written information on system operations and errors. In the top middle is a single Shaper/ADC window, with its control tab highlighted. In the top right is the system configuration screen, described in more detail below. Across the middle are windows showing data readout and collection information, followed by a single scope trace and a Shaper/ADC spectrum. The second-from-the-bottom on the left is a VME configuration window. The window on the bottom left shows the dispatcher client statuses (see Appendix B). The middle-bottom window shows current and past rates through a horizontal meter and strip chart. The run control window is second-from-the-bottom on the right, where system-wide runs are configured, started, and stopped. The window on the bottom right displays current system errors.

The NCD electronics are monitored and controlled by ORCA, the Object-oriented Real-time Control and Acquisition system running on MacOS-X operating systems⁴⁹. In typical NCD DAQ operation, ORCA has an interface screen as shown in Figure 3-15. ORCA was first implemented with the SNO NCDs, but it was designed for much more

general usage. ORCA was written with the MacOS-X Cocoa application framework development environment in Objective C++.

ORCA's architecture stems from the Model-View-Controller designs common with object-oriented languages. With ORCA, the view is the human interface to the piece of hardware being controlled. Views can display output both on the screen and on speakers while allowing input from the mouse and keyboard. The actual software modules for each component are the models. Each model is an independent object-oriented software module that represents a hardware object, data analysis package, control module, or data acquisition task. Every model interfaces with each view available to it through a controller. The controllers may pass DAQ data to the views associated with it as well as passing user information given through a particular view back to the linked models.

3.7.1 DAQ Control

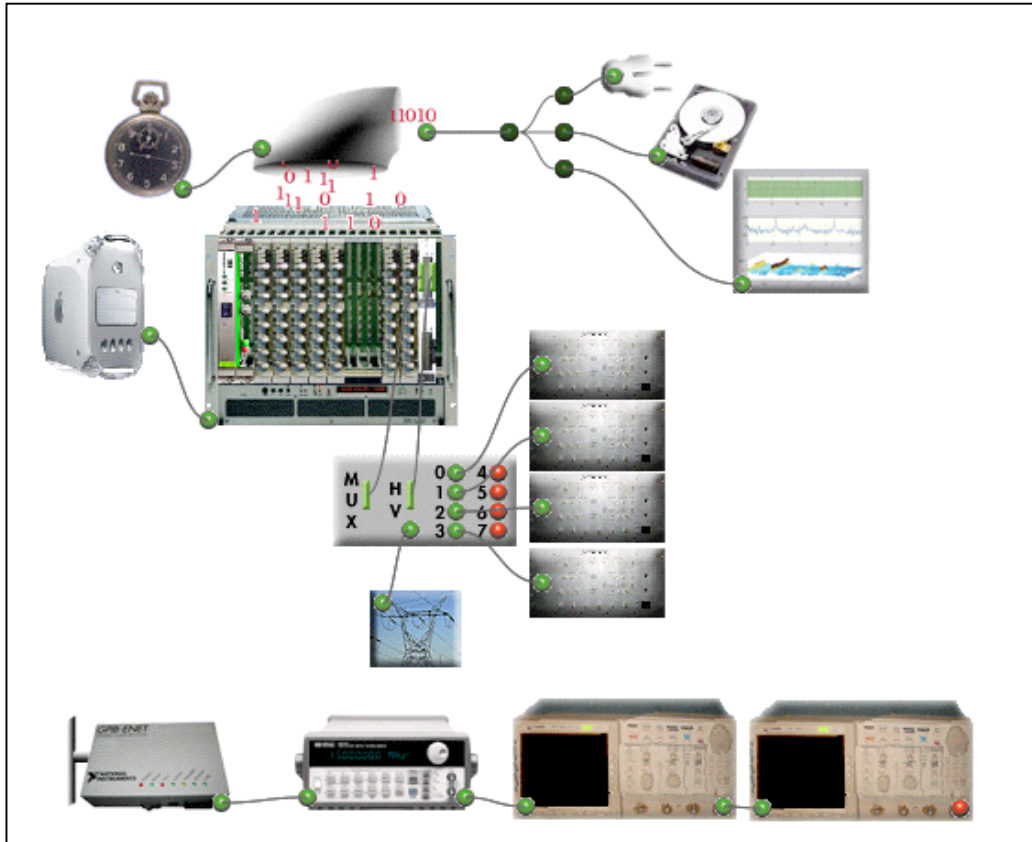


Figure 3-16:ORCA data acquisition system iconic view. The stopwatch represents run control and is connected to the data read-out object. The data read-out object collects data from the shared VME memory and routes it out to a dispatcher server, disk, and control screen. The Mac is set up to control the VME crate. The VME crate has a half dozen shaper/ADC cards, an NCD trigger card, and two IP408 modules in a single carrier board. The IP408s, in turn, are connected to the MUX and HV Control boards below them. There are 4 MUX crates shown connected to the right, and one HV system. Separately, across the bottom, the GPIB bus shows connectivity between the GPIB controller, the pulser distribution system, and the two digital scopes.

ORCA, when used for NCD DAQ, has a logical configuration screen similar to what is shown in Figure 3-16. Each model in the configuration screen has a graphical representation, usually clearly identifying its type and function. Setting up logical data connections between models is as simple as a click-and-drag operation. Opening up the

views for each device may be accomplished by double-clicking on any individual element. Introducing new hardware or control elements is done by opening up a view of the catalog, which contains an index of available modules.

3.7.2 Data Construction

SNO's main data stream is in CERNLlib ZDAB format, so all data leaving the NCD DAQ system is converted before it is merged. A block diagram of the data construction process is shown in Figure 3-17.

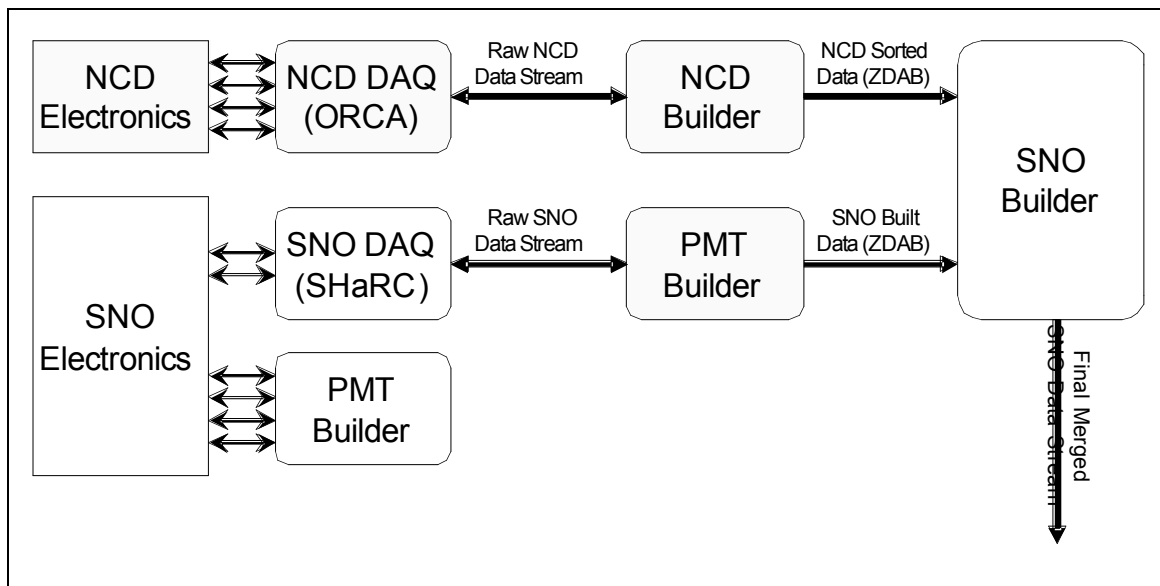


Figure 3-17: Building and merging SNO and NCD data streams

Each first level builder, shown in separately as NCD Builder and SNO Builder, operates independent of the other. This allows for single stream operation during maintenance of either system. The second level builder, SNO Builder, combines the built ZDAB data from both the NCD and PMT arrays, outputting a single data stream.

Chapter 4 Neutrons and NCDs

The detection of neutrons is the primary function of the Neutral Current Detectors. Aside from muon-following neutron events, most free neutrons within the D₂O have been liberated by neutrinos through neutral-current interactions. A smaller fraction of neutrons are liberated by high-energy gamma rays and alpha particles.

Once neutrons are liberated, the D₂O provides an ideal moderator to thermalize them with minimal absorption probability. The ³He-filled NCDs have a high thermal neutron absorption cross-section, making them effectively opaque to the neutrons meandering about through the D₂O.

The most frequent liberator of neutrons within SNO is solar neutrinos, a result of SNO's cleanliness. The ability to identify neutrons liberated by the neutral-current interaction is what makes SNO uniquely sensitive to the total active neutrino flux. The determination of this neutral current flux is the strength and the core of the experiment; therefore it is vitally important that all of the factors that stand between the liberation of a neutron and the identification of the same be well quantified.

4.1 Other neutron sources in SNO

The Sun is not the only source of neutron-liberating neutrinos—supernovae are expected to liberate many neutrons over a short period. For galactic-core supernovae, a hundred or more neutrons may be expected within a 10-second interval. The strength of SNO over other supernova neutrino detectors is its ability to record the total active neutrino flux, which again relies on SNO's neutral-current sensitivity.

Neutron flux within SNO's D₂O is not just a function of neutrino flux. The deuteron can be broken apart by 2.2 MeV of energy, and neutrons can be separated from other nuclei as well. The primary sources of backgrounds within the D₂O are the 2.6-MeV gamma ray in the ²³²Th decay chain, the 2.44-MeV gamma in the ²³⁸U chain, and (α,n) reactions from both the ²³⁸U and ²³²Th chains.

The rest of the mine has many neutrons passing through it due to ²³²Th and ²³⁸U decay chain (α,n) reactions, spontaneous U fission, and cosmic-ray muon interactions in the rock walls. SNO's D₂O is shielded by the light water that surrounds the vessel; less than one wall-contaminant neutron per year makes it into the fiducial volume²¹. In cooldown (see section 1.7.2, above), however, the light water was not present to absorb the neutrons, rather, the NCDs were directly exposed to the wall. The norite walls of the cavity provided a constant neutron signal for the NCD array to test its functionality.

4.2 Neutrons from supernovae during cooldown

Norite rock is an igneous silicate, made of mostly SiO₂, a rather poor material in terms of neutrino-neutron cross section when viewed as a potential supernova-neutrino-to-neutron converter. However, norite contains non-negligible amounts of aluminum, calcium, iron, sodium, and magnesium. There is a non-negligible amount of ²³²Th and ²³⁸U contamination within norite as well; early measurements⁶³ found 5.3 ppm ²³²Th and 1.3 ppm ²³⁸U.

Convoluting the energy-dependent neutrino-neutron cross sections of the major constituents of norite with the expected neutrino energy spectra in core-collapse events yields a net neutron production rate of roughly 8×10^{-6} n/g for a distance of 1 kpc.

The total mass of norite, to which the cool-down array is effectively sensitive:

$$M_s = \frac{r(L \cdot \varepsilon_L)}{(l \cdot \varepsilon_l)R_{(\alpha,n)}}$$

Where r was the rate of neutron detection with the first-phase cool-down array, l was the length of the first-phase array, and L is the length of the complete array. ε_L is the geometric efficiency of the first-phase cool-down array, and ε_l is the geometric efficiency of the final cooldown NCD array. Geometric efficiencies were calculated assuming a homogeneous and isotropic distribution of rays on the outside of the counter array, and assuming that NCDs were fully opaque to incident neutrons. $R_{(\alpha,n)}$ is the known rate of (α ,n) creation within the norite walls of SNO, deduced from the composition of the norite and the various (\square ,n) cross sections.

Calculating a potential total effective visible norite mass from the early (1998) cooldown data following the above prescription:

$$M_s = \frac{7000d^{-1}(775m \times 0.40)}{(16.5m \times 0.60)2.8g^{-1}yr^{-1}},$$

we found that M_s was approximately 3×10^7 grams with the early cooldown data. This implies that only ~ 4 neutrons in 10 second would be detected for a galactic-center supernova at 8 kpc. Unfortunately, this is not much higher than the (α, n) background rate due to norite itself. The predicted signal-to-background depends on the supernova model that is used; a one-second burst-model supernova would provide a ratio 7:3, while the data from 1987a seems to favor a ten-second burst, and thus a ratio of 34:30. In either case, the chance of triggering on most galactic supernova events is negligible without additional signal-to-background improvements.

Additional cooldown data were taken from December 23rd, 2002 to January 6th, 2003 and produced a somewhat different result. With 51 channels of electronics running a total of 409.5 meters, the entire array saw an average of 49508 neutrons per day. The calculated geometric array efficiency, given the configuration of active strings as shown in Figure 4-1, is 55%.




Figure 4-1: View of NCD configuration at the end of the counters in-rack for neutron measurement. X's are ^4He counters, and hollow circles represent unmonitored tubes.

Rates per counter varied between 163 and 91 neutrons per counter-meter per day. The lowest rate of 516 neutron counts per day was from a string in the middle of the array. The highest rate of 915 neutron-window events per day is from a string sticking out of the bottom on a rack by itself. This implies that individual string efficiency varied from 35% to 70% depending on position in the array, consistent with geometric and single-counter expectations.

For the array, 1350 events per square meter per day were detected as neutrons. However, only ADC signals between channels 140 and 165 were counted as neutrons. This corresponds to an approximate energy window from 700 MeV to 825 MeV. By examining counter spectra, the events captured in this energy window are only representative of $65 \pm 2\%$ of the total identifiable neutrons. Most of the remaining 35% fall below the proton-triton peak and this window. Adjusting for the 65% recognition efficiency yields an expectation of 2070 ± 40 neutrons per square meter per day resulting in a neutron capture in the NCDs. The alpha contamination in the neutron window is small and counter-dependent, but given expectations of roughly 8 alpha events per day, we can expect less than a 1% effect.

There is also the efficiency of each NCD for neutron interaction, given an incident neutron. As the mean free path of a thermalized neutron within the inner gas is about 2.5 centimeters, there is a non-negligible fraction that will exit through the counter without being detected. Numerically integrating the depth function,

$$\frac{2}{\pi} \int_0^{\pi} \left[R^2 + r_0^2 - 2Rr_0 \cos(\theta) \right]^{\frac{1}{2}} d\theta,$$

for an even distribution across all incident solid angles, we find that the mean alpha-parent distance from anode wire for detection within the NCDs is slightly over 1.2 cm. The fraction of incident thermal neutrons that range out of the detector is 30%. Adjusting accordingly, the thermal neutron flux within the SNO control room is 2700 ± 150 neutrons per meter per day.

However, we must also account for neutron absorption due to non- ^3He in the area. In an early SNO internal document⁶³, a 50% non-detector absorption number is recommended, but in that measurement a substantial amount of lead surrounded a moderator around a BF_3 counter. If we make the assumption that neutrons are isotropic and homogenous throughout the volume of space defined by the NCD cooldown rack and air between the rack and the near wall, we can estimate the absorption of neutrons.

Given 416 kg of natural-isotopic-abundantly distributed nickel, with a cross-section of 4.46 barns per atom, we find a 15.3% chance of interaction due to collision with the nickel counter bodies.

The 95 kg of iron in the steel Unistrut that makes up the support structure for the NCD cooldown rack will also absorb neutrons. With an average cross-section of 2.57 barns per atom, homogeneous thermal neutrons have a 3.8% chance of interacting inside the Unistrut mass.

Additionally, there is a mass of 18.3 kg of nitrogen in the 1.3 atm air surrounding the counters. Naturally abundant nitrogen has a cross-section of only 1.8 barns, so the chance of a collision with any of the nitrogen in the surrounding air is only 0.2%. There are also about 50 kg of mixed supplies beneath the rack. We assume that this contributes another 1% to the pre-NCD absorption probability.

In total, for absorption on non-sensitive array elements in the cooldown setup, we adjust our 2700 by 20%. This implies that the total flux of neutrons inside the control room is around 3250 ± 200 neutrons/meter²/day. This implies that the total array efficiency for capturing neutrons in the cooldown setup is 0.23 ± 0.02 for neutrons that are thermalized within the NCD racks. Despite the increased counter density, this is lower than D₂O as nearly half of the ³He was in tubes that were not powered.

There are a few effects due to non-thermalized neutrons that must be taken into consideration.

First, the efficiency was calculated under the assumption that neutrons could only travel straight towards the NCD counter bodies. Although the NCDs are mostly opaque to thermalized neutrons, fast neutrons can travel straight through them. The control room may expect to have more neutrons slowed down and captured when there is a significant body of metal and nickel surrounding the area.

Secondly, there was a “water wall” installed⁶⁴ into the control room to moderate and absorb fast neutrons for tests on a subset of counters. It is important to note that the largest number of neutrons per meter was measured by the counters within the water

wall. This indicates that thermalization within the water wall was a significant effect. Additionally, those counters that extended down towards the water wall had higher numbers of neutrons.

Finally, neutrons within the room may bounce off the walls until they are absorbed. In this case, the number of non-shadowed proportional counters will have their neutron flux increased by the mirroring effects of the wall.

With regard to the supernova sensitivity, the last cooldown data indicated around 49,500 neutrons per day were detected by the NCD array. If all of these neutrons are assumed to come from neutrons within norite, then the effective neutrino target mass can be calculated independent of neutron transport questions by the simple ratio of rates:

$$M_s = \frac{r}{R_{(\alpha,n)}},$$

where r is the rate of detected neutrons and $R_{(\alpha,n)}$ is the rate of alpha-n

neutron production in the target mass.

The production rate of neutrons from (α,n) reactions in sand is 0.3 neutrons $\text{ppm}^{-1} \text{ gm}^{-1} \text{ year}^{-1}$. As we detected 1.79×10^7 neutrons/year, we should be sensitive to 6.0×10^7 ppm grams. Given a total assay of 1.2 ± 0.2 ppm ^{238}U and 5.3 ± 1.1 ppm ^{232}U , the array was sensitive to $9.3 \times 10^6 \pm 1.6 \times 10^6$ grams of norite. This was done with a little more than half of the full array. The prediction for a full array, should the individual counter efficiencies prove similar, is approximately 1.4×10^7 grams, or less than one-half of the geometrically predicted value. The difference between the predicted and measured neutron rate as a function of added counter illustrates the more gas-like behavior (as opposed to straight-path movement) of the neutrons within the control room cavity.

Consequently, a galactic-core supernova could not be expected to produce more than a few neutrons that could be detected, certainly not enough to trigger a supernova warning.

One possible improvement would involve the insertion of a moderator, such as paraffin or water, into the NCD array, which would amplify both signal and background, effectively increasing the statistical significance of an event. Repositioning the NCD array, thereby increasing the geometric efficiency by minimizing detector ‘shadowing’ within the control room would make a second improvement. Finally, as was pointed out by Hargrove⁴², a substantial amount of natural Pb would dramatically increase the sensitivity of the NCDs to supernovae, since Pb has a neutrino-neutron cross-section by mass almost 20 times higher than norite. When this is coupled with Pb’s long neutron absorption length and very low neutron background, even 10 tons of strategically placed Pb with moderator would more than quadruple the expected neutron flux from supernovae without substantially altering the background. With a lead target mass on the order of 100 tons, we would expect dozens of neutrons to be detected.

