FIRST MEASUREMENT OF THE FLUX OF SOLAR NEUTRINOS FROM THE SUN AT THE SUDBURY NEUTRINO OBSERVATORY

Peter Wittich

A DISSERTATION

in

Physics and Astronomy

Presented to the Faculties of the University of Pennsylvania in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy

2000

Eugene W. Beier

Supervisor of Dissertation

Randall D. Kamien

Graduate Group Chairperson

To my mother

Acknowledgements

I would like to thank the many people who taught me an immense amount and who helped me get through my time in graduate school. The SNO group at Penn is an extraordinary place to work and learn. My advisor Gene Beier deserves my thanks for the guidance he gave me. Doug Cowen was a great source of ideas and made working at Penn more pleasant. Bill Frati was always willing to listen to my thoughts and I thank him for the times we spent batting ideas about.

Richard Van de Water and Josh Klein deserve my strongest thanks and gratitude for their years of guidance and instruction, from the first days in the basement of DRL to the final throes of this thesis. Mark Neubauer and Doug McDonald, the two SNO graduate students who suffered with me through years of AV construction delays and the darkest days of the PMT connector, were great companions in our DRL laboratory and great people off of whom to bounce ideas. While Ignacio Taboada was only on SNO for a short period of time, he helped create the good intellectual atmosphere that define the graduate student laboratory for me, and has been as important a person in the office as my SNO comrades. Both the post-docs and the graduate students made life in Sudbury and commissioning the detector a lot more fun, including late nights at Eddie's and Buzzy Brown's and early mornings going underground on the cage.

Working with the Penn High Energy Physics Engineering Group taught me how to work in a lab and how to think critically about making measurements. Rick Van Berg, in particular, always took time to answer my endless stream of questions, even when they strayed off-topic. Mitch taught me about electronics, but more abstractly, about thoroughness in problem-solving. Paul Keener taught me about computing and, in doing that, methods for approaching problems that I found to be more generally useful. Godwin Mayers, Chuck Alexander, Jim Cook, and Ron Pierce always came through in a pinch and most always, even when faced with the impossible, with a smile or at most a quickly forgotten frown. I would like to thank all the members of the HEP Engineering group for making my time in graduate school more pleasant and more efficient.

Outside of DRL and physics, Mac Stetzer suffered with me as a roommate for many years and confirmed that I was not going crazy when the graduate school doldrums hit. My dear friends Andreas Nahr and Christine Chung listened to my endless moans and never told me to stop complaining, even when I surely deserved it. My sister and father's confidence in me never wavered, even in times where mine was at new lows. Most of all, finally, I'd like to thank Laura for giving me the strength to finish the degree and for being there for me and with me.

ABSTRACT

FIRST MEASUREMENT OF THE FLUX OF SOLAR NEUTRINOS FROM THE SUN AT THE SUDBURY NEUTRINO OBSERVATORY

Peter Wittich

Eugene W. Beier

The Sudbury Neutrino Observatory (SNO) is a second generation solar neutrino detector. SNO is the first experiment that is able to measure both the electron neutrino flux and a flavor-blind flux of all active neutrino types, allowing a model-independent determination if the deficit of solar neutrinos known as the solar neutrino problem is due to neutrino oscillation. The Sudbury Neutrino Observatory started taking production data in November, 1999. A measurement of the charged current rate will be the first indication if SNO too sees a suppression of the solar neutrino signal relative to the theoretical predictions. Such a confirmation is the first step in SNO's ambitious science program. In this thesis, we present evidence that SNO is seeing solar neutrinos and a preliminary ratio of the measured *vs* predicted rate of electrons as induced by ^{8}B neutrinos in the $\nu_e + d \rightarrow p + p + e$ charged-current (CC) reaction.

Contents

1 Introduction **1**

1.1 Solar Neutrino Problem . . . 1

1.1.1 The Evidence . . . 3

1.1.2 Why the SNP is interesting . . . 6

1.2 Sudbury Neutrino Observatory (SNO) . . . 9

1.3 My contribution to SNO . . . 11

1.4 Reader's Guide . . . 13

2 Physics of the Standard Model, Standard Solar Model, and beyond **15**

2.1 Electroweak Physics . . . 15

2.1.1 Interactions . . . 16

2.1.2 ν mass in electroweak physics and physics beyond the standard model . . . 17