Such a detector assembly would represent a very cost effective, low maintenance supernova monitor. It has both charged-current and neutral-current sensitivity via the neutron multiplicity: energetic electron neutrinos can produce two neutrons from ^{208}Pb . At the conclusion of SNO’s operations with heavy water on December 31, 2006, transfer of the NCD array into an 80-metric ton lead matrix that was salvaged from the Deep River cosmic-ray neutron monitor has been proposed with the name “HALO” (Helium and Lead Observatory).

Chapter 5 Approaches to NCD Backgrounds

The Neutral Current Detection array (NCDs) provides the unique potential for event-by-event event recognition within SNO. If events are properly identified, the NCDs should produce an accurate neutron flux measurement for the D₂O volume. This neutron flux measurement can be used directly in determining the flavor-independent neutrino flux. Unfortunately, naturally occurring radioactive isotopes from the ²³²Th and ²³⁸U decay chains found on and in the NCDs can interfere with this process in two ways. Primarily, they interfere by creating signals inside NCDs that may be misidentified as neutron capture events. Secondly, highly energetic α decays may provide enough energy to liberate neutrons from the D₂O itself. We begin this section with a discussion of ways to determine these backgrounds.

THORIUM				A = 4n				Ra 228 5.75 a	←	Th 232 1.4·10 ¹⁰ a
								Qβ 0.039 0% 0.0145 40%	↘	Ac 228 6.13 h
	Qβ 0.569 12%	Pb 212 10.64 h	←	Po 216 158 ms	←	Rn 220 55.6 s	←	Ra 224 3.64 d	←	Th 228 1.913 a
	Tl 208 3.053 m	←	Bi 212 1.0092 h	Qβ 2.246 63% 1.5 10%						
	Qβ 1.8 48.8% 1.52 22.7% 1.29 24%	Pb 208 stable	←	Po 212 298 ns						

Figure 5-1: ²³²Th Decay Chain. Isotopes are kept in vertical columns. Horizontal rows represent changes of equal neutrons and protons, hence alpha decays occur on the horizontal. (Figure courtesy J. Farine)

The ²³²Th decay chain, shown above in Figure 5-1, has a number of distinguishing characteristics in both energy and time that we present in this chapter. The decay of highest concern to the neutron background signal immediately follows the beta decay of ²⁰⁸Tl which always leads to the emission of a 2.6 MeV gamma. An attempt at using this gamma to identify contamination in the tubes with a coincidence of NCDs and SNO PMTs is presented in section 5.1. The Monte-Carlo of expected response is given immediately following in section 5.2.

Further inspection of the ²³²Th decay chain shows an identifiable short-time triple coincidence between the decays of ²²⁴Ra → ²²⁰Rn + α, ²²⁰Rn → ²¹⁶Po + α and ²¹⁶Po → ²¹²Pb + α. We will examine these coincidences in section 5.3, below. There is also an interesting high-energy beta released in the (²¹²Bi, ²¹²Po + □) decay, which is almost immediately followed by a hard 8.78 MeV alpha decay in (²¹²Po, ²⁰⁸Pb + α). Although this beta decay is essentially simultaneous when given the drift times within the tube, there may be

sufficient geometrical differences in the ionization trails to glean some information in the resulting digitized current signals.

URANIUM - RADIUM $A = 4n + 2$										Th 234 24.10 d	← 4.196 77% 4.149 23%	U 238 4.49 · 10 ⁹ a
			Bi 214 Q _β 3.26 19% 1.51 40% 1.02 23%							Q _β 0.199 73% 0.104 21%	Pa 234* 1.18 m 4.7 h	
	Pb 214 Q _β 1.03	Pb 214 26.8 m	← 6.00	Po 218 3.05 m 99.98% 0.02%	← 5.49	Rn 222 3.825 d	← 4.78 94.5% 4.60 5%	Ra 226 1602 a	← 4.69 76% 4.62 23%	Th 230 7.52 · 10 ⁴ a	← 4.78 72% 4.72 28%	U 234 2.48 · 10 ⁵ a
	Tl 210 1.32 m	← 5.51 39% 5.49 54%	Bi 214 19.8 m 0.04% 99.96%	← 6.7 90%	At 218 1.6 s							
	Q _β 2.34 19% 1.87 56% 1.32 25%	Pb 210 22.3 a ~100%	← 7.69	Po 214 162 μs								
		Q _β 0.061 19% 0.015 81%	Bi 210 5.013 d ~100%	Bi 210 Q _β 1.16					← E _α MeV RI%		α% BR β%	
		Pb 206 stable	← 5.31	Po 210 138.376 d							β Q _β MeV RI%	

Figure 5-2: ²³⁸U Decay Chain. (Figure courtesy J. Farine)

The ²³⁸U chain is shown in Figure 5-2. Unfortunately, there is no analogue to the unique timing of the triple alpha decay present in the ²³²Th decay chain for ²³⁸U. However, the 3.05-minute half-life of ²¹⁸Po is short enough that it can be reasonably separated from the next shortest double-alpha decay in its two potential sequences. The next shortest alpha to alpha decay half-life is approximately 47 minutes between ²¹⁸Po → ²¹⁴Pb + α and ²¹⁴Po → ²¹⁰Pb + α. It is also interesting to note that the extremely short 162 μs half-life of ²¹⁴Po is preceded by a relatively high-energy beta decay.

5.1 Beta-Gamma and Gamma-Gamma Coincidence

An abridged level diagram of the final decay of the ^{232}Th chain, $^{208}\text{Tl} \rightarrow ^{208}\text{Pb} + \beta^-$ (2.6 MeV) + γ (LE) + β^- is shown in Figure 5-3.




Figure 5-3: $^{208}\text{Tl} \rightarrow ^{208}\text{Pb}$ Decay Scheme

Although the low energy gammas are not directly detectable in the proportional counter, there is potential to observe the electrons from Compton scattering within the nickel walls. The entire spectrum of emitted gammas during the relaxation of the ^{208}Pb nucleus is shown in Table 5-1.

Table 5-1: Common Gammas from ^{208}Tl decay

E_γ (keV)	I_γ (%)	Decay mode
211.40	0.178	γ
233.36	0.307	γ
252.61	0.69	γ
277.351	6.31	γ
510.77	22.6	γ
583.191	84.5	γ
722.04	0.201	γ
763.13	1.81	γ
860.564	12.42	γ
927.6	0.131	γ
982.7	0.203	γ
1093.9	0.40	γ
2614.533	99	γ

Some of the beta-decay electrons can also be detected, depending primarily on the beta track through the counter gas. The ionization energy deposited by the beta tracking through an NCD counter alone is usually too small to trigger the NCD electronics. A table of beta energies emitted during the decay of ^{208}Tl are shown in Table 5-2.

Table 5-2: Betas Emitted from ^{208}Tl Decay.

E_{β} endpoint (keV)	I_{β} (%)	Decay mode
520.5	0.052	5 β^{-}
618.1	0.017	5 β^{-}
642.54	0.044	5 β^{-}
704.83	0.08	2 β^{-}
820.59	0.227	12 β^{-}
875.69	0.168	11 β^{-}
1005.3	0.007	3 β^{-}
1040.04	3.09	17 β^{-}
1054.56	0.043	4 β^{-}
1081.2	0.63	6 β^{-}
1292.56	24.5	6 β^{-}
1525.89	21.8	4 β^{-}
1803.26	48.7	12 β^{-}

In order to reliably observe the effect of these events in the NCD bodies, we use a trigger supplied by the de-excitation of ^{208}Pb . In this section, we examine the NCD counter response when a 2.6-MeV gamma is detected.

The ^{232}Th decay rate from the NCDs at any position, \mathbf{x} , is simply:

$$c(\mathbf{x}) = \frac{n}{\epsilon_{NCD\gamma}(\mathbf{x}) \cdot \epsilon_{SNO\gamma}(\mathbf{x})},$$

where n is the detected rate of PMT-NCD gamma-gamma coincidences, $\epsilon_{NCD\gamma}$ is the NCD's efficiency of detecting a low energy gamma from ^{208}Tl decay, and $\epsilon_{SNO\gamma}$ is the

SNO PMT efficiency of 2.6 MeV gamma detection. The false neutral-current rate induced by the ^{208}Tl contamination can be calculated:

$$F = \int c(\mathbf{x})\eta(\mathbf{x})d\mathbf{x} = n \int \frac{\eta(\mathbf{x})}{\varepsilon_{NCD\gamma}(\mathbf{x}) \cdot \varepsilon_{SNO\gamma}(\mathbf{x})} d\mathbf{x},$$

where $\eta(\mathbf{x})$ is the efficiency of a 2.6 MeV gamma in creating a false neutral (i.e. a photodisintegration neutron) current event at a given location in the heavy water.

The gamma-gamma and beta-gamma coincidence tests were studied a stainless-steel prototype counter, 184 centimeters in length and 2.5 cm in radius. The wall of the cylinder is 250 μm thick and the center copper anode wire has a 50 μm diameter. These coincidence tests used a nonstandard 60:20:20 $^4\text{He}:^3\text{He}:\text{CF}_4$ gas mixture at 2.5 atmospheres. An ‘IQ’ model preamp was used, which provided basic current integration, and this stainless steel counter was held at a gas gain of 100 by maintaining 2250V on the anode.

The steel prototype counter was calibrated using an americium-beryllium neutron source, and the lower part of the proportional counter spectrum is shown on Figure 5-4. The calibration relied on the proton and triton ‘edges’ of the spectrum which lie at 191 and 572 keV. As the reaction $^3\text{He}(n,p)^3\text{H}$ releases 763 keV, momentum conservation dictates that the fractional distribution of energy will be 3:1. Thus, the proton is given 572 keV of kinetic energy, while the triton emerges with 191 keV. The ‘edges’ result when one of the particles is completely captured within the active detector volume, while the other is completely absorbed in the wall.




Figure 5-4: Prototype NCD Neutron Spectrum with reduced threshold on the 184-cm steel prototype counter and IQ preamp.

The coincidence tests were run with a NaI crystal detector with its axis perpendicular to the proportional counter (see Figure 5-9). The NaI crystal used is 135 mm in diameter, 110 mm in height, and was calibrated using both ^{60}Co and ^{232}Th sources. The ^{232}Th source is 3 classic thorium-nitrate-dipped Coleman lantern mantles, with a total activity of approximately 5×10^{-8} Ci, stacked and triple bagged. The mantles are flattened out, so that the basic shape is rectangular, with slightly skewed edges. The stacked mantles have a surface that is 3 cm by 6 cm, so we do not expect them to exhibit perfect r-squared dependence.

The NaI-proportional counter coincidence measurement uses only a lower bound on the energy deposited in the NaI. This measurement was set up to show the ability of the proportional counter and sodium iodide counters to detect associated particle decays from the ^{232}Th source. For this test, the source was placed directly between the NaI

crystal and the proportional counter. A block diagram of the electronics is shown as Figure 5-5.

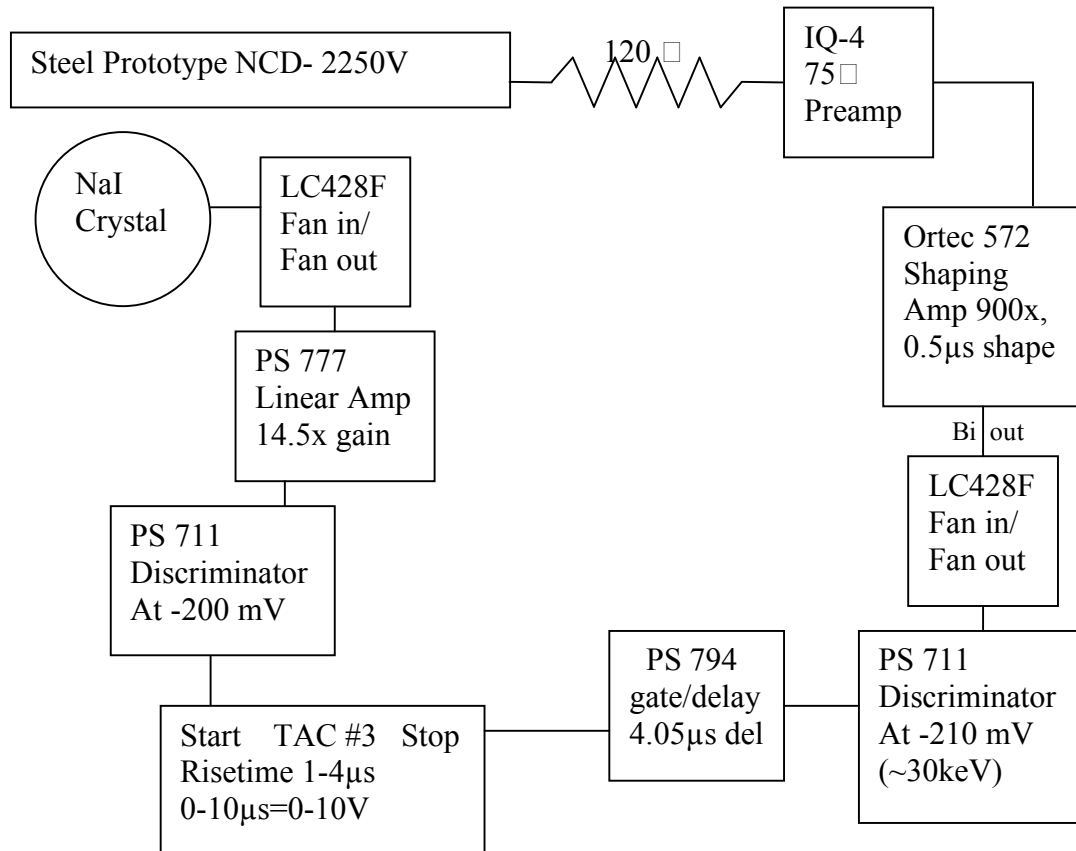


Figure 5-5: Beta-Gamma Coincidence TAC Spectrum Setup

The gamma-gamma coincidence measurement TAC (Time to Amplitude Converter) timing spectrum between an over-threshold signal in the NaI and the signal in the proportional counter is shown below in Figure 5-6. The strong time correlation of events demonstrates their association. The time delay between the start of the NaI pulse and the start of the proportional counter pulse is a function of detector gas drift speed and also of electronics.




Figure 5-6: TAC spectrum between NaI and NCD events. The time correlation corresponds to the combination of drift time and trigger threshold differences in the detectors.

We needed to verify that signals seen in coincidence with the NaI crystal were indeed associated decays and were neither background neutrons nor cosmic background events. This was demonstrated when we took a spectrum of the event energies measured in coincidence with the NaI 2.5 to 2.8 MeV triggers. A block diagram of the electronics used in this configuration is shown in Figure 5-7.

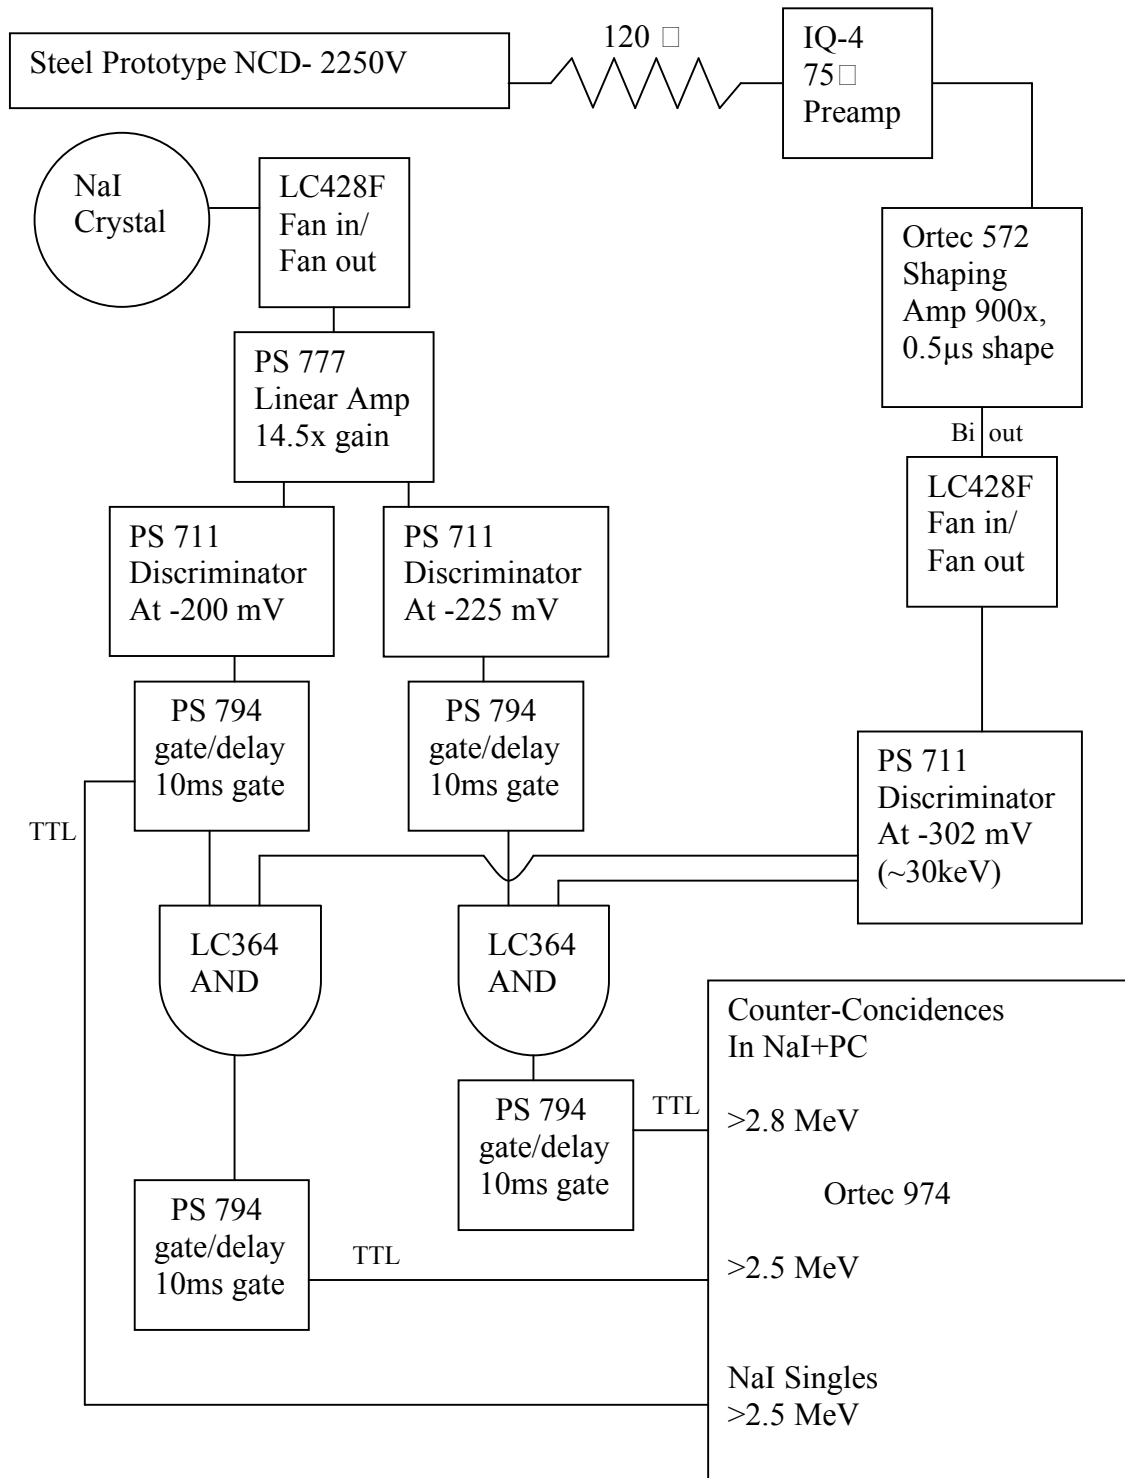


Figure 5-7: Beta-Gamma coincidence setup

The data plot, shown in Figure 5-8, clearly shows that the energies involved with the majority of the coincidences is well below the >191 keV expected for neutrons and has the general spectral shape expected for electrons.