2.2 Neutrino oscillation . . . 19

2.2.1 Vacuum Oscillations . . . 19

2.2.2 Matter-enhanced oscillations . . . 22

2.2.3 MSW resonance . . . 23

2.3 Standard Solar model . . . 28

2.4 In conclusion 32

3 SNO Signals and Backgrounds 33

3.1 Strategy to Test the Oscillation Hypothesis 33

3.2 How SNO Implements this Strategy 35

3.3 Backgrounds 36

3.4 Detecting the neutrinos 37

3.4.1 CC and ES 38

3.4.2 NC 41

3.4.3 Summary 47

4 The SNO detector 48

4.1 Overview 48

4.2 The acrylic vessel 51

4.3 The photo-tubes 51

4.4 Data Acquisition hardware and software 53

4.4.1 Trigger data path 55

4.4.2 PMT data path 56

4.4.3 Data acquisition software 57

4.5 Water systems 57

4.5.1 Light Water System 58

4.5.2 Heavy Water System 58

4.6 SNO calibrations 59

4.6.1 Electronics calibration 59

4.6.2 Optical calibration 63

4.6.3 Energy response 64

4.7 SNOMAN: the SNO Monte Carlo package 64

4.8 Summary 65

5 Data flow **67**

5.1 Physics goals motivate electronics design 67

5.1.1 Solar neutrinos 68

5.1.2 Supernova neutrinos 69

5.2 An event at SNO 69

5.2.1 SNO DAQ Hardware Overview 70

5.2.2 The event 72

5.3 Near-line data analysis and automated processing 77

5.3.1 List of SAS tasks 77

5.3.2 SunDAQ: Electronics Commissioning Software 78

5.4 Summary 79

6 First analysis of SNO data **80**

6.1 Background Estimation 80

6.1.1 Instrumental backgrounds 81

6.1.2 Data cleaning 81

6.1.3 Types of Instrumental backgrounds 83

6.1.4 Other Cuts 87

6.1.5 Biases in Data cleaning 89

6.1.6 Default Mask and Data Reduction 94

6.1.7 Instrumental Backgrounds Summary 102

6.1.8 Physics Backgrounds 102

6.2 Energy Scale and Resolution near threshold 106

6.2.1 Energy calibration 107
6.3 Event Reconstruction 117
6.3.1 Position resolution 118
6.3.2 Angular resolution 124
6.4 Data Set and Run Selection 130
6.5 In Conclusion 132

7 Physics Results **133**
7.1 Data characteristics 133
7.1.1 Angular distribution 134
7.1.2 r^3 distribution 150
7.1.3 N_{Hit} distribution 156
7.1.4 Data sanity checks 160
7.1.5 The data 162
7.2 Charged Current flux 162
7.2.1 Outline of Method 162
7.2.2 Fits to Data and MC distributions 172
7.2.3 Systematic Errors 175
7.3 Conclusion 181

A Sequencer State machine **182**
A.1 Purpose 182
A.2 Physical Implementation 183
A.3 Description of functionality 183
A.3.1 Finding active channels 186
A.3.2 Processing an active channel 188

A.3.3 Control functionality . . . 191
A.3.4 Test functionality . . . 191

B The FIFO controller **194**
B.1 Functionality and Purpose . . . 194
B.2 SNO FEC memory model . . . 194
B.3 Implementation . . . 196
B.3.1 Main Functionality . . . 196
B.3.2 Diagnostic functionality . . . 199

C Timing of the SNO front end **201**
C.1 Integrator and discriminator timing . . . 203
C.2 TAC and trigger timing . . . 204

List of Figures

1.1 Solar neutrino energy spectrum, pp chain 4

1.2 Neutrino fluxes as reported by various experiments. 5

2.1 Neutrino interactions in matter 22

2.2 MSW resonant enhancement of neutrino mixing the presence of matter. 24

2.3 Behavior of θ_m as a function of $A_R \equiv \frac{A}{\Delta m^2 \cos 2\theta}$. 26

2.4 ν_e survival probability as a function of energy for different neutrino oscillation schemes. 27

2.5 $\Delta m^2, \sin^2 2\theta$ overlay of all neutrino experiments. 29

2.6 Nuclear reactions of the pp chain. 31

3.1 Spectrum for Monte Carlo CC and ES electrons. 39

3.2 Convolution of the CC and ES response shapes with the ^{8}B spectral shape. 40

4.1 Cut-away of the SNO detector 49

4.2 SNO photo tube and light concentrator. 52

4.3 Block diagram of the SNO electronics. 54

4.4 Time Walk of PMT signals 61

6.1 SNO PMT tube failures, as a function of time. 90
6.2 Bias in r for data cleaning cuts. 91
6.3 Bias with N_{Hit} for Data Cleaning Cuts. 92
6.4 Inward-going bias in for Data Cleaning Cuts. 93
6.5 Raw N_{Hit} spectrum for all Set A data. 97
6.6 N_{Hit} spectrum for all Set A data with the Data Cleaning Group's flasher cuts applied. 98
6.7 N_{Hit} spectrum for all Set A data, after flashers, bursts, and events with OWL hits have been removed. 99
6.8 N_{Hit} spectrum as in Fig. 6.7, with pickup cuts applied. 100
6.9 Raw N_{Hit} spectrum compared to cleaned spectrum, with 540 cm fiducial volume cut applied. 101
6.10 High energy γ-ray estimation using the ^{16}N source 104
6.11 Number of Čerenkov photons emitted per unit length, as a function of energy. 108
6.12 MC prediction of the mean number of hit photo-tubes for a simulated electron 109
6.13 Number of hit photo-tubes for 6.1 MeV γ-rays and electrons, as predicted by SNOMAN. 112
6.14 Decay of ^{16}N 114
6.15 Mean response of SNO to ^{16}N γ-rays as a function of radial position. 116
6.16 Position resolution of the FTE(G) fitter 122
6.17 Comparison between data and MC ^{16}N position reconstruction accuracy. 123
6.18 Determining the angular resolution using the ^{16}N 124
6.19 Overlay of $\cos\theta_{\mathrm{fit}}$ for the FTP, FTE(G) and FTT fitters. 126

6.20 Angular resolution, as measured by the FTE(G) fitter. An additional cut 100 $cm < r_{\mathrm{fit}} < 200$ cm was applied to minimize the effects of misreconstruction. We have included a third exponential to model the backward peaking, though this has little effect on the actual resolution measurement. 128
6.21 Angular resolution check with the SNOMAN MC, full source ^{16}N run. 129

7.1 $\cos\theta_{\odot}$ distribution. 135
7.2 $\cos\theta_{\odot}$, with $r < 750$ cm. 136
7.3 Alternative fits to $\cos\theta_{\odot}$ distribution. 138
7.4 Monte Carlo SSM Simulation of 80 days of neutrino data 139
7.5 $\cos\theta_{\odot}$ in the light water for this data set. 140
7.6 Distribution of $\cos\theta_{\odot}$ for neck events from the data cleaning sample. . 142
7.7 Event times in hours for the neck data sample shown in Fig. 7.6. Times listed are in UTC, so the big peak at 3 corresponds to 10 PM Sudbury time. 143
7.8 Distribution of $\cos\theta_{\odot}$ for neck events and their times in hours. 144
7.9 Z direction cosine of reconstructed neck events from tape FD0000. . . 145
7.10 $\cos z$ for neutrino data sets A & B. 146
7.11 $\cos\theta_{\odot}$ distribution generated using the FTP Path fitter. 147
7.12 $\cos\theta_{\odot}$ for changed detector coordinates to off coast of Newfoundland. 148
7.13 $\cos\theta_{\odot}$ for changed detector coordinates to Berlin, Germany. 149
7.14 Distribution of reconstructed radial position in units $(\frac{r_{\mathrm{fit}}}{r_{\mathrm{AV}}})^3$ for events in the cleaned data set. An N_{Hit} cut of 45 tubes is included in this plot. 151

7.15 Radial distribution of reconstructed events for a simulation of neutron captures on deuteron. 152
7.16 Simulation of reconstructed radial distribution of events for electrons. 153
7.17 MC r^3 distribution with N_{Hit} cuts. 155
7.18 Effect of concentrator cutoff *vs* rapidly rising backgrounds 156
7.19 r^3 distribution with an additional cut $N_{\mathrm{Hit}} > 60$. 157
7.20 Overlay of data and MC N_{Hit} spectra. 158
7.21 Comparison of equal volume slices in the central part of the detector *vs* a slice in the light water. 159
7.22 Remaining DAMN bits in the selected, reconstructed data. 160
7.23 Distribution of reconstructed hits in $\cos\theta$ and ϕ. 161
7.24 ^{16}N data-MC scaling. 164
7.25 Fraction $\mathcal{F}$ of the spectrum above a threshold N_T. In the region around 60 hit PMTs, the total measured flux changes by 2.1% for each unit increase in N_{Hit} threshold. 165
7.26 Fitted Neutron contribution to the N_{Hit} spectrum. The shaded spectrum is a simulation of $CC - ES$-induced electrons, and the points are the data. The usual cuts were applied to the data and simulation. We obtained the scale factor by a χ^2 fit of the data and simulation at $N_T > 70$, where the neutron contamination is known to be small. The estimated number of neutrons based on the spectral differences below 70 hits is 105.9 ± 17.3. 167
7.27 Angular resolution of the CC spectrum. 169
7.28 Angular resolution of the ES spectrum. 170
7.29 Final $\cos\theta_{\odot}$ distribution 173

7.30 Monte Carlo $\cos\theta_{\odot}$ distribution for 1285 live days. 174

A.1 SNO data bitmap for PMT data 183

A.2 Schematic description of sequencer timing 184

A.3 A typical example of the non-linear response of the MAXIM 120 12 bit ADCs. 192

B.1 FIFO memory layout 197

C.1 Timing diagram for the SNO front end. Representative values of the timing are given in Tab. C.1. 202

List of Tables

3.1 Capture percentages of neutrons by isotope, for pure D_2O, for salt, and for NCDs. 42

4.1 Types of energy scaling sources in SNO for electrons 64

5.1 Size of the SNO data stream 68

6.1 Header information used to select runs. 131

6.2 Live time and time cut by data cleaning cuts. 132

7.1 Live time and time cut by data cleaning cuts, without tape FD0014. . 134

7.2 ^{16}N Data/MC scalings at various positions. 163

7.3 Backgrounds subtracted from the CC data total counts. 173

7.4 Systematic errors for the CC rate ratio. 177

A.1 PMT data bundle definitions 185

C.1 Representative values for SNO FE timing 201

Chapter 1

Introduction

The Sudbury Neutrino Observatory has been built to try to resolve a problem that has vexed particle and nuclear physicists for the last 30 years: the solar neutrino problem. This name is given to the apparent deficit between the predicted and measured fluxes of neutrinos emitted by the sun. SNO will test the currently favored solution to the solar neutrino problem, neutrino oscillations. If it turns out that neutrino oscillations explain the deficit, SNO will have helped solve a riddle that has taught scientists about both the sun and other main sequence stars and about neutrinos, arguably the least well understood of all fundamental particles.

1.1 Solar Neutrino Problem

In 1970, Ray Davis and his collaborators set out to measure the neutrino flux from the sun. They did so to verify the Standard Solar Model (SSM) of John Bahcall and Roger Uhlrich [1], a set of theoretical calculations that attempt to describe the nuclear fusion processes that power the sun. The SSM seeks to predict the sun's measurable characteristics by modeling the nuclear interactions at the center which fuel the sun. These fusion processes produce a stream of electron neutrinos deep at the core of the sun, where the fusion takes place. The neutrino flux allows a direct measurement of the processes in the center of the sun. The light that is generated in the core region of the sun takes $\sim 10^4$ years for photons to reach the outer edge of the sun and hence the surface of last scattering is near the outside. Thus, the light from the sun typically reflects the physics of the surface of the sun. In contrast, the fusion processes at the center also produce neutrinos which stream out of the sun unchanged. Light from the sun provides the experimenter with information about processes near the surface of the sun, but little information about what goes on deep inside the sun. Neutrinos, on the other hand, give information about the processes at the center of the sun. Ray Davis hoped to test the solar model's predictions of the processes at the center by measuring the sun's electron neutrino flux.

The experiment Davis and his collaborators undertook was ambitious. In a gold mine in Homestake, South Dakota, they built the first in what would be a long line of solar neutrino detectors. Using a large vat of chlorine-based dry cleaning fluid, the experiment detected neutrinos via inverse beta decay with chlorine.