Figure 5-8: Energy Spectrum of Associated NCD Events

Since the 250-micrometer thick steel wall of the prototype counter is not thick enough to stop the higher-energy beta decays, we placed a 0.32-centimeter thick piece of aluminum in between the source and the proportional counter. Running tests before and after the aluminum wall was placed between source and detector gave us the associated beta-decay detection rate.

In order to determine the efficiency of an entire ${}^3\text{He}$ proportional counter in detecting lower-energy gammas associated with the 2.6 MeV ${}^{208}\text{Tl}$ decay, we must

understand its positional dependence. This was accomplished by placing the ^{232}Th source at different positions along the detector, counting the number of simultaneous proportional-counter events over 20 keV that arrived in conjunction with a NaI event of 2.5 to 2.8 MeV, and subtracting out the expected background. The efficiencies were measured at the locations indicated in Figure 5-9, below, providing a map of efficiency vs. source position. These efficiencies are measured with the 0.32" aluminum shield between source and counter.




Figure 5-9: Position of Thorium Tests Along NCD Body

Positions 1, 2, and 3 yielded $0.211 \pm 0.006\%$, $0.132 \pm 0.006\%$ and $0.097 \pm 0.010\%$ efficiencies, respectively. Position 4 yielded $0.032 \pm 0.008\%$ efficiency, but is on the far side of 1.27 cm of steel.

Taking into account the energies of the secondary gammas, we can calculate the penetration depth in steel⁶⁵. Approximating the steel counter to be 85% iron and 15% chromium, and ignoring any ^{208}Pb decay with a likelihood of less than 1%, we find that the penetration depth of the secondary gammas varies between 0.90 and 0.51 centimeters.

Adding together the probabilities of each secondary decay penetrating past the 1.27 centimeters of steel, we find that only a fraction make it to the counter, T :

$$T = \frac{Ae^{-\alpha(1.27\text{cm})} + Be^{-\alpha(1.27\text{cm})} + Ce^{-\alpha(1.27\text{cm})} + De^{-\alpha(1.27\text{cm})} + Ee^{-\alpha(1.27\text{cm})}}{A + B + C + D + E}$$

$$T = \frac{6e^{-0.9(1.27)} + 23e^{-0.65(1.27)} + 85e^{-0.61(1.27)} + 12e^{-0.55(1.27)} + 1e^{-0.51(1.27)}}{127}$$

$$T = 0.456$$

This means that, in addition to the reduced solid angle, over half of the gammas incident on the steel endcap are absorbed. The end regions are of particular concern, as they are made from a suite of different materials and there can be no in-situ background determination from alpha measurements. Applying this to the measured efficiency implies that, for a detector without a half inch of steel at the end, the efficiency of associated gamma production should be around 0.06%. To compare the efficiency of gamma detection off the end of the counter with the efficiency along its side, we examine the differences in solid angle. The solid angle subtended by inner counter walls visible to the gamma at position 4 on Figure 5-9 is:

$$\Omega = \iint_S \frac{\mathbf{n} \cdot d\mathbf{A}}{r^2},$$

where \mathbf{n} is the unit vector in the direction from the gamma source to the wall element $d\mathbf{A}$. In spherical coordinates, neglecting the solid angle of the more distant end of the detector,

$$\Omega = \iint_S \sin \theta d\theta d\phi$$

$$\Omega = 2\pi(1 - \cos \theta)$$

$$\Omega = 2\pi \left(1 - \frac{R}{\sqrt{R^2 + r^2}} \right)$$

Where R is the distance from the source to the counter wall, and r is the radius of the counter. Given the source at position 4 is 4 cm from the end, we have about 15% of the solid angle of an event dead-center on the counter. A graph of position vs efficiency for detection of secondary gammas on a single counter is shown as Figure 5-10.




Figure 5-10: Single-counter gamma detection efficiency as a function of distance from end. Data is from 92-cm steel prototype counter.

Extrapolating these data, we can graph out what we would expect for contaminants along a string of counters. The values listed are calculated by taking the efficiency of detection for the secondary gamma on the closest tube, and adding the solid-angle corrected detection efficiency from the next most distant counter. The results are shown in Figure 5-11. At the ends of strings, the efficiency would drop off to values as shown at the extremes of Figure 5-10.




Figure 5-11: Expected gamma detection efficiency along a string of 92 cm counters with just under 20 centimeters between successive counters.

Note that every potential contaminant position has greater than 0.1 percent efficiency, with a maximum variation around a factor of two. This result is compared with Monte-Carlo in section 5.2.

5.2 Monte-Carlo

*Anyone who considers arithmetic methods of producing random digits is, of course, in a state of sin. **John von Neumann***

The response of NCDs was simulated in SNOMAN, the SNO Monte Carlo and ANalyzer. SNOMAN⁶⁶, integrates both simulation and data analysis into one code. SNOMAN is primarily written in Fortran-77 and uses ZEBRA⁶⁷ for memory management, allowing data structures that would otherwise be unavailable. SNOMAN, in Monte Carlo mode, is an EGS4⁶⁸-based simulation that contains the geometries for every object that is or will be brought into contact with the SNO detector. The output of a SNOMAN Monte Carlo can be recorded and analyzed just like real SNO data, without the long waiting process. SNOMAN's analysis packages allow it to reconstruct events within the detector.

The NCDs in SNOMAN are treated as an array of objects, each of which possesses 5 distinct regions and 2 additional associated regions. A list of the different regions, their associated identifying SNOMAN names, and their ZEBRA numerical representations are displayed in

Table 5-3.

Table 5-3: SNOMAN NCD Regions

Region	SNOMAN Name	ZEBRA ID
NCD Nickel Body	NCD_OVL	110
NCD ³ He-CF ₄ Gas	NCD_LIVL	111
NCD Inter-counter Dead Volume	NCD_DIVL	112
NCD Counter -edge endcaps	NCD_EC	113
NCD Inner Anode Wire	NCD_WIRE	114
NCD Cable	NCC	120
NCD Anchor Point	NCC_ANC	121

SNOMAN's geometry routine requires that regions be considered hierarchically. This means that a particle can be within a large number of regions, but the particle is only interacting with the material in the foremost region. Additionally, all regions must have a list of bordering regions into which the particle can move. To minimize the number of checks placed on a moving particle, all regions are created with overlap zones. In the case of SNOMAN's NCDs, the overlap regions and hierarchy are shown in Figure 5-12.




Figure 5-12: Varying transparency side-view of SNOMAN NCD Geometry Regions and Overlaps, regions not to scale.

The NCDs are set up in SNOMAN as a variable array of NCD Outer Shells. The location and contents of each string may be adjusted by parameters within SNOMAN's *geometry.dat* file. In each string, the existence and length of up to five counters is set, and the total length of string must reflect these values. *NCDarray.for* simplifies the creation of the array data for SNOMAN by translating a small set of parameters specified within the program into the detailed description of each array element expected in

geometry.dat. The output of *NCDarray.for* is, by default, stored into *NCDarray.dat* for later copying into *geometry.dat*. The NCDs as an array may be entirely turned off or turned on by adjusting the SNOMAN ‘ncd’ symbol by running the @geom_ncd command. Due to the nature of SNOMAN’s geometry code, only one material may be specified per region. If a subset of the array needs to be simulated, then *snoman.dat* must be adjusted appropriately. The cables for the NCD array, put into SNOMAN by S.J. Brice, may be selected in each simulation by enabling the NCD Cables in SNOMAN code.

The Monte Carlo simulation for beta-gamma and gamma-gamma coincidence measurements between the SNO PMTs and the NCDs was done with the originally envisioned full array (96 strings) in place. Each NCD in the array was set to the lengths determined by the deployment plan of 1996. Two sets of simulations were run, one for ^{214}Bi , the other for ^{208}Tl . For each of the elements, a million decays were run, with three different contaminant distributions. In each simulation, SNO’s expected PMT response in the NHIT (Number of firing photomultiplier tubes) spectrum above a predetermined number was used as the ‘trigger’. If the trigger fired, the energy deposited within the active counter gas region was plotted.

The isotope ^{214}Bi is the only daughter of the ^{238}U chain which has a gamma decay with energy in excess of the 2.2 MeV required for neutron liberation in the D_2O . Since the Q-value of the decay is not much higher than 2.2 MeV, the energy of any potential associated beta and gamma-decays is very low, and unlikely to propagate far through the NCD walls. It is also important to note that 36% of the ^{232}Th progeny, ^{212}Bi result in a

2.6 MeV gamma. This can be compared to the ^{238}U progeny, ^{214}Bi , for which only 3% of the decays result in high-energy gammas. A histogram of the energy deposited, in keV, within the active counter gas volume, when a ^{212}Bi decay resulted in a >10 NHIT response from SNO, is shown as Figure 5-13.




Figure 5-13: Energy deposited within NCD inner volume when gammas from ^{212}Bi decay resulted in >10 NHit in PMT array. There were 536 total energy-depositing (in the NCD gas) coincident events.

^{208}Tl decays to ^{208}Pb with a 2.6 MeV gamma emission 99% of the time. This ^{232}Th daughter is the only one to surpass the neutron liberation threshold in the D_2O , and does it with a substantial amount of energy to spare. As a result, each ^{208}Tl decay is

about twice as likely to liberate a neutron from the deuteron than the 2.4 MeV gamma emitted by ^{214}Bi . The histogram of the energy deposited within the active NCD volume, when more than 10 PMTs fire, is shown in Figure 5-14.




Figure 5-14: Energy deposited in counter during >10 NHit event. This simulation had many fewer total events than the corresponding ^{212}Bi (Figure 5-13).

While the gamma-gamma coincidence technique is clearly a viable means for assessing the amount of Th in counters, especially that which may be hidden in the

endcaps, in practice the thresholds of the NCDs are set above 100 keV to reduce the data rates. At this threshold, it is impossible to detect most gamma-induced events.

5.3 NCD contaminant alpha-alpha self-analysis

Alphas that deposit a relatively small fraction of their kinetic energy within a counter can mimic a certain subset of ${}^3\text{He}(n,p){}^3\text{H}$ reactions. Given the relatively short penetration depth in nickel, very few alphas from contaminants within the bulk will range into the counter gas. However, alphas from contaminants within the counter gas, on and included in the immediate inner NCD surface have a high probability of creating detectable events. As was previously noted^{69,70}, differences in track lengths between alpha and neutrons of the same energy might allow the identification of a subset of neutrons free of alpha background. Unfortunately, the gas gain used in deployed NCDs is too high—the ionization saturation shields the electric field from the center anode, causing some alphas to draw out their signal risetime into the neutron region because they have higher energy than their collected charge indicates.

Independent of gas gain and saturation concerns, the large amounts of energy deposited by some alphas as they pass through the counter gas provide near-absolute particle identification. In addition, the major contaminant decay series both have distinguishing time-coincidence alpha-alpha decays. Using these definitive identifiers, we are able to perform alpha-alpha coincidence analyses to estimate the ${}^{232}\text{Th}$ and ${}^{238}\text{U}$ contamination.

Take the rate of ‘triple-coincidence’ ^{232}Th decays within the nickel body of the detector to be $T = \int \rho_{Th}(x) dx$, where ρ is the decay rate density in decays per second per unit depth, and x is the depth from the surface. Let the rate of ^{238}U decays in the ‘double-coincidence’ be $\rho_U(x)$. Let us also label the place of a decay in its decay sequence with a subscript, such that $U_1(x)$ would give us the total rate of detected 5.6 MeV alphas from the ^{238}U decay chain progeny $^{222}\text{Ra} \rightarrow ^{218}\text{Po} + \alpha$ at a depth x .

The time-coincident decays in the thorium and uranium chains have different energies and time separations as follows:

^{238}U ‘double coincidence’: 5.49 MeV alpha from ^{222}Ra to ^{218}Po , half-life of 182 seconds before ^{218}Po decays to ^{214}Pb with a 6.00 MeV alpha.

^{232}Th ‘triple coincidence’: ^{224}Ra emits a 5.69 MeV alpha while decaying to ^{220}Rn , with a half-life of 55.6 seconds before emitting a 6.23 MeV alpha on its way to ^{216}Po , which, in turn, emits a 6.78 MeV alpha with a half-life of 0.15 seconds, resulting in ^{212}Pb .

The probability, $P(x, E)$ of detecting any given alpha of initial energy E from a decay is a function of depth. The total rate of detection is the depth integral of the probability of detection multiplied by the rate of decay, both as a function of depth:

$$U_{12} = \int P(x, E_1) P(x, E_2) \rho_U(x) dx$$

$$Th_{123} = \int P(x, E_1) P(x, E_2) P(x, E_3) \rho_{Th}(x) dx$$

$$\begin{aligned}
Th_{2x} &= \int P(x, E_1)P(x, E_2)\bar{P}(x, E_3)\rho_{Th}(x)dx \\
&+ \int P(x, E_2)P(x, E_3)\bar{P}(x, E_1)\rho_{Th}(x)dx \\
&+ \int P(x, E_1)P(x, E_3)\bar{P}(x, E_2)\rho_{Th}(x)dx
\end{aligned}$$

Note that Th_{23} is distinguishable from U_{12} in the regime when the total number of detected Th_{23} decays, C_{Th} , is larger than the uncertainty of the expected number of detected U_{12} decays C_U , that both fall within the same time period. Since the Th_{23} decay has a much shorter time separation, if one half-life is used to sample the C_{Th} rate, then the statistical uncertainty in the C_U rate is:

$$\begin{aligned}
C_{U=Partial} &= C_U \left(1 - \frac{1}{2} \frac{0.15s}{182s} \right) \\
&= (.00057)C_U \\
\Delta C_{U=Partial} &= \sqrt{.00057C_U}
\end{aligned}$$

Therefore, if the total number of detected C_{Th} decays within that Th_{23} half-life is of comparable magnitude to the uncertainty in the Uranium chain-contribution, the amount of thorium contamination should be separable.

$$C_{Th} > 0.024C_U$$

If we approximate the stopping power of nickel to be constant, then our detection probability is roughly a straight line from maximum detectable depth with zero probability to 0 depth with 0.5 probability of detection as shown below. The reason for this distribution is simple. Assume azimuthal symmetry and take the parent particle's distance from the counter gas to be 'd'. The angle that the particle travels, taken as a

difference from the straight-line path normal towards the counter surface is ' θ '. The total distance the particle can travel and still carry sufficient energy to be detected within the counter volume is ' R '. This somewhat simplified situation appears like Figure 5-15:




Figure 5-15: Simplified counter cross-cut with a decay at a depth d inside the nickel wall (above) traveling at an angle θ to the inside counter surface (below). The circle at R represents the maximum distance the alpha particle can travel before dropping below detection threshold.

As the alpha flux into any element of solid angle is equal, any angle θ has a probability $dP(\theta) = \frac{1}{2} \sin \theta d\theta = \frac{1}{2} d(\cos \theta)$ of containing the alpha. Likewise, all azimuthal rings on the surface of this sphere should expect a probability of $\Delta P = \frac{1}{2} (\cos \theta_1 - \cos \theta_2)$. Since we are looking for the chance that an alpha makes it to the surface, given $R > d$, the minimum angle $\theta_2 = 0$, and the maximum:

$$R = \int dx$$

$$= \int_0^{E_0} \frac{dE}{\left(\frac{dE}{dx}\right)}$$

$$\theta_2 = \frac{1}{2} \left(1 - \frac{x\eta}{E_0} \right),$$

where $\eta = \frac{dE}{dx}$. Figure 5-16 is a graph of this probability as a function of the fraction of the depth of the parent particle to the maximum possible observable depth.




Figure 5-16: Probability of Alpha Detection at Fraction of Max Depth. Typical depths for alpha decays in nickel are around 10nm.

The detection of one alpha modifies the chance of detecting the second, in that the distribution of expected particles is multiplied by the probability of detection. A depth-independent bulk distribution, for example, would turn into a profile of particles with density distribution matching that probability profile. These profiles do not take in to consideration the recoil of the parent particle, which carry almost 2% of the alpha's energy, and have predicted movement within the nickel of up to 15 nm. Under most circumstances, this is a negligible effect so it is not accounted for in this section. These profiles also do not account for the decreasing \square at low energies. Neglecting these effects, the chance of detecting two out of two alphas, as in ^{238}U , is:

$$\begin{aligned} P_{12} &= \frac{1}{4} \left(1 - \frac{x\eta}{E_1} \right) \left(1 - \frac{x\eta}{E_2} \right) \\ &= \frac{1}{4} \left(1 - \frac{x\eta}{E_1} - \frac{x\eta}{E_2} + \frac{x^2\eta^2}{E_1E_2} \right) \end{aligned}$$

All 3 of the ^{232}Th alphas can be detected only if:

$$\begin{aligned} P_{123} &= \frac{1}{8} \left(1 - \frac{x\eta}{E_1} \right) \left(1 - \frac{x\eta}{E_2} \right) \left(1 - \frac{x\eta}{E_3} \right) \\ &= \frac{1}{8} \left(1 - \frac{x\eta}{E_1} - \frac{x\eta}{E_2} - \frac{x\eta}{E_3} + \frac{x^2\eta^2}{E_1E_2} + \frac{x^2\eta^2}{E_1E_3} + \frac{x^2\eta^2}{E_2E_3} - \frac{x^3\eta^3}{E_1E_2E_3} \right) \end{aligned}$$

In a similar manner, 2 of the 3 ^{232}Th alphas:

$$P_{2x} = \frac{1}{8} \left[\left(1 - \frac{x\eta}{E_1}\right) \left(1 - \frac{x\eta}{E_2}\right) \left(1 + \frac{x\eta}{E_3}\right) + \left(1 - \frac{x\eta}{E_1}\right) \left(1 - \frac{x\eta}{E_3}\right) \left(1 + \frac{x\eta}{E_2}\right) + \left(1 - \frac{x\eta}{E_3}\right) \left(1 - \frac{x\eta}{E_2}\right) \left(1 + \frac{x\eta}{E_1}\right) \right]$$

$$= \frac{1}{8} \left(3 - \frac{x\eta}{E_3} - \frac{x\eta}{E_2} - \frac{x\eta}{E_1} - \frac{x^2\eta^2}{E_1E_2} - \frac{x^2\eta^2}{E_1E_3} - \frac{x^2\eta^2}{E_2E_3} + \frac{3x^3\eta^3}{E_1E_2E_3} \right)$$

Graphically, the expected depth distribution of all four of the cases mentioned above, assuming equal energy alphas, is shown on Figure 5-17.