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

A small fraction of the neutrinos streaming from the sun interact in the so-called Chlorine experiment, changing one chlorine atom into an argon atom with each interaction. At fixed intervals, the argon atoms were extracted from the chlorine vat and counted. The number of argon atoms thus counted constitute a direct measure of the sun's electron neutrino flux.

Davis and his collaborators were surprised by what they found. The measurements they made were consistently lower than what was predicted by the standard theoretical models. Initially, a simple explanation seemed most logical: either the experimental apparatus was not understood, or the theoretical calculations were wrong. In either case, physicists thought the measured solar neutrino deficit was either due to an experimental problem, a nuclear physics problem or an astrophysics problem.

For instance, on the experimental side, the extraction efficiency for counting and detecting a handful of argon atoms from a vat of chlorine might have been overestimated, leading to an apparent deficit where, in fact, none existed. The nuclear cross sections for the fusion processes might have been overestimated, leading to an overestimate of the number of expected neutrinos. Finally, the solar model might be wrong, e.g., it might predict the wrong solar core temperature. The neutrino flux measured in the Chlorine experiment is very sensitive to the solar core temperature, so an error in this prediction would lead to a false discrepancy.

1.1.1 The Evidence

Over time, both experimenters and theorists refined their work, but the discrepancy between theory and experiment remained. Fig. 1.1 shows the predicted solar neutrino spectrum of the *pp* chain, which produces the bulk of the solar neutrinos. Experimenters have used three types of experiments, the Chlorine, Gallium, and water Čerenkov, to measure the solar neutrino flux. In Fig. 1.1, we see that the energy ranges that each of the experiments can measure is quite different. All experiments have measured a deficit relative to the theoretical predictions. Fig. 1.2 shows the results of the three classes solar of neutrino experiments to date. The data points

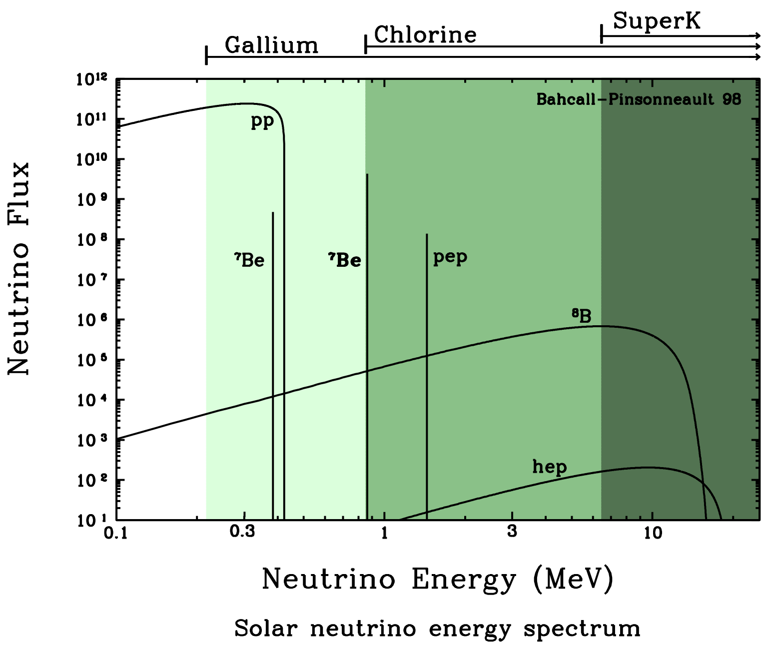

Figure 1.1: Solar neutrino spectrum, pp chain. Courtesy of John Bahcall [2]. The approximate energy thresholds of several detector types are shown on top. SNO's expected energy reach overlaps with Super Kamiokande, though the lower edge is expected to extend below 5 MeV.

represent an average value at which the experiments are sensitive. The two Gallium experiments, SAGE and Gallex, are mainly sensitive to the abundant, low-energy *pp* neutrinos. Both SAGE and GALLEX (and now, its successor, GNO) observe *pp* and *pep* neutrinos through the following reaction

$$\nu_e + {}^{71}\mathrm{Ga} \rightarrow {}^{71}\mathrm{Ge} + e^- \tag{1.2}$$

Both experiments have measured a reduced flux relative to theoretical predictions, as indicated in the Figure and in [3, 4].

In the third category, the water Čerenkov detectors, are Kamiokande and Super-Kamiokande, along with SNO. Unlike SNO, their detectors do not distinguish between

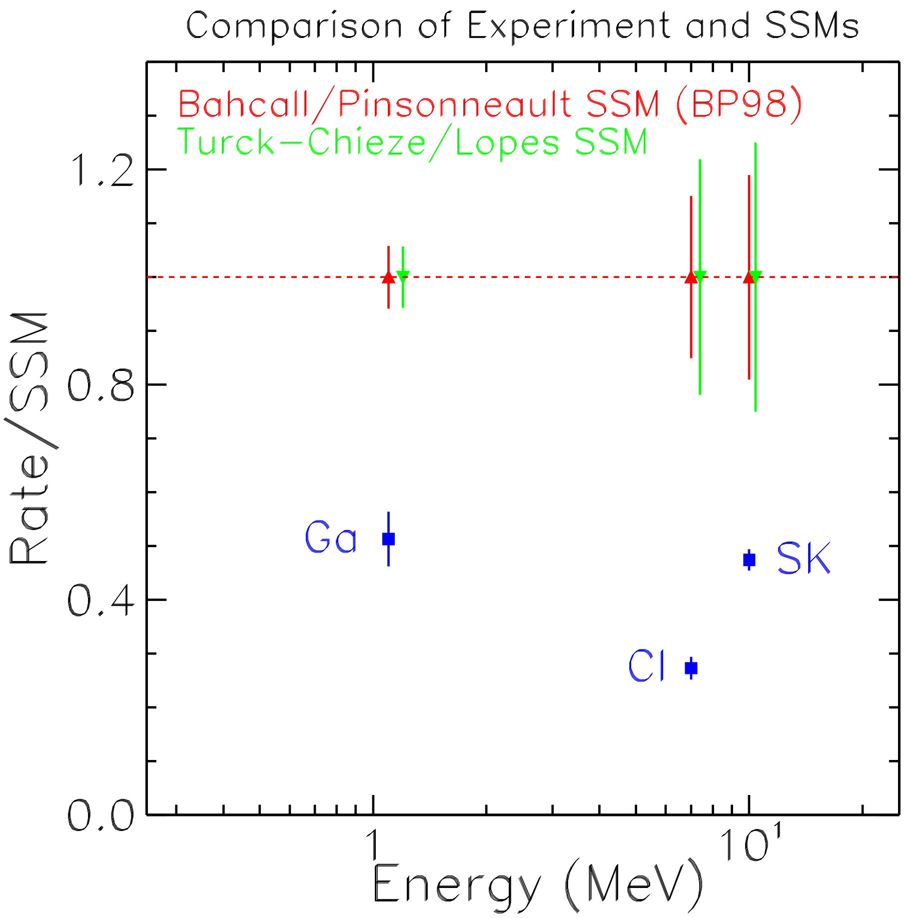

Figure 1.2: Neutrino fluxes as reported by other experiments. The value plotted on the x-axis is the average energy value at which the experiment is sensitive to solar neutrinos.

electron and other types of neutrinos. Both experiments have measured a deficit of solar neutrinos [5, 6].

In all experiments, the ratio of measured rate *vs* theoretical prediction is below 1/2, but the experiments measure a suppression that appears to depend on which part of the energy spectrum they sample.

When it became clear that neither experimental error nor a theoretical mistake could explain the discrepancy, the mismatch between theory and experiment became known as the *solar neutrino problem* (SNP.) This shift emphasized the belief that the discrepancy was a window into a more fundamental particle physics problem. Finally, Hata, Bludman and Langacker showed that an incorrect solar model *could not* account for all the discrepancies in the data [7, 8].

1.1.2 Why the SNP is interesting

At the time of Davis's initial discoveries, the extent of the solar neutrino problem was not understood. The standard solar model was in its infancy, and the Chlorine experiment itself was not well understood. These uncertainties raised the expectations that a simple solution to the solar neutrino problem would be forthcoming, either in the form of an updated solar model, or in the form of a new experimental result that more in line with the theoretical model. In time, the opposite occurred.

The standard solar model was refined and tested in other ways. The only place where it failed was the comparison with the neutrino flux prediction. A number of other theorists developed their own solar models, which agreed with Bahcall's SSM but disagreed with the experimental calculations. The SSM was constrained by another prediction, the helioseismological measurements of the vibration of the sun's surface. The oscillation modes depend on the density profile of the Sun's interior. By mapping out this profile, the nuclear fusion regions are also constrained. The experimental measurements agree very well with the SSM's predictions.

Initial expectations of achieving a quick fix by changing the astrophysical models or discovering a flaw in the experimental techniques faded when careful work showed that neither the theories nor the experiments appeared to be at fault. It seemed that the solution to the solar neutrino problem could have much wider implications than just affecting the models of stellar evolution.

When both the theorists and the experimentalists had each verified their work, without reaching an agreement, new models had to be developed. Solving the solar neutrino problem was only possible by invoking *new physics*, that is, physics beyond the standard model of particle physics (hereafter referred to as the standard model.) Earlier, Bruno Pontecorvo had suggested a model whereby neutrinos could change from one flavor to another [9]. The possibility existed that both the experimenters and the theorists were correct, but that neutrinos *oscillate*, or change from one flavor to another during their journey from the sun to the earth. If neutrinos oscillate, the solar neutrino problem is properly seen as a particle physics problem. The sun emits exclusively electron neutrinos, and the chlorine experiment is only sensitive to that flavor neutrino. However, the missing neutrinos had changed to some other undetected neutrino flavor while en route to the earth. The discrepancy between solar neutrino theory and experiment is due to our incomplete knowledge about the fundamental properties of the neutrino rather than problems in astrophysical models. The SNP's resolution could then therefore lead to new discoveries about the properties of neutrinos themselves.

Neutrino oscillation is currently the favored answer to the solar neutrino problem.

At base, oscillations can be expressed in a very simple fashion understandable to anyone with a rudimentary knowledge of quantum mechanics: the eigenstates of the neutrino wave functions in the free particle Hamiltonian are not eigenstates of the electroweak Hamiltonian. These off-diagonal terms lead to oscillations, just as the mixing between different projection of spins leads to spin precession. In the sun, neutrinos are created in fusion processes, which are governed by the weak force. Hence, neutrinos are created in electroweak eigenstates. On their path from the sun to the earth, however, the Hamiltonian governing the motion of the wave functions is the free particle Hamiltonian. The electroweak eigenstates can be projected onto the free particle wave eigenstates,

$$|\nu_e(t)\rangle = \alpha(t)|\nu_1(0)\rangle + \beta(t)|\nu_2(0)\rangle \tag{1.3}$$

The different free particle eigenstates have different eigenvalues and propagate at different velocities. Therefore, if at a time t_0 we start with one flavor eigenstate $|\nu_y\rangle$, at some later time t_1 there is a non-zero probability of finding a different neutrino flavor eigenstate $|\nu_{x \neq y}\rangle$. Even if only one flavor of neutrino is present initially, others can appear at a later time.