Figure 5-17: Probability of Multiple Alpha Detections

In order to evaluate how the detected amount of coincident decays corresponds to the actual contaminant amount, we can approximate the density of contaminants to be, as a function of depth, x :

$$\rho(x) = B + S\partial(x)$$

Where B represents the bulk contamination of the nickel, and is not a function of depth, and S represents the surface contamination.

During decay detections, time separation and energy deposited within the counter gas can be used to identify the source of alphas. Consecutive event analysis through time-separation is useful so long as expected coincidence rates are understood, and backgrounds can be properly accounted for.

The (^{224}Ra \diamond ^{220}Rn \diamond ^{216}Po \diamond ^{212}Pb) triple-decay is the most recognizable and rare. Even if the ^{232}Th contamination were expected to be as high as the ^{238}U , the triple detection is considerably rarer: Extracting probability and density distributions from above, we see that the rate of triple-alpha detection is (again assuming \square is not a function of x):

$$\begin{aligned} Th_{123} &= \int_0^D P(x, E_1)P(x, E_2)P(x, E_3)\rho_{Th}(x)dx \\ &= \int_0^D \frac{1}{8} \left(1 - \frac{x\eta}{E_1} - \frac{x\eta}{E_2} - \frac{x\eta}{E_3} + \frac{x^2\eta^2}{E_1E_2} + \frac{x^2\eta^2}{E_1E_3} + \frac{x^2\eta^2}{E_2E_3} - \frac{x^3\eta^3}{E_1E_2E_3} \right) (B_{Th} + S_{Th}\partial(x)) dx \\ &= \frac{1}{8} S_{Th} + \frac{B_{Th}}{8} \left[x - \frac{x^2\eta}{2E_1} - \frac{x^2\eta}{2E_2} - \frac{x^2\eta}{2E_3} + \frac{x^3\eta^2}{3E_1E_2} + \frac{x^3\eta^2}{3E_1E_3} + \frac{x^3\eta^2}{3E_2E_3} - \frac{x^4\eta^2}{4E_1E_2E_3} \right] \Bigg|_0^D \\ &= \frac{1}{8} S_{Th} + \frac{B_{Th}D}{8} \left[1 - \frac{D\eta}{2E_1} - \frac{D\eta}{2E_2} - \frac{D\eta}{2E_3} + \frac{D^2\eta^2}{3E_1E_2} + \frac{D^2\eta^2}{3E_1E_3} + \frac{D^2\eta^2}{3E_2E_3} - \frac{D^3\eta^2}{4E_1E_2E_3} \right] \end{aligned}$$

Where $B_{Th}D$ is the total bulk thorium alpha-decay rate. D is the maximum depth that a particular energy alpha can travel through the nickel and still be detected upon entry into the NCD gas. In cases where multiple ranges are multiplied, we take the shortest range of the group. In the case of the triple decay, this corresponds to the 5.8 MeV alpha decay of ^{224}Ra . Since we are assuming constant energy loss, $D = \frac{E}{\eta}$, the probability of detection becomes:

$$\begin{aligned} Th_{123} &= \frac{1}{8}S_{Th} + \frac{B_{Th}D}{8} \left[1 - \frac{1}{2} - \frac{1}{2} - \frac{1}{2} + \frac{1}{3} + \frac{1}{3} + \frac{1}{3} - \frac{1}{4} \right] \\ &= \frac{1}{8}S_{Th} + \frac{B_{Th}D}{32} \end{aligned}$$

Note that this is substantially different from expecting contaminants within the bulk to be, on average, halfway between the surface and the total depth. The efficiency for single alpha detection at the midway point is roughly one half of surface efficiency, so we would expect to detect one out of four for each alpha. Detecting all three decays consecutively would be a 1 out of 64 rather than 32. Using the same technique, we can show that the probability of detecting a ^{238}U daughter double-decay is:

$$\begin{aligned} U_{12} &= \frac{1}{4}S_U + \frac{B_U D}{4} \left[1 - \frac{1}{2} - \frac{1}{2} + \frac{1}{3} \right] \\ &= \frac{1}{4}S_U + \frac{B_U D}{12} \\ U_{12} &= \frac{1}{4}S_U + \frac{B_U D}{12} \\ Th_{23} &= \frac{1}{8}S_{Th} + \frac{5B_{Th}D}{96} \end{aligned}$$

Likewise, we can determine that detecting any two of the three decay alphas from ^{232}Th decay daughters is:

$$\begin{aligned} Th_{2x} &= \frac{3}{8}S_{Th} + \frac{B_{Th}D}{8} \left[3 - \frac{1}{2} - \frac{1}{2} - \frac{1}{2} - \frac{1}{3} - \frac{1}{3} - \frac{1}{3} + \frac{3}{4} \right] \\ &= \frac{3}{8}S_{Th} + \frac{5B_{Th}D}{32} \end{aligned}$$

We can also look at the mean depth of the decay source, given a triple decay detection, by using a weighted average:

$$\bar{x}_{123} = \frac{\int_0^D P_{123} x dx}{\int_0^D P_{123} dx}$$

In the case of a triple decay, this equates to $\frac{D}{5}$, the average expected depth of the ^{232}Th daughter that undergoes the characteristic triple-decay and is detected in the NCD.

In the Uranium double decay, we expect the contaminant to be at $\frac{D}{4}$, if both progenies are detected. If two of the three ^{232}Th decays are detected, the expected depth increases to $\frac{7D}{25}$. A single decay can be trivially calculated to have come from, on average, $\frac{D}{3}$.

One problem that stands out in detecting two of the three decays in the ^{232}Th chain is that the 55-second half-life of ^{220}Rn needs to be distinguishable from the 182-second half-life of ^{218}Po . In order to recognize the shorter 55-second half-life double-decay, the total events must rise above the expected uncertainty of the 182-second half-

life decays. Although more sensitive measures such as decay-time curve fitting can be used, a basic way of measuring this effect is to take a 55-second window after a single alpha, in which we expect half of the detectable 55-second decays to appear, and $1 - 0.5^{\frac{55}{182}} = 18.9\%$ of the 182-second second alphas. If the uncertainty on that 0.189 is higher than the half of the 55-second expected events, then we cannot resolve the two decay chains easily. The same is true for distinguishing the 0.15 second double-alpha from the 182-second double-alphas. In a 2-second window, $1 - 0.5^{\frac{2}{182}} = 0.8\%$ of the 182-second decays should come through, compared to almost 100% of the 0.15 second half-life consecutive decays.

NCD coincidence event energy can also be used to help determine contaminant levels. The energy an alpha decay can put into the gas is a function of both the depth of the contaminant and the angle of the decay. Assuming an isotropic distribution of decay paths, we can find the expected distribution of alpha energy within the counter gas. If the decay can be recognized by the timing of the coincident events, then at the very least, a maximum depth can be established by assuming that the alpha traveled straight out of the nickel. Note that surface decays should deposit all of their energy, unless they range out into the counter wall. The expected alpha energy distributions, for alphas originating at varying depths within the counter walls, plotted for a 5.8 MeV alpha decay, are plotted below. Each of the distributions is normalized to equal area on Figure 5-18. In this fashion, we can see the expected energy spectrum for a 5.8 MeV alpha originating at each depth listed.




Figure 5-18: Alpha Energy Profile for Varying Source Depth

The result is not surprising—alphas starting deep inside the bulk will only be seen at low energies in the counter, but with a somewhat level energy distribution. Alphas beginning near the surface will deposit almost all of their energy within the counter; there is little solid angle for detection with significant energy loss. Resolving surface and bulk contamination levels can be done by analyzing the double and triple-coincidence data. Since the triple and 0.15-second double-decays in the ^{232}Th chain are both linear

combinations of surface and bulk contaminant levels, simple algebra can be used to determine the depth/surface contaminant ratio:

$$Th_{123} = \frac{1}{8}S_{Th} + \frac{B_{Th}D}{32}$$

$$Th_{23} = \frac{1}{8}S_{Th} + \frac{5B_{Th}D}{96}$$

$$B_{Th}D = 48(Th_{23} - Th_{123})$$

In order for this analysis to work, the profile of the distribution of contaminants must be as stated above, a delta function at the surface plus a constant level throughout the bulk. ^{238}U levels cannot be determined using this method alone, but an upper limit on the levels is obtained by assuming that all coincidence events originated within the bulk.

Recoil of the parent atom is ignored for purposes of this measurement. The parent nucleus can carry as much as 100 keV away from any one alpha decay, which will take the atom up to 15 nm through the nickel. This effect isn't enough to make a noticeable difference for bulk decays, but will effect surface contaminant detection. If the alpha decay heads into the bulk, a surface contaminant may be launched into the counter gas, and thereby have a greatly increased fraction of detected second decays. This might increase the amount of single-alpha detections per surface contaminant, when the second alpha follows the undetected first. This effect would only be noticeable in a coincidence on the Thorium triple decay, where the efficiency of detecting the last two decays would be increased. If we only care to get an upper limit on possible contaminant levels, we may simply ignore the effect of nuclear recoil.

5.3.1 Alpha-Alpha Coincidence Cooldown Data

Using NCD runs 6458 through 6984, 40.5 days during the summer shutdown of 2003, June 27 through August 7, we first apply this technique to determine the maximum levels of uranium and thorium within the nickel bodies of the NCDs. During this period, more than a year before deployment, we used 304 meters of NCDs for data. A full shaper/ADC spectrum of the energy deposited inside one full counter's gas is shown as Figure 5-19, and the curved line is the fit neutron peak:




Figure 5-19: NCD Cooldown Underground Spectrum

The energy deposited in a Shaper/ADC can be roughly approximated to be ADC channel*5 keV. These data include 'picket fence' events, in which the shaper places events in random bins which are evenly divisible by 256. Removing those events from

the data set³, as well as background testing runs 6707-6710, 7543 events remain over 36 strings. Making a conservative cut at channel 250 requires that, in order to be detected, an alpha must deposit over 1.25 MeV of energy into the counter gas. We use the approximation that the least energy in the coincident ²³²Th triple-decay chains is 5.8 MeV and ²³⁸U is 5.6 MeV. Using SRIM-2003.20⁷¹ to determine alpha energy loss in nickel, the 5.8 MeV alpha can travel 9.3 μ m and the 5.6 MeV alpha can travel 8.8 μ m before dropping below 1.25 MeV.

For the double-decay in the uranium chain, $\frac{1}{12}$ of the doubles should arrive with 1.25 MeV or more energy from a maximum depth of 8.8 μ m. This 8.33% needs to be adjusted before it can be checked against a Monte-Carlo expectation⁴⁴ of 3.50% efficiency. The Monte-Carlo was taken for all decays within 22 μ m of the surface, whereas the 1.25 MeV threshold requires the alpha decay begin no more than 8.8 μ m from the surface. To account for this difference in parent distribution, we may simply decrease the theoretical efficiency of detection by the fraction of added volume, as there is no chance for a double-alpha event to be recorded from beyond the penetration depth of the lower-energy alpha. This reduces the fractional efficiency to $0.0833 \left(\frac{8.8}{22} \right) = 3.32\%$. Furthermore, the Monte-Carlo assumed a cutoff of 1 MeV, which

³ “Picket Fence” events have been eliminated from data sets by fixing the DAQ. See section 2.6.5.

adds another $0.3 \text{ } \mu\text{m}$ to the range: $0.0833\left(\frac{9.1}{22}\right) = 3.45\%$, in good agreement with simulation. Finally, the calculated value is still expected to be low, as the detection probability as a function of depth for more energetic decay tends to be slightly higher.

For the triple-decay in the thorium chain, we expect $\frac{1}{32}$ from a maximum depth of $9.3 \text{ } \mu\text{m}$, which corresponds to a $22 \text{ } \mu\text{m}$ efficiency of $0.03125\left(\frac{9.3}{22}\right) = 1.32\%$. To compare this result with Monte-Carlo, the range must be adjusted upward by increasing the alpha energy by 1.15 MeV . This crosses over the Bragg peak, and increases the range to $11.3 \text{ } \mu\text{m}$. The value is then 1.61% compared to 1.71% , this difference being due to the significantly higher ranges of the last two events.

During the cooldown runs 6458 through 6984, 276 doubles events occurred within a 366-second time window. The background rates above channel 250 in each counter create an expected accidental coincidence background rate for doubles within this time window of 186 ± 16 counts. If all 90 of the double coincidences within this time window were due to ^{238}U -chain daughters within the bulk of the nickel walls, this only represents only 75% of the time-coincidences, so that 112.5 double-events are expected to have both left over 1.25 MeV of energy. Given that we only expected 8.33% efficiency from the $8.8 \text{ } \mu\text{m}$ sensitive region, 1351 double-decays were expected to have happened. The total number of events, $\square N$, is related to the uranium-chain decay rate $\square U$, and the amount of time $\square t$, and total number of contaminants in the $8.8 \text{ } \mu\text{m}$ of nickel wall, N , by:

$$\Delta N = N\lambda \Delta t.$$

Solving for the number of ^{238}U contaminants in the 8.8 μm cylindrical shell yields:

$$N = \frac{\Delta N}{\lambda_U \Delta t} = \frac{1351}{\left(\frac{\ln(2)}{(4.47 \times 10^9)(3.156 \times 10^7)} \right) (40.50)(86400)} = 7.86 \times 10^{13},$$

which is a total contaminant mass of 1.87×10^{16} amu. This compares to the total mass of the 304.5 meter-long 8.8 μm -thick cylindrical nickel shells with radius of 2.54 centimeters of 3.81 kg. The maximum ^{238}U contaminate mass ratio is 8.1 parts per trillion by mass, but the ^{232}Th double-decays that infiltrate this region have not been removed. Even in the worst-case scenario, however, this beats the design goal of 24 ppt.

For the ^{232}Th decay chain, in the full 40.50-day data set, only 4 candidate triple-decays were found. The time window allowed between the first and second alpha was 366 seconds, over 6 of the 55-second ^{220}Rn half-lives. Additionally, up to 1.5 ± 0.5 seconds were allowed for the time separation between the 2nd and 3rd alphas, for a minimum of over 6 half-lives. If all four of the candidate events are from ^{232}Th decay daughters in the bulk of the counter walls, they represent 310 ^{232}Th chain decays, accounting for a 2% loss from the time cut on the Poisson distribution. Computing the contamination fraction yields $5.2_{-1.7}^{+2.4}$ ppt ^{232}Th by mass inside the nickel. Deadtimes in the Shaper/ADC system contribute a negligible amount (0.1% effect on the 158 ms separation). Given this level of ^{232}Th contamination, we would expect that 36 double-

decays of all sorts will have been included on the previous ^{238}U level. Correcting for this, the adjusted mass fraction of ^{238}U is $4.8^{+2.2}_{-1.6}$ ppt.

These two measurements are in complete agreement with in-situ values of $5.7^{+1.0}_{-0.9}$ ppt ^{232}Th and $2.8^{+0.6}_{-0.8}$ ppt ^{238}U from [44].

5.3.2 In-Situ Alpha-Alpha Coincidence Measurements

The NCDs were deployed over the period from December 2, 2003 to April 21, 2004. During deployment, only 40 of the 96 NCD strings were placed within the D_2O volume, and their current configuration is shown as Figure 5-20.




Figure 5-20: Deployed NCD Configuration— ^4He counters I2,I3,I6,I7 are highlighted, and counters with time-correlated electronics backgrounds are bordered by dashed lines.

Counters I2, I3, I6, and I7 are filled with ^4He , and thus do not capture neutrons. The NCDs with dashed lines surrounding them in Figure 5-20 are counters with noise issues, and will be ignored for some of the following discussion.

The data used for the in-situ test is the Open Data set, taken between November 27, 2004 and January 3, 2005. The data come from 26.164 days of active triggering, spread out over 122 runs. Blindness cuts, where an unknown fraction of the data is removed from analysis, were not imposed on these runs.

The data was scanned through by a ROOT⁷² C++ script that looked for time-correlated events on a single NCD channel. Each channel was examined for total event rate so that background alpha rates could be determined. A cut was imposed, that both of the Shaper/ADC events in the coincidence deposit more than 1 MeV of energy. The code finds the corresponding digitized event for each Shaper/ADC event that is found to be part of a time-coincidence. Each digitized event is de-logged and saved into a file for later analysis. After a subset of runs is analyzed, each de-logged digitized event is hand-analyzed to remove spurious events which contain no potential. The process of scanning digitized events is accelerated by an interactive ROOT program, wherein the user is shown a set of digitized events. After each set of correlated events appear, the user must determine the potential validity of the events. If the digitized records are unclear or not available for an event, the user is supposed to accept the event. A page of such correlated events is shown as Figure 5-21.




Figure 5-21: Digitized Event Scan Sheet for a series of runs. Three pair of events line each row.

The basic livetime of each run is determined by taking the difference between the first and third from last events in the run. The deadtime is estimated by multiplying the total number of events occurring across all channels throughout the data set by the associated Shaper/ADC lockout time. The runs spanned 123 runs over a total time of 26.164 days. The total energy spectrum for each of the 40 ADCs are shown as

Figure 5-22. Note that channels(strings) 3(I2), 10(I3), 20(I6) and 30(I7) are all ^4He .

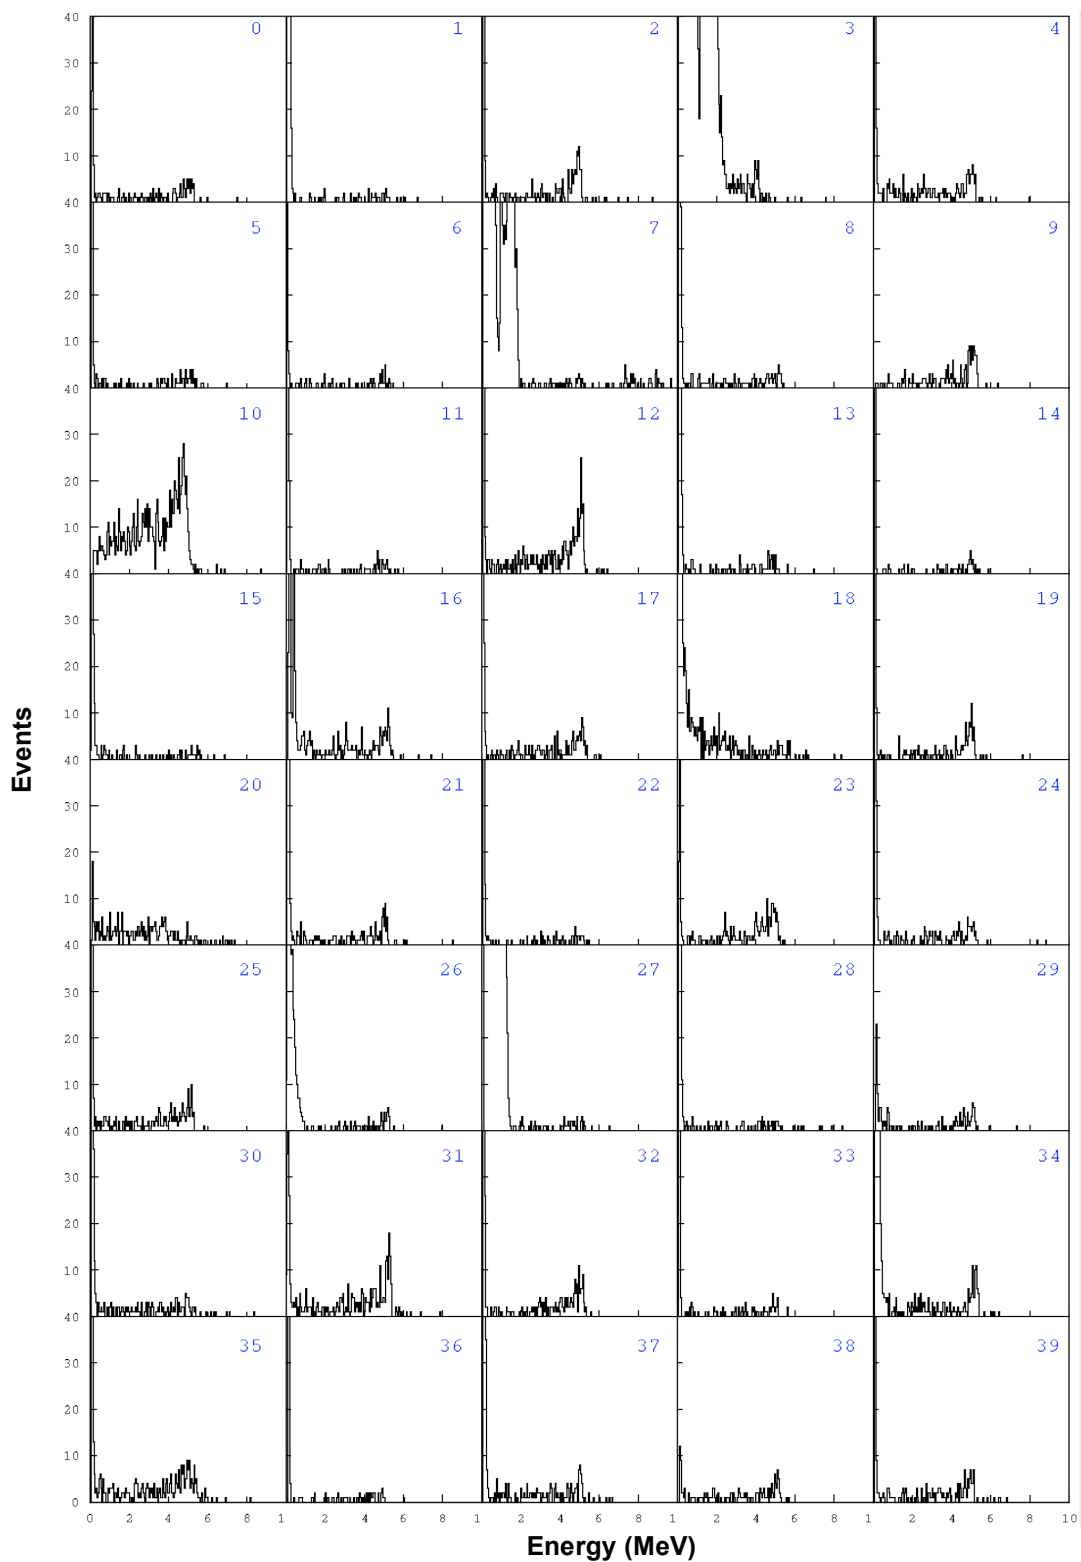


Figure 5-22: Cleaned ADC Spectra for all channels during open data set

In the following analysis, channels 3(I7), 7(M7), 18(K5), and 27(M3) were excluded due to non-alpha noise in the alpha energy range. Specifically, they had at least one of their 50 keV bins above 1 MeV but not above 3 MeV with more than 50 entries in the data set. These bins are too high-energy for neutrons, and alphas should have a higher occupancy in the >3 MeV bins.

Each channel has an independent alpha background that depends on the contaminant level of the nickel body, internal counter walls, feedthroughs, as well as the counter gas. As such, all analyses on expected background must put forth individual counter coincidence expectation values. The expectation and the uncertainty in the expectation must be individually applied to the results from each counter before the results are summarized for the array as a whole.

Through the duration of the 26.16-day run, all time-correlated events on single channels were captured. In total, only 2 candidate triple-events were detected, but a smaller (2 half-life) time window was placed on the data for the decays. The scope trace for each event was examined for consistency after the events were collected. If the scope trace was unclear or absent, the event was still counted.

349 meters of counter are used for the in-situ alpha analysis. Given the run duration, using the technique described in section 5.3.1, the ^{232}Th contamination is found at a level of $6.7^{+5.0}_{-2.7}$ ppt inside the nickel counter bodies. This is consistent with the cool-down measurement, but not as precise, due to smaller statistics.

Looking lower in the energy spectrum opens up the possibility of seeing a time-correlated beta decay along with an alpha. Unfortunately, despite a thorough search, no definitive correlated betas have been seen. This is partially due to the lack of convenient time-signatures for observation. The ^{238}U chain has a 162 μs half-life separation between a >1 MeV beta decay and a following 7.7 MeV alpha decay. Unfortunately, the ADC/shaper deadtime is 230 μs , so only 24% of the alphas would escape the deadtime. No such events were seen in the Open Data set.

In a similar fashion, the ^{232}Th chain has a 298-ns separation between a $\sim 2\text{MeV}$ beta followed promptly by an 8.78MeV alpha. These two pulses will be integrated together inside the counter gas as their separation is shorter than the ion drift time.

5.4 In-situ neutron background measurements

In addition to the alphas measured in the open data set, we may also examine the neutron energy region itself. First, we will look to determine the amount of non-neutron events that appear in this energy window. We will then show that the NCD hardware is capable of detecting the neutron signal from solar neutrinos above this background with an acceptable statistical accuracy. This section will demonstrate that the electronics as a whole are truly capable of measuring what they were made to measure.

The NCDs were designed to positively identify neutron events through the digitization channel. We will not present a method here that analyzes the digitized data as the event-by-event analyses are still under development. Instead, due to the

impressive capabilities of this array, we can and will present an analysis based on shaper/ADC data alone.

First off, we must determine which channels contain useful shaper/ADC values. Figure 5-23 presents the lower-energy shaper/ADC spectra for all 40 channels in the open data set. Examination of this plot, without looking for a neutron peak in individual channels, provides a method of cutting out noisy channels.

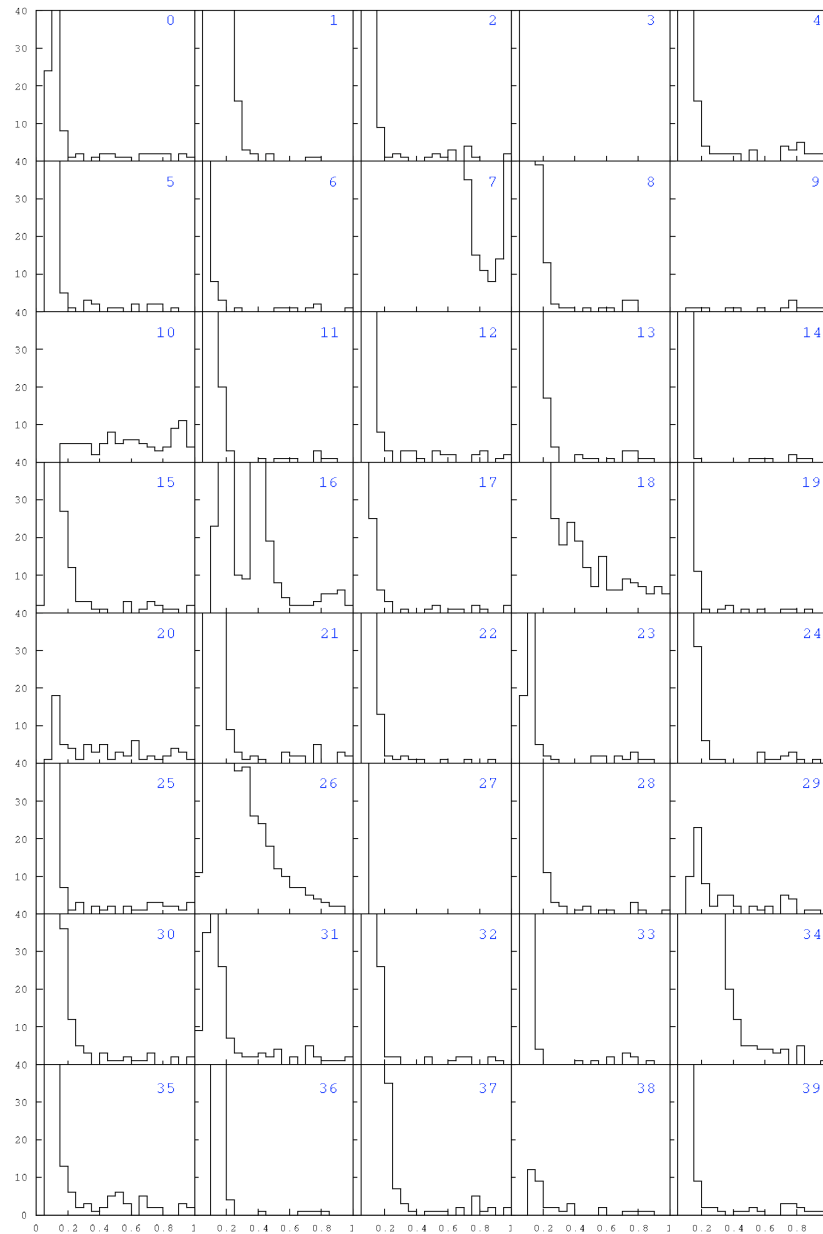


Figure 5-23: Total open data set Shaper/ADC events below 1MeV

Any channel with an excessive noise, exhibited more than one bin whose deviation from the system-wide mean has a less than 1×10^{-5} probability of being due to a pure statistical fluctuation. All energy bins between 300 keV and 6 MeV were checked, and can be seen in lower energies on Figure 5-23 and in higher energies on

Figure 5-22. Additionally, all ^4He channels are rejected, as they will not contribute to a measure of neutron flux. A summary of the channels cut is in Table 5-4.

Table 5-4: Summary of neutron run channels cut

Channel	Notes:
1(M8)	Low Energy Noise
3(I7)	Helium-4, Noisy
7(M7)	Low and High Energy Noise
10(I6)	Helium-4
16(J5)	Low Energy Noise
18(K5)	Changing Gain ⁴
20(I3)	Helium-4
26(J3)	Low Energy Noise
27(M3)	Low Energy Noise
30(I2)	Helium-4
34(L1)	Low Energy Noise

All but the cut channels are energy-calibrated and then summed together for one cumulative Shaper/ADC spectrum, as shown in Figure 5-24. The neutron peak is clearly visible and its intensity can be evaluated, which we will do after we quantify the alpha curve. It is also important to note that there is only one overflow energy value, and very little above 5.5 MeV, despite the presence of higher-energy alphas in the ^{238}U and ^{232}Th chains. This indicates that the vast majority of alpha events come from an amount of ^{210}Po that is above the secular equilibrium concentration of ^{238}U . The ^{210}Po is presumably supported by ^{210}Pb , the longest-lived (22.3 year half-life) progeny of ^{222}Rn .

⁴ Channel 18 (String K5) has seen some of its neutron peak shift over its deployment lifetime. As the neutron peak represents one of the best energy calibrations possible on the strings, the inclusion of string K5 will reduce the accuracy of the array.

5.5 Neutrons from Solar Neutrino Interactions

Using shaper/ADC data over the open data set, the neutron flux can be quantified. By excluding those few channels (listed in Table 5-4), we produce a calibrated Shaper/ADC spectrum with an obvious neutron and ^{210}Po peak, demonstrated in Figure 5-24.




Figure 5-24: Sum of 29 calibrated channel spectra over the Open Data Set. The small overlay plot is a neutron spectrum taken from the N2 string during a run in which activated ^{24}Na was circulated throughout the D_2O .

In order to verify that the peak between 700 and 800 keV is not due to non-neutron effects, we can examine the ^4He tubes, where there should be no 764 keV neutron peak. The ^4He spectrum, taken from channel(string) 10(I6), 20(I3), and 30(I2) over the open data set, is shown in Figure 5-25. Channel 3 was left out due to the high

noise on the channel. The graph on Figure 5-25 shows no indication of a peak between 700 and 800 keV, but does not have enough statistical significance to exclude one.




Figure 5-25: Calibrated ^4He counter spectra summed over the Open Data set

Unfortunately, there are only 3 usable ^4He strings, so the statistics during the open data set include a signal at with greater than 10% probability. In addition, the absence of a clear ^{210}Po peak indicates inaccuracies with matching the gains of individual counters in ^4He strings. Despite these issues, with the addition of the full data set and more than 10 times the data presented, this same plot should show a distinct lack of excess of events near the ^3He counter neutron region.

If we only use the two 50-keV bins from 700-800 keV to count neutrons, we need to quantify what fraction of the total captured neutron events are represented in those two

bins. In October of 2005, SNO uniformly distributed ^{24}Na within the D_2O , providing a homogeneous distribution of neutrons produced through photodisintegration by the 2.75-MeV γ from ^{24}Na . Neutrons captured during this period run provide an in-situ calibration spectrum for each channel. Figure 5-26 shows the spectrum from string N2 (25) during the ^{24}Na run. This calibration spectrum provides estimate of the fraction of NCD neutron events within the 700-800 keV range, adequate for the present purposes. A more detailed analysis would make use of an appropriately weighted average of the neutron spectra from all active strings.




Figure 5-26: The N2 string, chosen as a representative string, during ^{24}Na calibration runs. Figure is from [73].

From the ^{24}Na runs, the expected fraction of neutrons inside the 700 to 800 keV window is approximately 60%. If indeed, there are ^3He events in the 700-800 keV window from neutrons, we must estimate what our background due to alphas should be to separate them. Fitting the full energy spectrum of good channels between 850 keV and 2.1 MeV yields a basic expectation for alpha-contamination within the neutron-peak window with good statistical accuracy. In Figure 5-27, we fit a second-order polynomial to the plot.




Figure 5-27: Shaper ADC spectrum with 2nd degree polynomial fit to alphas between 2.5 MeV and 850 keV, prediction extrapolated over neutron peak

The fit converged on the polynomial:

$$Counts = 24.4(\pm 8.6) - 6.6(\pm 10.8)E + 3.1(\pm 3.2)E^2$$

In the two bins that effectively make up the neutron window, one expects 21 ± 2 alphas to contaminate each bin. For the total range of 700-800keV, this is 42 ± 3 expected alphas, and we measure 111 events. Given that the source of events in the neutron

window is Poisson-distributed, we arrive at 111 ± 11 events. The net difference and uncertainty from expectation is 69 ± 11 , the number of neutrons in the peak.

The total array of ^3He NCDs will capture 0.280 of the neutrons liberated in the water⁷⁴. The length of counter included in this analysis was 290.5 meters. The total length of the ^3He array is 362 meters, for a neutron counter volume of 0.802 of the full array. Approximating the efficiency to scale as the volume of counter, we expect a 0.225 capture efficiency for neutrons in the selected counters. The 100 keV-wide neutron peak between 700 and 800 keV contains approximately 60% of all of the neutrons that capture within a ^3He counter. So we expect that our measured peak contains 13.5% of the total neutron flux. That implies that the total number of neutrons is 510 ± 87 in the 26-day open data set. Since the data set used is 26.164 live days of data we conclude that 19.4 ± 3.3 neutrons are liberated in the vessel every day.

The background, without contribution from NCD contaminants, is expected to be 0.78 neutrons per day²⁰. The background due to bulk ^{232}Th and ^{238}U contamination within the nickel counter bodies is an additional 0.6 neutrons per day. This implies that 18.0 ± 3.3 neutrons are liberated due to neutral-current neutrino flux plus NCD-related backgrounds and external backgrounds daily, and that the subset of NCDs used for this analysis captures 2.43 of them in its peak.

To compare this with the expected neutrino flux as a check on the system, and not as an accurate determination of neutral current flux: 18.0 ± 3.3 neutrons per day

corresponds to a neutral current neutrino flux of 7.0 ± 1.2 million neutrinos per second per centimeter-squared.

The flux recently reported from SNO's salt phase is 4.94 ± 0.43 in the same units. The difference between the numbers is 2.1 ± 1.3 , a not unreasonable agreement at this level of analysis, but also indicative of additional NCD-related backgrounds.

Chapter 6 Conclusions

The ability to quantify flavor-independent neutrino flux allowed the Sudbury Neutrino Observatory to make fundamental advances in the understanding of neutrinos. Once it was convincingly shown that solar neutrinos were still arriving at the earth, but transformed along the way, more interest and research focused on quantifying the neutrino itself. The field of neutrino studies is in a period of great interest and expansion, with a large number of new experiments being proposed and built.

Since 2002, the NCD electronics we have described herein have been running an array of proportional counters. Once NCD data taking began in November of 2004, production data has been coming out of the electronics and data acquisition systems. The NCD array and its electronics have proven themselves quite reliable with a very high livetime fraction, 98% supernova-live and 62% neutrino-live⁷⁵, throughout the last year and a half. The performance and data quality of the NCD array is clearly very high, as has been demonstrated by the ability to observe a signal of the magnitude expected for solar neutrinos with rudimentary analysis.

6.1 Verification of NCD Electronics

Each NCD electronics subsystem has a well-quantified deadtime. The fast data path, through the Shaper/ADCs, can record events 230 μs apart, without partially formed or falsely tagged records. This allows it to take data rates into many kilohertz, which is sufficient for the NCDs alone to act as a supernova detector, should sufficient neutrons be created. The high neutron multiplicity of SNO's standardized ^{252}Cf fission neutron sources would challenge the response of slower systems. The robust stability and reliability of the Shaper/ADC system will allow them to serve in many other roles after SNO's use has waned. In particular, the ^3He array and its electronics appear well-suited as the active element in a Pb-based supernova detector.

The multiplexers have a software-dependant deadtime of less than a millisecond. Although this could be decreased still further, as the system is built to allow for more digitizers, it is sufficient for the needs of SNO.

The digitizing scopes may take slightly more than half a second to acquire and transmit an event, but the electronics that back them allow for quick consecutive events by switching between scopes. This effectively reduces deadtime to less than a millisecond per event, except in quick multiple events. The system can always be improved by updating the method of data transfer or adding more digitizers. Not coincidentally, this improvement has not been necessary for the NCDs in SNO due to extreme cleanliness and purity.

6.2 NCD array backgrounds

We have shown alpha-alpha time coincidence to be possible within the electronics through both of the major electronics channels. Shaper/ADC values can easily be summarized and presented for coincidence analysis. Digitized events may also be correlated and individually presented for examination and analysis. Individually identified alpha decays, the parent and daughter nucleus known, can be selected from the data at any time in the future, and decay sequences extending over hours can be observed.

^{238}U levels within the counters are measured to be $4.8^{+2.2}_{-1.6}$ parts per trillion, a contaminate level that would yield $0.24^{+11}_{-0.8}\%$ for the expected NC signal. ^{232}Th contamination was shown to be $5.2^{+2.6}_{-1.7}$ parts per trillion, corresponding to $2.0^{+1.0}_{-0.7}\%$ of the expected NC signal. The large fractional errors in this section are due to limited data, high cleanliness, and an uncertain contaminant distribution.

We have demonstrated that the Neutral Current Detection electronics system can detect neutrons above the backgrounds, and hence the neutrino neutral-current interaction rate in the vessel. With just over 3 weeks of data on 80% of the tubes, the neutron flux through the D_2O was measured to be 19.4 ± 3.3 neutrons per day, as compared to the expectation of 14.3 ± 1.1 neutrons per day. This analysis was done entirely separate from the event-by-event recognition provided by the digitization channel.