In this mold, neutrinos can only oscillate if at least one neutrino has a non-zero mass. This conflicts with the standard model, in which the neutrino is described as a massless particle. Additionally, the standard model posits that only the left-handed chiral projection of the neutrinos exist . If neutrinos turn out to be massive, the standard model must be amended to include massive neutrinos. Since massive neutrinos would allow you to take a left-handed neutrino into a right-handed neutrino via a change in reference frame, the new models would have to incorporate this feature.

If neutrino oscillation is discovered, all models that cannot reproduce these features must be ruled out. In this way, neutrino oscillations provides an interesting avenue into "physics beyond the standard model."

1.2 Sudbury Neutrino Observatory (SNO)

The Sudbury Neutrino Observatory [10] is a second generation solar neutrino experiment. SNO is a water Čerenkov detector that is unique in its use of heavy water as the target medium. The idea for a solar neutrino detector with a heavy water target was first proposed by Herbert Chen in 1985 as a direct approach to solving the solar neutrino problem [11].

SNO can make a *model independent* measurement of the neutrino flux for the different active neutrino flavors. The experiment is sensitive to three interactions relevant to solar neutrinos.

$$\nu_e + d \rightarrow e^- + p + p, \qquad \text{CC} \tag{1.4}$$

$$\nu_x + d \rightarrow n + p + \nu_x \qquad \text{NC} \tag{1.5}$$

$$\nu_x + e^- \rightarrow \nu_x + e^- \qquad \text{ES} \tag{1.6}$$

where $\nu_x = \nu_{e,\mu,\tau}$. Interaction (1.4) is called the *charged current* (CC) interaction. In it, the recoil electron is observed via its Čerenkov radiation. This interaction allows a direct measurement of the electron neutrino flux, unlike (1.6), in which both electron neutrinos and others are measured. In the *neutral current* (NC) interaction (1.5), all neutrino flavors e, μ, τ interact with equal strength. In the NC interaction, the free neutron is captured and indirectly detected.[1] The ability to measure the flux of all neutrinos and the flux of electron neutrinos separately allow SNO to make a model-independent analysis of the solar neutrino problem. If the ratio CC/NC of the charged current flux to the neutral current flux is not

$$CC/NC = \text{“1”}, \tag{1.7}$$

that is, the flux of electron neutrinos from the sun is not the same as the flux of all combined neutrino flavors, then neutrinos oscillate.[2] This conclusion does not depend on the particulars of any solar model. It is just a measurement of the ratio of all neutrino flavors to electron neutrino flavors coming from the sun. Note that both the CC and NC reactions are only possible due to the presence of heavy water target.

The third interaction, Eq. 1.6, is called the *elastic scattering* (ES) interaction. Čerenkov detectors with a light water target are only sensitive to this interaction. For an ES event, all neutrino flavors can participate in the elastic scattering process of a neutrino with an electron. At solar neutrino energies, the cross section for the interaction where $x = e$ is approximately six times larger than the cross section for $x = \mu, \tau$, however, and therefore, for the null hypothesis of no neutrino oscillations, measuring the number ES interactions is mainly a measure of the electron neutrino flux.

[1] SNO has three neutral current detection strategies. The simplest involves capture of n on pure D_2O. Two alternative methods, using a chlorine salt and using special ^{3}He neutral current detectors (NCD), will be described in Ch. 4.

[2] Oscillation into sterile neutrinos, i.e., neutrinos that do not participate in weak interactions, will not be detected by a CC/NC ratio. However, the suppression of the NC signal relative to the SSM prediction can be interpreted as strong evidence for oscillations into sterile neutrinos, and the energy spectrum of the CC electrons might provide an additional diagnostic.

If SNO measures a CC/NC ratio that is consistent with neutrino oscillations, the goal of the SNO collaboration is to measure the quantities that parameterize the oscillation. The type of oscillation models presently favored predict that neutrinos will oscillate in a manner dependent on their energy. As a result, the electron neutrino energy spectrum at the earth changes due to energy-dependent disappearance of ν_e . The electrons measured in the CC interaction (1.4) have an energy closely related to the energy of the incident neutrino. Therefore, any distortions of the neutrino energy spectrum will be readily apparent in the measured CC spectrum at SNO. This CC spectral distortion measurement will allow SNO physicists to determine the oscillation parameters, if neutrinos do indeed oscillate. The elastic scattering interaction (1.6) can be used similarly. In contrast, the electron energy provides one only with a lower bound of the resulting neutrino energy, due to the shape of the electron energy spectrum in elastic scattering. This makes it much harder to see oscillation-induced distortions in the ES interaction.

1.3 My contribution to SNO

My tenure with SNO encompassed the construction and commissioning phase of the experiment as well as the first solar neutrino data taking. The Penn SNO Group's main responsibility during the construction phase consisted of building the front end and trigger electronics system for the experiment. Led by the experienced hands of Rick van Berg, Mitch Newcomer and Gene Beier, a group of five post-docs and graduate students designed and built this system.

I worked closely with postdoctoral researcher Richard Van de Water and fellow graduate student Doug McDonald in developing the front end signal processing subsystem. The front end measures the signals from SNO's photo multiplier tubes (PMTs) and emits trigger primitives to form the detector trigger. Additionally, the system is highly configurable and self-calibrating.