The current phase of the Sudbury Neutrino Observatory with the Neutral Current Detection array is expected to provide the best precision measurement yet of the mixing angle θ_{12} . The ratio of neutral-current to charged-current rates is a rather direct measure

of the value of $\sin^2(\theta_{12})\cos^4(\theta_{13})$. Given the data presented, it is reasonable to expect the accuracy of this measurement will depend primarily on the statistics of neutrino flux and runtime.

BIBLIOGRAPHY

- J. Pons et al., *Astrophysical Journal*, 513, 780 (1999)
- S. Reddy et al., *Physical Review C*, 59, 2888 (1999)
- A. Mezzacappa et al., *Astrophysics Journal*, 495, 911 (1998)
- A. Burrows et al., *Physical Review D*, 45, 3361 (1992)
- A. Burrows et al., *Astrophysical Journal*, 450, 830 (1995)
- W. Haxton, personal correspondence
- W. Haxton, unpublished
- G. Fuller, et al., *Nuclear Physics A*, 606, no. 1-2, 167-172 (1996)
- R. Hoffman et al., *Astrophysical Journal*, 460, 478 (1996)
- G. Tammann, W. Loeffler, A. Schroeder; *Astrophysical Journal*, 92, 487 (1994)
- Los Alamos Science, "Celebrating the Neutrino" (1997)
- G. Fuller et al., *Astrophysical Journal*, 389, 517 (1992)
- G. Fuller et al., *Astrophysical Journal*, 453, 792 (1995)
- W. Ormand et al., *Physics Letters B*, 308, 207 (1993)
- C. Caso et al., *European Physical Journal C3*, 1 (1998)
- D. Spergel et al., *Physics Letters B*, 200 no 3, 366 (1998)
- H. Bethe, *Astrophysical Journal* 490, 765 (1997)

- R. Bionta et al., Physical Review Letters 58 no 14 (1987)
- K. Sato et al., Physical Review Letters 58 no 25 (1987)
- K. Heeger et al., Physical Review Letters 77 no18, pp. 3720-3 (1996)
- A. Poon, Dissertation, University of British Columbia (1999)
- Q. Ahmad, Dissertation, Brown University (2002)
- R. Mohapatra et al., Massive Neutrinos in Physics and Astrophysics, World Scientific Pub. Co. Inc. (1991)
- G. Knoll, Radiation Detection and Measurement, Wiley (1979)
- J. Byrne, Neutrons, Nuclei and Matter, Institute of Physics Publishing (1995)
- ¹ W. Pauli, Address to the Physical Institute of the Swiss Federal Institute of Technology, (1930)
- ² H. Bethe, C. Critchfield, Phys. Rev., **54**, 238, (1938)
- ³ J. Bahcall et al., Astrophysical Journal Letters, 621, L85 (2005)
- ⁴ J. Formaggio et al., Phys. Rev. C, v69, n1, p15801-1-4 (2004)
- ⁵ R. Davis, Physical Review Letters, 12, 303, (1964)
- ⁶ J. Bahcall, Physical Review Letters, 12, 300, (1964)
- ⁷ R. Davis et al., Search for Neutrinos from the Sun, Phys. Rev. Lett. 20, 1205-1209 (1968)

- ⁸ B. Pontecorvo, Inverse Beta Processes and Nonconservation of Lepton Charge, JINR Preprint P-95, Dubna (1957)
- ⁹ S. Mikheyev et al., Resonance enhancement of oscillations in matter and solar neutrino spectroscopy, Sov. J. Nucl. Phys. 42 913-917 (1985)
- ¹⁰ L. Wolfenstein, Neutrino oscillations in matter, Phys. Rev. D17, 2369-2374 (1978)
- ¹¹ K. Hirata et al., Observation of ^8B solar neutrinos in the Kamiokande-II detector, Phys. Rev. Letters, v63, n1, p16-19 (1989)
- ¹² A.I. Abazov et al., Proc. Int Conf. Neutrino Physics and Astrophysics (Neutrinos 1990), Nuclear Phys. B Proc. Suppl. 19 (1991) 77.
- ¹³ P. Anselmann et al., Nuclear Phys. Proc. Suppl. 31, 117 (1993)
- ¹⁴ M. Altmann et al., Phys. Lett. B 490, 16 (2000)
- ¹⁵ Y. Fukuda et al., Evidence for Oscillation of Atmospheric Neutrinos, Phys. Rev. Lett. 81, 1562-1567 (1998)
- ¹⁶ SNO Collaboration, The Sudbury Neutrino Observatory, Nuclear Instruments and Methods in Physics Research, A449, pp 172-207 (2000)
- ¹⁷ Q. Ahmad et al., Measurement of charged current interactions produced by ^8B solar neutrinos at the Sudbury Neutrino Observatory, Phys. Rev. Lett., 87:071301 (2001)
- ¹⁸ J. Bahcall et al., The Astrophysical Journal, 555, 990-1012 (2001)
- ¹⁹ S. Fukuda et al., Phys. Rev. Lett., 86, 5651 (2001)

²⁰ Q. Ahmad et al., Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory, Phys. Rev. Lett. Vol 89 #1, 011301 (2002)

²¹ Q. Ahmad et al., Measurement of the Total Active ^8B Solar Neutrino Flux at the Sudbury Neutrino Observatory with Enhanced Neutral Current Sensitivity, arXiv:nucl-ex/0309004 (2003)

²² B. Aharmim et al., Electron energy spectra, fluxes, and day-night asymmetries of ^8B solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory, Phys. Rev. C 72, 055502 (2005)

²³ J. Bahcall et al., Solar Neutrinos Before and After KamLAND, J. High Energy Phys. JHEP02, 009, (2003)

²⁴ K. Eguchi et al., First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance, Phys. Rev. Lett., Vol 90, #2, 021802 (2003)

²⁵ T. Araki et al., Measurement of Neutrino Oscillation with KamLAND: Evidence of Spectral Distortion, Phys. Rev. Lett. 94, 081801 (2005)

²⁶ G. Fogli et al., Global Analysis of three-flavor neutrino masses and mixings, In Press, Progress in Particle and Nuclear Physics (2006) arXiv- hep-ph/0506083

²⁷ SNO Collaboration, SNO response to ARC re Completion of the SNO experiment (2005)

- ²⁸ E. Akhmedov, Proceedings of the 2004 Nobel Symposium on Neutrino Physics, *Physica Scripta* (2005); hep-ph/0412029
- ²⁹ M. Kobayashi et al., *Prog. Theor. Phys.* 49, 652 (1973)
- ³⁰ L. Chau et al., Comments on the Parameterization of the Kobayashi-Maskawa Matrix, *Phys. Rev. Lett.* 53, 1802 (1984)
- ³¹ C. Martoff et al., *Comput. Phys. Commun.* 72, 96 (1992)
- ³² IMB Collaboration, Observation of a neutrino burst in coincidence with supernova 1987A in the Large Magellanic Cloud, *Phys. Rev. Lett.* 58, 1494 (1987)
- ³³ Kamiokanda Collaboration, Observation of a neutrino burst from the supernova SN1987A, *Physical Review Letters* 58, 1490 (1987)
- ³⁴ E. Alekseev et al., Possible detection of a neutrino signal on 23 February 1987 at the Baksan underground scintillation detector, *JETP Lett.* 45, 589 (1987)
- ³⁵ H. Bethe et al., *Astrophys. J.* 295, 14 (1985)
- H. Bethe, *Rev. Mod. Phys.* 62, 801 (1990)
- A. Burrows et al., *Astrophys. J.* 450, 830 (1995)
- H. Janka et al., *Astron. Astrophys.* 306, 167 (1996)
- M. Rampp et al., *Astrophys. J. Lett.* 539, L33 (2000)
- A. Mezzacappa et al., *Phys. Rev. Lett.* 86, 1935 (2001)
- M. Liebendoerfer et al., *Phys. Rev. D* 63, 103004 (2001)

- ³⁶ Kamiokanda Collaboration, Phys. Rev. Letters, 58, 1490 (1987)
- ³⁷ IMB Collaboration, Phys. Rev. Letters, 58, 1494 (1987)
- ³⁸ E. Alexeyev et al., Detection of the neutrino signal from SN 1987A in the LMC using the INR Baksan underground scintillation telescope, Phys. Lett. B 205, 209 (1988)
- ³⁹ T. Loredó et al., Bayesian Analysis of Neutrinos Observed from Supernova SN 1987A, Phys. Rev. D 65, 063002 (2002)
- ⁴⁰ J. Heise, A Search for Supernova Neutrinos with the Sudbury Neutrino Observatory, PhD Thesis, University of British Columbia (2001)
- ⁴¹ R. Boyd. Et al., OMNIS, The Observatory for Multiflavor Neutrinos from Supernovae, proceedings of NNN99, AIP Conf Proc 533, (2000)
- ⁴² C. K. Hargrove et al., A lead astronomical neutrino detector: LAND, Astroparticle Physics 5, 183-196 (1996)
- ⁴³ Y. Giomataris et al., A Network of Neutral Current Spherical TPC's for Dedicated Supernova Detection, Phys. Lett. B. 634, 23-29 (2006)
- ⁴⁴ L. Stonehill, Deployment and Background Characterization of the Sudbury Neutrino Observatory Neutral Current Detectors, UW Dissertation (2005)
- ⁴⁵ T. Burritt, NCD internal mechanical specifications and personal correspondence.
- ⁴⁶ N. Oblath, UW Dissertation, in progress.
- ⁴⁷ M. Zhou et al., J. Am. Chem. Soc. 121, 12188 (1999)

- ⁴⁸ The NCD Analysis Group, Optimizing the Neutral Current Detector Deployment Length, SNO Collaboration internal document, November 21, 2002
- ⁴⁹ M. Howe et al., IEEE Transactions on Nuclear Science, 51, 3, 878 (2004)
- ⁵⁰ J. Boger et al., The Sudbury Neutrino Observatory, Nucl. Inst. And Meth, A449, p 172. (2000)
- ⁵¹ T. Van Wechel et al., Mk II SNO NCD Preamplifier, SNO-STR-2006-001
- ⁵² M. Browne, PhD Thesis, North Carolina State University (1999)
- ⁵³ IP408 I/O Pack, Acromag, <http://www.acromag.com/manuals/ip408.cfm> (1995)
- ⁵⁴ AVME 9660, Acromag, http://www.acromag.com/pdf/AVME9630_60.pdf (2006)
- ⁵⁵ M. Neubauer, Evidence for Electron Neutrino Flavor Change through Measurement of the ^8B Solar Neutrino Flux at the Sudbury Neutrino Observatory, Dissertation, University of Pennsylvania (2001)
- ⁵⁶ University of Pennsylvania SNO Group, The SNO Trigger System, SNO Internal Document, October 12th (1997)
- ⁵⁷ W. Diethorn, US Atomic Energy Commission, New York Operations- report NYO-6628 (1956)
- ⁵⁸ Spellman MP-series power supplies- <http://www.spellmanhv.com/pdf/MP.pdf>
- ⁵⁹ SNO Collaboration, NCD ECA Unidoc, In Progress
- ⁶⁰ R.G.H. Robertson et. al, SNO-STR-05-003 (2005)

⁶¹ M. Howe et al., NCD System Livetime and Efficiency Measurements, SNO internal document (2004)

⁶² G. Cox, NCD Electronics Model, NCD Unidoc (2005)

⁶³ SNO Collaboration, Background Measurements at Creighton Mine, SNO internal document (1995)

⁶⁴ K. M. Heeger, Model-Independent Measurement of the Neutral-Current Interaction Rate of Solar ^8B Neutrinos with Deuterium in the Sudbury Neutrino Observatory, Dissertation, University of Washington (2002)

⁶⁵ XCOM:Photon Cross Sections, NIST, physics.nist.gov/PhysRefData/Xcom/text/XCOM.html

⁶⁶ SNO Collaboration, SNOMAN User Manual

⁶⁷ CN/ASD Group, ZEBRA Users Guide, Program Library Q100, CERN (1993)

⁶⁸ W. Nelson et al., The EGS4 Code System, Stanford Linear Accelerator, SLAC-265 (1984)

⁶⁹ P. Thornewell, ^3He Neutral Current Detectors for the Sudbury Neutrino Observatory, D. Phil thesis, Oxford (1997)

⁷⁰ P. Doe et al., Construction of an Array of Neutral-current detectors for the Sudbury Neutrino Observatory, SNO-STR-95-023 (1995)

⁷¹ J. Ziegler et al., SRIM: The Stopping and Range of Ions in Matter, SRIM.com

⁷² R. Brun et al., ROOT – An Object Oriented Data Analysis Framework, AIHENP conference, Lausanne (1996)

⁷³ S. Peeters, Preliminary results of reprocessed ^{24}Na data, SNO collaboration meeting (2005)

⁷⁴ M. Di Marco, NCD Data Analysis Without P_{SA} , SNO internal document (2006)

⁷⁵ J. Heise et al., SNO ARC Committee Review (2006)

Appendix A NCD Communication Tables

Table A-1: NCD Shaper/ADC Register Definition Table

Register Number	Read Write	Base + Address	Register Function
1-Mode Select	W	1D	Selects continuous or one-shot for each channel
2-Thresh. Address	W	1B	Set individual channel thresholds, in mV at input
3-Select Channel	W	13	Select which channel is to be adjusted or read
4-Threshold Read	W	19	Select a channel-threshold to read and initiate read
5-Fast Clear	W	3	Altera reset, starts conversion on all channels
6-Register Reset	W	1	Resets all registers except 3-select channel
7-Scaler Enable	W	9	Enable/Disable selected scalers
8-Not Used		F	Not Used
9-Scaler Clear	W	B	Clears all scalers
10-Enable Discrm	W	D	Connects one channels' discriminators to Dsub Test
11-Miscellaneous	W	1F	Set on/off all scalers, Set on/off Multiboard function
12-ADC Out Ch1	R	0	12-bit conversion of channel 1
13-ADC Out Ch2	R	2	12-bit conversion of channel 2
14-ADC Out Ch3	R	4	12-bit conversion of channel 3
15-ADC Out Ch4	R	6	12-bit conversion of channel 4
16-ADC Out Ch5	R	8	12-bit conversion of channel 5
17-ADC Out Ch6	R	A	12-bit conversion of channel 6
18-ADC Out Ch7	R	C	12-bit conversion of channel 7
19-ADC Out Ch8	R	E	12-bit conversion of channel 8
20-Gain Write	W	15	Write desired 8 bit channel gain (follows reg. 3)
21-Gain Read	R	13	Read 8 bit channel gain (follows reg. 3)
22-Thresh Val 1 -4	W	7	Set threshold for selected channel (follows reg. 2)
23-Thresh Val 5 -8	W	5	Set threshold for selected channel (follows reg. 2)
24-ThreshConvert	R	15	Read threshold initiated (follows reg. 4)

Table A-2 briefly lists all of the available monitors that can be read from the NCD Shaper/ADC, ranging in function from simple board ID to current individual channel scaler values.

Table A-2: NCD Shaper/ADC Monitor List

Monitor Number	Read Writ	Base + Addre	Register Function
1-Mode Read	R	00	Reads continuous or one-shot for all channels
2-Thresh. Address	R	1B	Reads current DAC Address
3-Revision No	R	11	Reads board revision number
4-Conversion Stat	R	1F	Returns the conversion status of all channels
5-Board Scaler -I	R	30	Returns first 16-bits of full board counter
6-Board Scaler -II	R	32	Returns bit 16-28 of full board counter
<i>7-Not Used</i>		16	<i>Not Used</i>
8-Ch1 Scaler	R	20	Returns bits 0-16 of the scaler/counter for channel 1
9-Ch2 Scaler	R	22	Returns bits 0-16 of the scaler/counter for channel 2
10-Ch3 Scaler	R	24	Returns bits 0-16 of the scaler/counter for channel 3
11-Ch4 Scaler	R	26	Returns bits 0-16 of the scaler/counter for channel 4
12-Ch5 Scaler	R	28	Returns bits 0-16 of the scaler/counter for channel 5
13-Ch6 Scaler	R	2A	Returns bits 0-16 of the scaler/counter for channel 6
14-Ch7 Scaler	R	2C	Returns bits 0-16 of the scaler/counter for channel 7
15-Ch8 Scaler	R	2E	Returns bits 0-16 of the scaler/counter for channel 8

Table 2-3 maps out the 32-bit parallel communication lines used between the MUX controller card and the IP408. A brief description of each section follows.

Table A-3: IP408 IO Mapping

IO Channel	Read Write	IO/Register Definition
0	R	Controller Status Bit (CSB)
1-15	R	Data Register (DR <14...0>)
16-18	R	Controller Return Register (CRR <2...0>)
19-21	W	Write Register (WR <2...0>)
22	W	Reset Bit (RB)
23	W	Bus Status Bit (BSB)
24-31	W	Computer Data Register (CDR <7...0>)

There are 9 controller registers which defined as follows:

1. Controller Address Register, **CAR**<7...0>, from **CDR**<7...0>.

The CAR is used for addressing the DACs and ADCs that set the threshold voltages and Log amplifier input offset voltages. The 3 most significant bits **CAR**<7...5> specify the **MUX_BOX_ADDRESS**<2...0>, which selects one of eight MUX boxes. **CAR**<4...1> specifies the **CHANNEL_ADDRESS**<15...0>, within a MUX box. **CAR**<0> is not used. **CHANNEL_ADDRESS**<11...0> selects threshold voltage **THRESHOLD**<11...0>. **CHANNEL_ADDRESS**<12> selects the **LOG_OFFSET** voltage. **CHANNEL_ADDRESS**<15...13> are defined as spare channels and are not currently used.

2. Data Register, **DR**<13...0>

The 14-bit DR contains data that is to be sent from the controller to the computer, whose format depends on what is being transmitted back to the computer. In the data register are the following registers:

A. ADC Value, **DR<7...0>**

The lowest 8 bits of the DR are used by the controller to carry the ADC of the threshold value on a particular channel.

B. MUX Hardware Identifier Register, **DR<7...0>**

Each MUX board has eight jumpers, all but two of which will be connected to ground, in a unique pattern, so as to uniquely identify the connected hardware.

C. Scope Trigger Routing, **DR<9...8>**

A 2-bit scope trigger routing register which informs the user if scope A <10> or B<01> are being used exclusively, or if both <11> or neither <00> scopes are to take triggers.

D. Scope Availability, **DR<11...10>**

A 2-bit scope availability register. <11> both scopes available to <00> neither scope available. This is a response from the controller and scopes, whereas the trigger routing is a response to computer request.

E. Event Register, ER<7...0>, from **DR<7...0>**

This 8-bit register specifies which Mux boxes have received an event.

F. MUX Channel Hit Register, **DR<11...0>**

This 12-bit register shows which individual channels within a MUX box have received an event. Identified as BOX_n_MUX_TriggerED<y> on schematics.

3. Controller Status Bit, **CSB**

A 1 on this line Indicates that controller has completed a command, a 0 indicates that the controller is either ready for a command, or has not yet processed a command.

4. Controller Return Register, **CRR<2...0>**.

Return codes from controller after execution of a command., with the register explaining the state of the MUXCB as shown in Table A-4:

Table A-4: MUXCB Controller Return Register Definitions

CRR	MUXCB Definition
0 0 0	Controller Reset/MUX Rearmed
0 0 1	Channel Address Register Written
0 1 0	DAC Value Written
0 1 1	ADC Data Written to DR<7...0>
1 0 0	MUX Status (ID) written to DR<7...0>
1 0 1	Hitmap of all MUX boxes written to DR<7...0>
1 1 0	Selected MUX data (channel hitmap) written to DR<7...0>
1 1 1	Error condition encountered by MUXCB

The following commands, shown in Table A-5 may be issued to the MUXCB on the **WR<2...0>**:

Table A-5: MUXCB Write Register Definitions

WR	MUXCB Definition
0 0 0	Controller Reset/MUX Rearmed
0 0 1	Write Channel Address Value in CDR<7...1>
0 1 0	Write DAC Value
0 1 1	Read ADC Value
1 0 0	MUX Status Query
1 0 1	MUX Rearm
1 1 0	Read Event Register
1 1 1	Read Selected MUX

The function of each command is as follows:

1. Write Channel Address value. **WR<2...0> = 001.**

This command loads the CDR into the Controller Address Register (CAR). $CAR<7...1> = CDR<7...1>$. This command is used to specify the MUX box and channel address for the Write DAC value and Read ADC value commands. The controller return code for this is command is **CRR<2...0> = 001**, CAR Written.

2. Write DAC value. **WR<2...0> = 010.**

This command sets the value of the DAC at the address set in the CAR. The value in the CDR is loaded into the DAC. The return code for this command is **CRR<2...0> = 010**, DAC Write Complete. Also **DR<14...8> = CAR<7...1>** at the completion of this instruction

3. Read ADC value. **WR<2...0> = 011.**

This command tells the ADC (on the DAC output) at the address set in the CAR to convert and upon completion to load its output into **DR<7...0>**, and **CAR<7...1>** is

loaded into **DR<13...7>**. The return code for this instruction is **CAR<2...1> = 011**, ADC Data written to DR. This command waits for the ADC to finish converting, so if the ADC hangs up or is non-existent, then the command will not finish and CSB will never be set. The ADC takes approximately 3 microseconds to convert so to detect this situation the software should have a 10-microsecond timeout and issue an error message and generate a MUX Reset to clear the command and set the CSB.

4. MUX Status Query. MUX Box **CDR<7...5>** and **WR<2...0> = 100**.

This command places the identity of a MUX unit (MUXID) specified by the CDR into the DR. The **MUX_BOX_ADDRESS<2...0>** of the desired MUX unit is placed into **CDR<7...5>**. The return code for this command is **CRR<2...0> = 100**, MUX status written to DR. The selected **MUXID<7...0>** is loaded into **DR<7...0>**.

5. MUX Rearm. MUX Box **CDR<7...5>** and **WR<2...0> = 101**.

Resets (closes) all MUX gates. Sets scope trigger routing bits. **CDR<0> = Scope A trigger on**. **CDR<1> = Scope B trigger on**. (**00** = neither scope triggers, **01** = Scope A triggers, **10** = Scope B triggers, **11** = trigger toggles between Scope A and Scope B.)

6. Read Event Register, **WR<2...0> = 110**.

Gets the ‘triggered’ status (has any channel in the MUX fired since their last ‘rearm’?) of each MUX board and places it on the **DR<7...0>**. Also reads the scope availability **DR<11...10>** and triggering status **DR<9...8>**.

7. Read Selected MUX, MUX Box **CDR<7...5>** and **WR<2...0> = 111**.

Reads the ‘triggered’ status (which channels in this MUX have fired since last ‘rearm’)

8. Controller Reset, **RB<1> = 1**

Resets the controller card to default state, with CSB=0. This should not effect threshold values on the DACs board.

Table A-6: HV Controller Board Line Definitions

IO Channel	Read Write	IO/Register Definition
0	R	Controller Status Bit (CSB)
1-15	R	Data Register (DR<14...0>)
16-18	R	Controller Return Register (CRR<2...0>)
19-21	W	Write Register (WR<2...0>)
22	W	Reset Bit (RB)
23	W	Bus Status Bit (BSB)
24-31	W	Computer Data Register (CDR<7...0>)

1. Controller Address Register, **CAR<7...0>**.

The CAR is used for addressing the DACs and ADCs that control the HV coarse and fine settings, as well as read back HV system current and voltages. The 3 most significant bits **CAR<7...5>** specify the HV_FUNCTION_ADDRESS<2..0>, which selects one of four HV controller options—Coarse HV setting (default to 0), Fine HV setting (default to 1), Voltage ADC read back (default to 4), and Current ADC readback. **CAR<4...1>** specifies the HV power supply channel CHANNEL_ADDRESS<15..0>, within the controller. **CAR<0>** is not used.

2. Data Register, **DR<13...0>**.

The 14-bit DR contains data that is to be sent from the controller to the computer.

A. ADC Value, **DR<7...0>**

The lowest 8 bits of the DR are used by the controller to carry the ADC of the voltage on a particular channel, previously determined by 1, above.

B. HV Status Register, **DR<8...0>**

1 x 9 bits. Each HV relay has a bit on this register, 0-7 corresponding with each of the eight high voltage relays in order. The 9th bit, **DR<8>**, holds the HV supply voltage state (1=on, 0=off).

3. Controller Status Bit, **CSB<1>**.

A 1 on this line indicates that controller has completed a command, see the MUXCB section above for a timing diagram of how the CSB is used.

4. Controller Return Register, **CRR<2...0>**.

Return codes from controller after execution of a command are shown as

Table A-7.

Table A-7: HV Controller Board Return Register Definitions

CRR	HVCB Definition
0 0 0	Controller Reset, HV Relays shut off
0 0 1	Channel Address Register Written to
0 1 0	DAC Write Complete
0 1 1	ADC Data Written to DR<7...0>
1 0 0	HV Relay and power status written to DR<8...0>
1 0 1	Error, this is not an expected CRR from a HVCB
1 1 0	HV Relays set as specified
1 1 1	Error condition encountered by HVCB

1. Write Channel Address value. $WR<2...0> = 001$.

This command loads the CDR into the Controller Address Register (CAR). $CAR<7...1> = CDR<7...1>$. This command is used to specify the HV function and channel address for the Write DAC value and Read ADC value commands. The controller return code for this is command is $CRR<2...0> = 001$, CAR Written.

2. Write DAC value. $WR<2...0> = 010$.

This command sets the value of the DAC at the address set in the CAR. The value in the CDR is loaded into the DAC. The return code for this command is $CRR<2...0> = 010$, DAC Write Complete. Also $DR<14...8> = CAR<7...1>$ at the completion of this instruction.

3. Read ADC value. $WR\langle 2\dots 0 \rangle = 011$.

This command tells the ADC (on the DAC output) at the address set in the CAR to convert and upon completion to load its output into $DR\langle 7\dots 0 \rangle$, and $CAR\langle 7\dots 1 \rangle$ is loaded into $DR\langle 13\dots 7 \rangle$. The return code for this instruction is $CAR\langle 2\dots 1 \rangle = 011$, ADC Data written to DR. This command waits for the ADC to finish converting, so if the ADC hangs up or is non-existent, then the command will not finish and CSB will never be set. The ADC takes approximately 3 microseconds to convert so to detect this situation the software should have a 10-microsecond timeout and issue an error message and generate a HV Controller Reset to clear the command and set the CSB.

4. Null Command. $WR\langle 2\dots 0 \rangle = 100$.

This command has no effect on the hardware.

5. Read HV Relays. $WR\langle 2\dots 0 \rangle = 101$.

Reads the status of all high voltage relays and copies them over to the $DR\langle 0\dots 7 \rangle$. The status of the power going to the HV supplies is read in and copied to $DR\langle 8 \rangle$.

6. Null Command, $WR\langle 2\dots 0 \rangle = 110$.

Runs a command that has no effect.

7. Set HV Relays, HV relays $CDR\langle 7\dots 0 \rangle$ and $WR\langle 2\dots 0 \rangle = 111$.

Sets the HV relays 'on' if the corresponding bit on the CDR is set to 1. Turns the HV relay 'off' if the corresponding CDR bit is 0. Note that if the HV Controller board reads no power back from the HV supplies, it will NOT turn on the HV relays.

8. HV Controller Reset, RB<1> = 1

Resets the controller card to default state, with CSB=0. This should not effect threshold values on the DACs board.

The NCD GTID Board contains the information offset by addresses as shown in Table A-8.

Table A-8: GTID Board Register Read Definitions

Register Number	Addr offset	Register Function
Read Board ID	<i>10</i>	Reads the GTID Board's serial ID
Read Status	<i>14</i>	Status gives trigger types, errors with status
Shaper GTID	<i>18</i>	Read Shaper GTID Register
MUX GTID	<i>1C</i>	Read MUX GTID Register
Lower Shaper Time	<i>20</i>	Returns the lowest 32 bits of Shaper Time
Upper Shaper Time	<i>24</i>	Returns the highest 23 bits of Shaper Time
Lower MUX Time	<i>28</i>	Returns the lowest 32 bits of MUX Time
Upper MUX Time	<i>2C</i>	Returns the highest 23 bits of MUX Time
Read Test Register	<i>30</i>	Reads a test register (TVW's Debug)

Writing to the GTID Board at different offsets from the base address will have different results, depending on the offset address used, as shown in Table A-9.

Table A-9: GTID Board Write Register Definitions

Register Number	Addr offset	Register Function
Register reset	00	Reset internal registers, not GTID or time
Shaper event input en	02	Enable the Shaper Event Input
Shaper busy output en	04	Enable the Shaper Busy Output
Time clock count en	06	Enable/disable Time Clock Counter
Time clock count reset	08	Reset Time Clock Counter
Shaper GT event reset	0A	Reset Shaper Status to prepare for new event
MUX GT event reset	0C	Reset MUX Status bits to prepare for new event
Reset count error	0E	Reset Count Error status bit
Test latch shaper GTID	10	Latch current GTID into Shaper GTID register
Test latch shaper time	12	Latch current time clock count into shaper time
Test latch MUX GTID	14	Latch current GTID into MUX GTID register
Test latch MUX time	16	Latch current time clock count into shaper time
Soft GTRIG	18	Simulate a GTRIG (SNO triggered event)
Soft SYNCLR	1A	Simulate a SYNCLR (reset lower 16 GTID)
Soft GTRID + SYNCLR	1C	Simulate both GTRIG and SYNCLR
Soft SYNCLR24	1E	Simulate a SYNCLR24 (reset upper 8 GTID)
NCD GT OR output en	20	Enable/disable shaper/MUX GT output
MUX-Shaper event reset	22	Reset MUX Shaper Event status bit
Load test register	30	Load the 32-bit Test Register
Load GTID counter	34	Load the 24-bit GTID counter
Load lower time counter	38	Load the lower 32 bits of the Time Counter
Load upper time counter	3C	Load the upper 24 bits of the Time Counter

Table A-10: Spellman MP3P24 pin functions

Pin	Color	HVIB pin #	Function
1-Sync	-	-	Not used
2-+Vin	Red	1	24V power to the power supply, relay
3-Vmon	Orang	2	EDM voltage monitor, 0 -10V read:0 -3kV
4-Local	e	-	Not connected to 5.
5-Control set -to	Blue	-	Connects to either 4 or 6 to specify V -set
6-Remote	Blue	-	Voltage will be controlled by pins 9,7.
7-+Vset	Green	3	0-10V corresponds to 0 -3kV output.
8-Imon	Yello	4	Current Monitor
9-Vset	Grey	5	The lower end of the differential V setting.
10-P-on (ground)	Black	6	Grounded through the HVIB.

Table A-11: NCD PDS Command Definitions

<u>Bits</u>	<u>Function</u>
0	Default 'Instruction Complete' Return signal from NCDPDOB
0-7	Potential 'Instruction Complete' channels, dependant on jumper settings
8-15	Unused
16	Strobe - assert channel map unto NCDPD System
17	Clear - clears all of the NCDPDC channels, disconnecting them
18	Lowest PDC Address bit (1= \pm 1)
19	PDC Address bit (1= \pm 2)
20	PDC Address bit (1= \pm 4)
21	Highest PDC Address bit (1= \pm 8)
22	Intended connect status of channel 0 on specified PDC
23	Intended connect status of channel 1 on specified PDC
24	Intended connect status of channel 2 on specified PDC
25	Intended connect status of channel 3 on specified PDC
26	Intended connect status of channel 4 on specified PDC
27	Intended connect status of channel 5 on specified PDC
28	Intended connect status of channel 6 on specified PDC
29	Intended connect status of channel 7 on specified PDC
30-31	Unused

Low voltage power requirements are shown in Table A-12:

Table A-12: Low-Voltage Power requirements for NCD electronics

Stand -Alone Boards	<u>Voltage</u>	<u>Current</u>	<u># in system</u>	<u>Total Current</u>
MUX/HV Controller	(+) 5 VDC	3 A	2	6.0 A
	(-) 5 VDC	0.1 A	2	0.2 A
	Iso +5 VDC	1.0 A	2	2.0 A
MUX/HV DACs/ADC	(+) 5 VDC	0.56 A	2	1.12 A
	(-) 5 VDC	0.17 A	2	0.34 A
	(+) 15 VDC	0.1 A	2	0.2 A
Multiplexer	(+) 12 VDC	1.1 A	8	8.8 A
	(-) 12 VDC	0.3 A	8	2.4 A
Preamplifier	(+) 12 VDC	0.03 A	96	2.88 A
	(-6) VDC	0.01 A	96	0.96 A
VME Boards				
VME Crate (empty)	(+)5VDC	1.5 A	1	1.5 A
617 Controller	(+)5VDC	3.5 A	1	3.5 A
620 Controller	(+)5VDC	3.2 A	0	
9660 Carrier W/1 IP408	(+)5VDC	0.3 A	1	0.3 A
MVME 147SB -1	(+)5VDC	5.7 A	1	5.7 A
MVME 167(32K)	(+)5VDC	3.6 A	0	
Shaper/ADC BOARD	(+)5VDC	2.0 A	13	26 A
	(+)12 VDC	.05 A	13	.65 A
	(-)12 VDC	.05 A	13	.65 A

Appendix B Snostream

Snostream is a real-time data monitoring tool, which attaches to as a Dispatcher client to the data stream as it flows from the builder. The hierarchical goals of Snostream, in order of highest priority to lowest, are to monitor the data stream for completeness, to examine data for integrity, to display data for real-time analysis, and to examine data for unusual events.

Data from SNO is built into logical pieces of information on the DAQ Sun machine and put into two separate data paths. Time permitting, the Builder will place the data onto a TCPIP information server called Dispatcher. The Dispatcher is a modified version of part of the ControlHost Distributed Data Handling Package made by the CASPUR collaboration. More information can be obtained at: <http://www.caspur.it>

A client which subscribes to the dispatcher can request or submit data of multiple types, depending on how it attaches to the Dispatcher server. The builder submits the data stream, in its raw C format, onto the Dispatcher server, with a tag telling the server what kind of information is contained in each packet. The Dispatcher then sends copies of each bit of data to every client which has requested that type of data. At SNO, the dispatcher connects to Snostream, Lazy, XSNOED, and a redispacher directly. The

‘Redispatcher’ copies the data stream to another dispatcher, to which non-essential clients can connect. A diagram of the data flow is shown as Figure B-1, where the italics beneath processes list their host computer names.

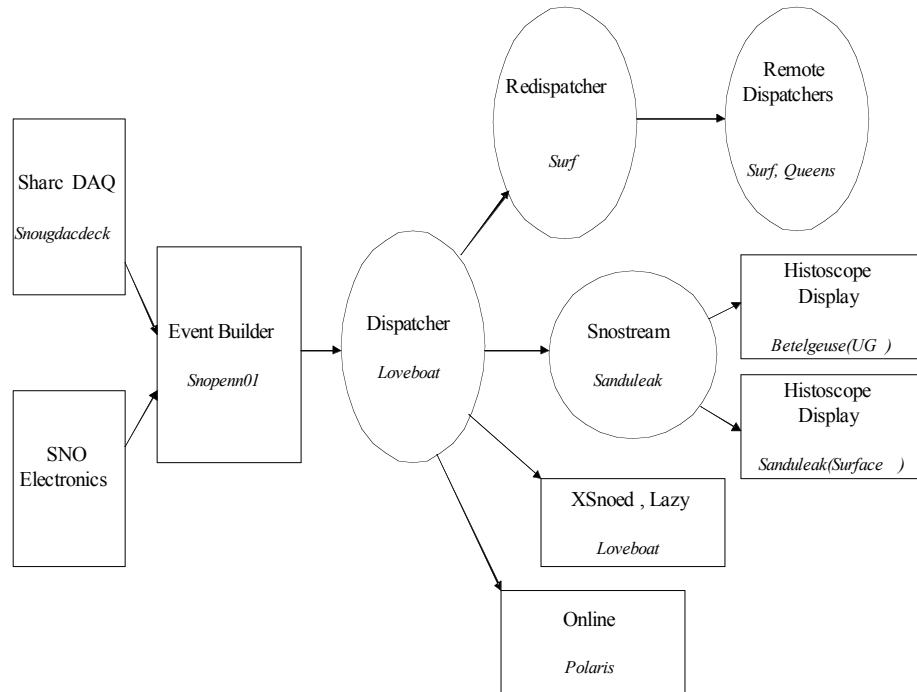


Figure B-1: SNO Live Data Flow

Snostream subscribes to CMOSDATA, RAWDATA, STATUS, PMTBNDL, CALBNDL, MONICMD, MTCBNDL, and RECHDR data types.

CMOSDATA is data returning from a SNO electronics board, carrying digitization rates for every one of the 32 channels. This data packet is of fixed length, and is sent regularly when SHARC has the ‘read scalers’ option selected. Upon completion of a full-board read of all channels, Snostream adds to the 2d-low level

histogram all of the data from that board. Later, if Snostream has spare cycles, it will sum up channels and provide more global statistics.

RAWDATA is the standard data type arriving from the detector. An event tagged as RAWDATA can actually be one of many different subtypes of information coming from the operation of the DAQ in standard running. When starting a run, the run-level information will be sent through the *Dispatcher* as a RUN_RECORD-labelled RAWDATA event. Any detector event that contains information from individual PMTs will be passed as RAWDATA with the subtype PMT_RECORD. Each PMT_RECORD carries the global event information, of standard length, immediately followed by the individual PMT data. The second part of a PMT_RECORD is of variable length determined by the number of PMTs fired. Again, as with CMOSDATA, the lowest level information is gathered from this event, and is analyzed globally when time permits.

STATUS is a primary mechanism for verifying operation of the *Dispatcher*. If Snostream is not swamped by data processing, Snostream will transmit a STATUS data packet to the *Dispatcher*, and then expects to read a STATUS event. If a STATUS event is not returned, Snostream tallies up a missed count and waits to send out another STATUS event to the *Dispatcher*. If this fails four times in a row, and Snostream has not specifically been requested to halt communication with the *Dispatcher*, a warning flag will go up, and Snostream will verify the operation of the *Dispatcher* by sending a highest-priority client list request. If the *Dispatcher* returns that Snostream is indeed connected and subscribed to the appropriate data types, Snostream will assume that the *Dispatcher* is swamped. If the *Dispatcher* does not list Snostream properly connected

and subscribed, Snostream will reconnect and subscribe. This status checking will continue, with its 'fail' counter reset.