When I started working with Penn, the front end consisted of three prototype Application-Specific Integrated Circuits (ASICs) and a prototype eight channel front end card. The scope of Penn's involvement in the SNO electronics had recently changed from supplying only two of three custom integrated circuits and no Printed Circuit Boards (PCBs) to building all PCBs for these chips. Penn's new responsibilities meant a lot of opportunities for an inexperienced student to contribute significantly. I was able to help in the digital design work, and was primarily responsible for the digital simulation of the front end card, the custom SNOBus backplane and the interface to standard VME architecture. In addition, I was able to gain valuable debugging experience with the first set of prototype chips and boards. In time, a lot of hard work produced a system of over 2000 printed circuit boards of twelve types, and I learned how to set up and complete production testing of these electronics.

At the same time, working on the rest of the system, Josh Klein, a post-doc, and Mark Neubauer, a fellow graduate student, with the help of the HEP electronics group, designed and built the trigger system for the experiment,

During the final stages of the electronics debugging, the Penn group decided that it was necessary to write its own data acquisition (DAQ) software to finish large scale testing of the electronics. The architect of the Penn DAQ effort was Paul Keener. I was able to learn from Paul and help him guide the rest of the group in this DAQ effort. The DAQ system provided me an opportunity to re-examine my electronics knowledge from the point of view of a software person. The Penn DAQ was a fully capable system, and took the first SNO PMT data. Large parts of it were incorporated into the official SNO DAQ system.

The electronics were successfully installed by September 1998. At this point, SNO was into its first data taking shifts, and I was able to be among the first to take data with this new detector. In this stage, the goal was to learn about this detector we had just built. During this time, I was able to help identify and find a number of unexpected instrumental backgrounds, such as ‘flashers,’ and help identify problems in the electronics, such as a case where failures in commercial ICs were causing the channel time-to-amplitude converters to produce discontinuities in their output.

Now, in production running, I have set up an analysis framework in order to be able to handle SNO’s incoming data stream. This framework has allowed me to make the measurement in this thesis, an early estimate of the flux of solar electron neutrinos as measured by the charged current reaction.

1.4 Reader’s Guide

This thesis begins with an introduction to the solar neutrino problem and its implications for particle physics (Ch. 1–2). The next chapter, Ch. 3, describes an abstract detector design to test the oscillation hypothesis to the SNP. A detailed description of the SNO detector follows in the next chapter, Ch. 4, which will show how SNO implements this abstract design in a concrete detector. Ch. 5 provides a detailed but non-technical description of the path of data from the PMTs to long term storage, which will allow a general description of some of the technical work I have done for SNO. The goal of this chapter is to describe my contribution to SNO and to place it in the wider detector-level context.

Next I will discuss the charged current rate measurement as the first SNO data analysis. The emphasis will be on describing what we have learned about this new detector and how the data must be handled in order to extract meaningful physics. This includes how to reduce the amount of instrumental and physics backgrounds to me to make a high threshold charged current rate measurement (Ch. 6). The methods used for event-by-event physics parameter extraction such as position fitting and energy estimation will figure prominently.

Ch. 7 will describe the results of this measurement and place it in the wider context of the solar neutrino problem. A summation will conclude the thesis. Finally, in the appendices, I will describe some of the technical aspects of the electronics for reference purposes.

Chapter 2

Physics of the Standard Model, Standard Solar Model, and beyond

In this chapter, I will briefly describe in outline form the electroweak model of particle physics as it is relevant to SNO. I will then discuss neutrino oscillations, with particular focus on matter-enhanced oscillations as these are relevant for solar neutrinos. Finally, the last part of the chapter deals with the standard solar model, which is the basis for solar neutrino physics.

2.1 Electroweak Physics

The standard model of particle physics consists of a description of the strong force and a description of the electroweak interaction. The electroweak force is a unification of the electromagnetic force of Maxwell and the weak interaction that describes nuclear decays. The first subsection outlines the features of the force itself, while the second deals with some of the issues of neutrino mass in the standard model.

2.1.1 Interactions

The weak interaction was first discovered by the observation of radioactive decay. Since then, it has developed into one of the most successful theories in particle physics, and, with the electromagnetic force, has been unified with the electromagnetic force in the electroweak theory of Weinberg, Glashow and Salaam.

The name 'weak' comes from the fact that the weak force is a lot weaker than the strong and electromagnetic forces, as evidenced by the low decay rates. A general scaling between the strengths of the forces, based on the typical decay rates of those interactions, is Strong : EM : Weak $= 1 : 10^{-2} : 10^{-6}$. The weak force is mediated by two force carriers, the neutral Z_0 and the charged $W^{\pm}$. These bosons are massive ($m_Z \sim 90\,\mathrm{GeV}, m_W \sim 80\,\mathrm{Gev}$), leading to a force of limited range.[1]

The weak force has two additional properties that are interesting. It violates parity and charge conjugation symmetry, as the force only couples to ν_L and $\overline{\nu}_R$, that is, the left-handed neutrino and the right-handed anti-neutrino. In the standard model, all particles interact via the weak force. The neutrino is unique in that it *only* interacts via the weak force, as it has no electric charge and does not carry color.

Finally, the weak force does not mix leptons of different generations, but it *does* mix quarks across generations. The latter statement is equivalent to saying that the weak interaction eigenstates are not the eigenstates of the mass matrix and the strong interaction. The mixing matrix is called the CKM matrix, and the neutrino oscillation formalism is similar to the CKM formalism.

[1] $\delta E \delta t \geq \hbar, c\delta t \sim \hbar c/\delta E, \delta E = m_W, \hbar c \simeq 197\,\mathrm{MeV\,fm}, ct \sim 2 \times 10^{-3}\,\mathrm{fm}$ as an upper limit of the interaction length of the force.