If the STATUS data type is properly returned, Snostream will begin counting how many times it has enough free processing and network time. If the count gets to an adjustable level (3 by default), Snostream will take the time out to do some global-data calculations. If Snostream passes through the higher-level calculation loop a few times, meaning that it has a significant amount of free CPU time, it will update the display, and take action to reduce the CPU drain.

PMTBNDL and MTCBNDL data types are used for testing and debugging of the data flow only. Real data from SNO does not carry these flags, but rather, fake generated data used to verify system operation and decoding. Snostream will respond to these events like it would to the individual subtypes of RAWDATA, and results can be examined for data stream integrity verification.

MONICMD packets carry remote commands for all Snostream clients. Snostream can be suspended from the dispatcher by sending a 'SUSPEND' request inside this data, and will resume data acquisition upon receiving a 'RESUME' command. Note that a suspended Snostream client will still continue to read MONICMD commands. A Snostream client can be removed from the data stream by sending it a 'STOP' MONICMD is intended for use by other software routines, as complete remote control can be gained by attaching to any Snostream process with a Histoscope client.

CALBNDL is meant for data that has already had calibrations applied to it, or is to be used in calibrating the detector. As of this time, however, this feature is not used.

RECHDR carries the information in either a trigger or run record header. These 'header' events inform Snostream about the run that is just beginning, from type of run to electronics settings.

Snostream uses Curses emulation for its non-graphical terminal-style interface with the user. The curses display screen shows the collective sums of errors and data integrity problems, as well as the last crate/board/channel to send such data. The curses interface screen appears in Figure B-2:


```

-SNOSTREAM -<localhost> V 1.2 -(s)ave (l)oad (CC)lear (Q)uit

+----- Data Stream Errors -----+
| Error: Last - Cr Bo Ch Ce Total: | Error: Last - Cr Bo Ch Ce Total: |
+-----+-----+
| Anomalous TAC  -- -- -- --  ----- | Anomalous QLX  -- -- -- --  ----- |
| Anomalous QHS  -- -- -- --  ----- | Anomalous QHL  -- -- -- --  ----- |
| CGT-ES-16      -- -- -- --  ----- | CGT-ES-24      -- -- -- --  ----- |
| CMOS-ES-16     -- -- -- --  ----- | LGI Select     -- -- -- --  ----- |
| NC/CC event    -- -- -- --  ----- | Missed Count   -- -- -- --  ----- |
| Bad Crate/Bo   -- -- -- --  ----- | GTID = 0       -- -- -- --  ----- |
| Non-Seq GTID   -- -- -- --  ----- | MTC Error bit  XX XX XX XX  ----- |
+-----+-----+

+----- SNOSTream General Settings (User Specified) -----+
| QLX Minimum    ---- | QLX Maximum     ---- |
| QHS Minimum    ---- | QHS Maximum     ---- |
| QHL Minimum    ---- | QHL Maximum     ---- |
+-----+-----+

```

Figure B-2: Snostream Control Screen

Snostream uses Histoscope, a Fermilab real-time histogram and graphical display package, for event display. More information can be obtained about Histoscope at http://www-pat.fnal.gov/histo_doc/. Snostream, through Histoscope, allows for multiple display connections to attach to a single Snostream process, so that multiple users can access the same information without starting multiple processes on critical machines. The main Snostream-Histoscope control window is shown as Figure B-3.

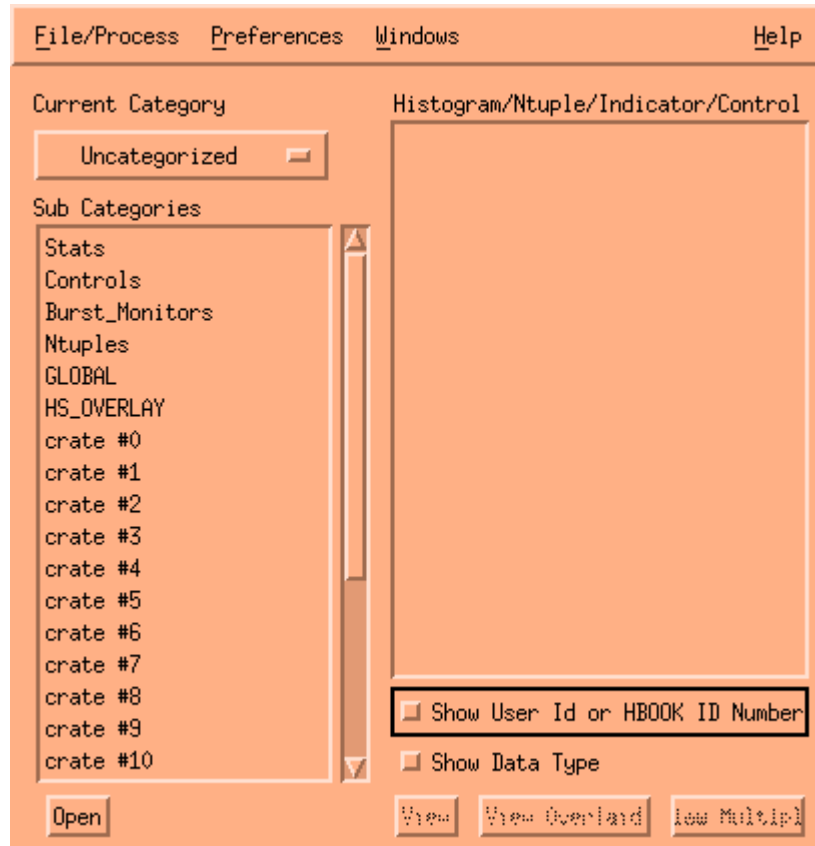


Figure B-3: Snostream's Graphical Control Window

Through Histoscope, Snostream is able to present all the information passing through the data stream in real time. Digitization rates are binned for display as both board-channel and crate-board 2-dimensional histograms. Likewise, error rates, binned into source channels, boards, and crates, can be displayed at either level. This allows for simple monitoring of one 2-dimensional histogram to spot any system-wide trends, and the user may then call up a more specific display for trend localization. A system-wide crate and board event count is shown as Figure B-4.




Figure B-4: Snostream Global Crate and Board Counts

Snostream also displays a global event rate as a function of time, using an 8 color strip chart flowing across the screen to display the history of detector-wide triggers. This provides, at a glance, the detector operator with the most important global history of the detector. Figure B-5 shows the Snostream event data passing to the left, with time listed on the horizontal axis. Snostream rebuilds this strip chart according to the data arrival times, and has a separate chart for rebuilt trigger times. If these two charts are compared, one can get a feeling for the status of the data path as well as the detector.




Figure B-5: Snostream Event Strip Chart

SNO has many more than the 7 trigger types shown on the strip charts. Extracting trigger information can be determined by examining the global trigger histogram, that sums up each trigger type individually. A short set of event triggers is shown below, in Figure B-6. A translation of trigger types is attached to Snostream display computers.




Figure B-6: Snostream Trigger Chart

Snostream allows for high-level analyses through Histoscope's use of Ntuples. Events may be summed or individually displayed, their characteristics selected for by a series of operator-determined cuts. Then, the operator can choose to display data in any manner, using histograms, 2- and 3-d scatter plots. The PMT locations in the PSUP that

are read in to Snostream at startup allow cuts and displays to be done dependant on PMT location. Up to four different data sets of Ntuples can be stored in active memory, in what Snostream loosely refers to as pseudo ntuples. Four pseudo ntuples exist in memory at any one given time, but only the active one can be effected. Sets of data may be recorded to disk, and loaded into, or on top of, the active Ntuple at any time. Cuts may be made on the pseudo-ntuples at any time, to provide for post-event cutting.

Snostream control is granted through buttons available through Histoscope. All record data may be reset, connection to the dispatcher can be manually reset, ntuple memory banks saved, loaded, or switched out. A 10-minute timed run, done at least once per shift for logging purposes, can be started at any point. Supernova warning windows can also be cleared through the use of the command buttons.

The first-level supernova detection was developed inside Snostream to provide the highest accuracy and fastest recognition without impacting the main SNO data acquisition. A supernova inside the detector would result in a high frequency of high-count PMT events over a period of many seconds, so the supernova monitoring routing looks for just that. In order to minimize resource usage, a modified Snostream records only the time of events that have the possibility of being a supernova neutrino. As each event time is logged into the 30-event-long circular buffer, Snostream compares the times of the current event and the last event stored in the buffer. If the difference is less than nominal 3 seconds, the supernova monitor informs the operator and initiates higher-level analyses.

Appendix C MUX Schematic

The NCD Multiplexer trigger schematic is shown as Figure C-1.

VITA

Charles Duba was born in Canberra, Australia due to the timing of a sabbatical. Although primarily raised between the glowing golden hills of Livermore, California, he spent a few critical years getting schooled in Germany and the Netherlands. The University of California at San Diego awarded him his Bachelor of Science degree in Physics in 1994. At the University of Washington in Seattle, he earned his Master of Science in Physics in 1996. In 2001, he became full-time faculty at the Digipen Institute of Technology, and he hopes to earn his Doctor of Philosophy in Physics from the University of Washington before the year's end.

-
- ¹ W. Pauli, Address to the Physical Institute of the Swiss Federal Institute of Technology, (1930)
- ² H. Bethe, C. Critchfield, Phys. Rev., **54**, 238, (1938)
- ³ J. Bahcall et al., Astrophysical Journal Letters, 621, L85 (2005)
- ⁴ J. Formaggio et al., Phys. Rev. C, v69, n1, p15801-1-4 (2004)
- ⁵ R. Davis, Physical Review Letters, 12, 303, (1964)
- ⁶ J. Bahcall, Physical Review Letters, 12, 300, (1964)
- ⁷ R. Davis et al., Search for Neutrinos from the Sun, Phys. Rev. Lett. 20, 1205-1209 (1968)
- ⁸ B. Pontecorvo, Inverse Beta Processes and Nonconservation of Lepton Charge, JINR Preprint P-95, Dubna (1957)
- ⁹ S. Mikheyev et al., Resonance enhancement of oscillations in matter and solar neutrino spectroscopy, Sov. J. Nucl. Phys. 42 913-917 (1985)
- ¹⁰ L. Wolfenstein, Neutrino oscillations in matter, Phys. Rev. D17, 2369-2374 (1978)
- ¹¹ K. Hirata et al., Observation of ^8B solar neutrinos in the Kamiokande-II detector, Phys. Rev. Letters, v63, n1, p16-19 (1989)

-
- ¹² A.I. Abazov et al., Proc. Int Conf. Neutrino Physics and Astrophysics (Neutrinos 1990), Nuclear Phys. B Proc. Suppl. 19 (1991) 77.
- ¹³ P. Anselmann et al., Nuclear Phys. Proc. Suppl. 31, 117 (1993)
- ¹⁴ M. Altmann et al., Phys. Lett. B 490, 16 (2000)
- ¹⁵ Y. Fukuda et al., Evidence for Oscillation of Atmospheric Neutrinos, Phys. Rev. Lett. 81, 1562-1567 (1998)
- ¹⁶ SNO Collaboration, The Sudbury Neutrino Observatory, Nuclear Instruments and Methods in Physics Research, A449, pp 172-207 (2000)
- ¹⁷ Q. Ahmad et al., Measurement of charged current interactions produced by ^8B solar neutrinos at the Sudbury Neutrino Observatory, Phys. Rev. Lett., 87:071301 (2001)
- ¹⁸ J. Bahcall et al., The Astrophysical Journal, 555, 990-1012 (2001)
- ¹⁹ S. Fukuda et al., Phys. Rev. Lett., 86, 5651 (2001)
- ²⁰ Q. Ahmad et al., Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory, Phys. Rev. Lett. Vol 89 #1, 011301 (2002)
- ²¹ Q. Ahmad et al., Measurement of the Total Active ^8B Solar Neutrino Flux at the Sudbury Neutrino Observatory with Enhanced Neutral Current Sensitivity, arXiv:nucl-ex/0309004 (2003)

²² B. Aharmim et al., Electron energy spectra, fluxes, and day-night asymmetries of ^8B solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory, *Phys. Rev. C* 72, 055502 (2005)

²³ J. Bahcall et al., Solar Neutrinos Before and After KamLAND, *J. High Energy Phys.* JHEP02, 009, (2003)

²⁴ K. Eguchi et al., First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance, *Phys. Rev. Lett.*, Vol 90, #2, 021802 (2003)

²⁵ T. Araki et al., Measurement of Neutrino Oscillation with KamLAND: Evidence of Spectral Distortion, *Phys. Rev. Lett.* 94, 081801 (2005)

²⁶ G. Fogli et al., Global Analysis of three-flavor neutrino masses and mixings, In Press, *Progress in Particle and Nuclear Physics* (2006) arXiv- hep-ph/0506083

²⁷ SNO Collaboration, SNO response to ARC re Completion of the SNO experiment (2005)

²⁸ E. Akhmedov, Proceedings of the 2004 Nobel Symposium on Neutrino Physics, *Physica Scripta* (2005); hep-ph/0412029

²⁹ M. Kobayashi et al., *Prog. Theor. Phys.* 49, 652 (1973)

³⁰ L. Chau et al., Comments on the Parameterization of the Kobayashi-Maskawa Matrix, *Phys. Rev. Lett.* 53, 1802 (1984)

³¹ C. Martoff et al., *Comput. Phys. Commun.* 72, 96 (1992)

³² IMB Collaboration, Observation of a neutrino burst in coincidence with supernova 1987A in the Large Magellanic Cloud, *Phys. Rev. Lett.* 58, 1494 (1987)

³³ Kamiokanda Collaboration, Observation of a neutrino burst from the supernova SN1987A, *Physical Review Letters* 58, 1490 (1987)

³⁴ E. Alekseev et al., Possible detection of a neutrino signal on 23 February 1987 at the Baksan underground scintillation detector, *JETP Lett.* 45, 589 (1987)

³⁵ H. Bethe et al., *Astrophys. J.* 295, 14 (1985)

H. Bethe, *Rev. Mod. Phys.* 62, 801 (1990)

A. Burrows et al., *Astrophys. J.* 450, 830 (1995)

H. Janka et al., *Astron. Astrophys.* 306, 167 (1996)

M. Rampp et al., *Astrophys. J. Lett.* 539, L33 (2000)

A. Mezzacappa et al., *Phys. Rev. Lett.* 86, 1935 (2001)

M. Liebendoerfer et al., *Phys. Rev. D* 63, 103004 (2001)

³⁶ Kamiokanda Collaboration, *Phys. Rev. Letters*, 58, 1490 (1987)

³⁷ IMB Collaboration, *Phys. Rev. Letters*, 58, 1494 (1987)

³⁸ E. Alexeyev et al., Detection of the neutrino signal from SN 1987A in the LMC using the INR Baksan underground scintillation telescope, *Phys. Lett. B* 205, 209 (1988)

-
- ³⁹ T. Loredo et al., Bayesian Analysis of Neutrinos Observed from Supernova SN 1987A, *Phys. Rev. D* 65, 063002 (2002)
- ⁴⁰ J. Heise, A Search for Supernova Neutrinos with the Sudbury Neutrino Observatory, PhD Thesis, University of British Columbia (2001)
- ⁴¹ R. Boyd. Et al., OMNIS, The Observatory for Multiflavor Neutrinos from Supernovae, proceedings of NNN99, AIP Conf Proc 533, (2000)
- ⁴² C. K. Hargrove et al., A lead astronomical neutrino detector: LAND, *Astroparticle Physics* 5, 183-196 (1996)
- ⁴³ Y. Giomataris et al., A Network of Neutral Current Spherical TPC's for Dedicated Supernova Detection, *Phys. Lett. B.* 634, 23-29 (2006)
- ⁴⁴ L. Stonehill, Deployment and Background Characterization of the Sudbury Neutrino Observatory Neutral Current Detectors, UW Dissertation (2005)
- ⁴⁵ T. Burritt, NCD internal mechanical specifications and personal correspondence.
- ⁴⁶ N. Oblath, UW Dissertation, in progress.
- ⁴⁷ M. Zhou et al., *J. Am. Chem. Soc.* 121, 12188 (1999)
- ⁴⁸ The NCD Analysis Group, Optimizing the Neutral Current Detector Deployment Length, SNO Collaboration internal document, November 21, 2002
- ⁴⁹ M. Howe et al., *IEEE Transactions on Nuclear Science*, 51, 3, 878 (2004)

-
- ⁵⁰ J. Boger et al., The Sudbury Neutrino Observatory, Nucl. Inst. And Meth, A449, p 172. (2000)
- ⁵¹ T. Van Wechel et al., Mk II SNO NCD Preamplifier, SNO-STR-2006-001
- ⁵² M. Browne, PhD Thesis, North Carolina State University (1999)
- ⁵³ IP408 I/O Pack, Acromag, <http://www.acromag.com/manuals/ip408.cfm> (1995)
- ⁵⁴ AVME 9660, Acromag, http://www.acromag.com/pdf/AVME9630_60.pdf (2006)
- ⁵⁵ M. Neubauer, Evidence for Electron Neutrino Flavor Change through Measurement of the ^8B Solar Neutrino Flux at the Sudbury Neutrino Observatory, Dissertation, University of Pennsylvania (2001)
- ⁵⁶ University of Pennsylvania SNO Group, The SNO Trigger System, SNO Internal Document, October 12th (1997)
- ⁵⁷ W. Diethorn, US Atomic Energy Commission, New York Operations- report NYO-6628 (1956)
- ⁵⁸ Spellman MP-series power supplies- <http://www.spellmanhv.com/pdf/MP.pdf>
- ⁵⁹ SNO Collaboration, NCD ECA Unidoc, In Progress
- ⁶⁰ R.G.H. Robertson et. al, SNO-STR-05-003 (2005)

⁶¹ M. Howe et al., NCD System Livetime and Efficiency Measurements, SNO internal document (2004)

⁶² G. Cox, NCD Electronics Model, NCD Unidoc (2005)

⁶³ SNO Collaboration, Background Measurements at Creighton Mine, SNO internal document (1995)

⁶⁴ K. M. Heeger, Model-Independent Measurement of the Neutral-Current Interaction Rate of Solar ^8B Neutrinos with Deuterium in the Sudbury Neutrino Observatory, Dissertation, University of Washington (2002)

⁶⁵ XCOM:Photon Cross Sections, NIST, physics.nist.gov/PhysRefData/Xcom/text/XCOM.html

⁶⁶ SNO Collaboration, SNOMAN User Manual

⁶⁷ CN/ASD Group, ZEBRA Users Guide, Program Library Q100, CERN (1993)

⁶⁸ W. Nelson et al., The EGS4 Code System, Stanford Linear Accelerator, SLAC-265 (1984)

⁶⁹ P. Thornewell, ^3He Neutral Current Detectors for the Sudbury Neutrino Observatory, D. Phil thesis, Oxford (1997)

⁷⁰ P. Doe et al., Construction of an Array of Neutral-current detectors for the Sudbury Neutrino Observatory, SNO-STR-95-023 (1995)

⁷¹ J. Ziegler et al., SRIM: The Stopping and Range of Ions in Matter, SRIM.com

⁷² R. Brun et al., ROOT – An Object Oriented Data Analysis Framework, AIHENP conference, Lausanne (1996)

⁷³ S. Peeters, Preliminary results of reprocessed ^{24}Na data, SNO collaboration meeting (2005)

⁷⁴ M. Di Marco, NCD Data Analysis Without P_{SA} , SNO internal document (2006)

⁷⁵ J. Heise et al., SNO ARC Committee Review (2006)