2.1.2 ν mass in electroweak physics and physics beyond the standard model

In the standard model of particle physics, the neutrino is taken to be massless. Unlike the photon, the neutrino's masslessness is not required by any basic symmetries. The assumption is based on the fact that no experiment has provided compelling evidence for a non-zero mass.

When we want to amend the standard model to include a massive neutrino, it is tempting to try to mimic the way mass is included for other particles [12]. For example, the mass part of the Lagrangian for electrons looks as follows:

$$-\mathcal{L}_m = m^2 \overline{e} e \tag{2.1}$$

At first glance, nothing prevents us from adding a similar neutrino mass term:

$$-\mathcal{L}_\nu = m^2 \overline{\nu} \nu \tag{2.2}$$

However, only left-handed neutrinos exist, and the above term necessarily includes left- and right-handed terms.

$$\nu = \frac{(1-\gamma_5)}{2}\nu + \frac{(1+\gamma_5)}{2}\nu = \nu_L + \nu_R \tag{2.3}$$

Therefore,

$$\begin{aligned} m^2 \overline{\nu}\nu &= m^2(\overline{\nu}_L + \overline{\nu}_R) \times (\nu_L + \nu_R) && (2.4) \\ &= m^2 \overline{\nu}_L \nu_R + m^2 \overline{\nu}_R \nu_L && (2.5) \end{aligned}$$

For this simple mechanism to work, right-handed neutrinos must exist. However, none have been observed, and the weak interaction in the standard model does not accommodate right-handed neutrinos. A method where we mimic the mass term of the charged leptons cannot be the correct approach.

A more sophisticated approach involving a general Majorana-Dirac[2] mass term can not only overcome this shortcoming but also explain the vast difference between the masses of the neutrinos and their charged leptonic partners [12–14]. Called the *see-saw mechanism*, this method introduces a heavy right-handed neutrino in addition to the massive left-handed neutrinos of the standard model. The right-handed neutrino has a mass on the order of the grand unification (GUT) scale ($m_{\mathrm{GUT}} \sim 10^{14}$ GeV), which explains why it has not been experimentally observed. The regular left-handed neutrino's mass is predicted in the see-saw model to have a mass related to the mass of the charged fermions:

$$m_\nu \simeq \frac{m_f^2}{m_{\mathrm{GUT}}} \ll m_f \tag{2.6}$$

where m_f can be either a fermion or a lepton. This model introduces right-handed neutrinos in a natural way and explains the large scale difference in the mass between the charged and the uncharged leptons. Without the experimental discovery of the heavy right-handed neutrino, it is hard to verify this model, but it does appear to have the right features.

[2] A Majorana particle is a particle that is its own antiparticle. A Dirac particle is one where this is not the case. In the standard model, the neutrino is a Dirac particle.

2.2 Neutrino oscillation

The standard model has been very successful in describing electroweak interactions. Precision measurements at accelerators around the world, especially at LEP, have tested its predictions and not found any disagreement. However, a set of experiments is finding that in at least one respect, the standard model is incomplete. While the standard model neutrino is massless and carries a quantum number, lepton number, which is a conserved quantity, evidence is emerging that this is not the complete story. Neutrinos appear to be able to oscillate, i.e., change from one flavor to another. This oscillation is only possible if neutrinos have mass. In a development that could herald the new direction of physics beyond the standard model, these experiments indicate that neutrino mass and lepton number conservation have to be reconsidered.

Below, I will briefly develop the theory of neutrino oscillations, and discuss the implications of non-zero neutrino mass and individual lepton number conservation.

2.2.1 Vacuum Oscillations

In its simplest incarnation, neutrino oscillation is due to the fact that the free particle neutrino hamiltonian and the weak interaction hamiltonian don't commute.

$$[H, V_{\mathrm{EW}}] \neq 0 \tag{2.7}$$

This means that the eigenstates of the free particle Hamiltonian, i.e., the mass eigenstates, are not the same as the eigenstates of the weak interaction, i.e., the flavor eigenstates.

The mass eigenstates are related to the flavor eigenstates of the standard electroweak model as follows (where we only use two lepton families for brevity).

$$\begin{aligned}
\begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} &= U \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} \\
&= \begin{pmatrix} u_{11} & u_{12} \\ u_{21} & u_{22} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix}
\end{aligned} \tag{2.8}$$

The transformation matrix U must be hermitian and unitary, therefore, it is a function of only one parameter, a mixing angle θ.

An experimental example that follows a neutrino from its creation point in a weak interaction, to detection at a remote particle detector is shown here. At time t_0, in the center of the sun or in a reactor core, an electron neutrino is created via a weak interaction.

$$|\nu_e\rangle = \cos\theta|\nu_1\rangle + \sin\theta|\nu_2\rangle$$

This electron neutrino then propagates to the waiting particle physics detector. Its motion is governed by the free particle equation.

$$|\nu_i(t)\rangle = e^{-(px+iE_it)}|\nu_i(0)\rangle \simeq e^{-i\frac{m_i^2}{2p}t)}|\nu_i(0)\rangle,$$

$i = 1, 2$ and $p \gg m, E \simeq p(1 + \frac{1}{2}(\frac{m}{p})^2)$. Our neutrino evolves with time as follows:

$$|\nu_e(t)\rangle = \cos\theta e^{-iE_1t}|\nu_1\rangle + \sin\theta e^{-iE_2t}|\nu_2\rangle$$

At some later time $t_1 > t_0$, the neutrino arrives at our particle detector. At this point, the wave function has the following transition probability that the neutrino will interact as a muon neutrino ν_μ.

$$\begin{aligned} P(\nu_e \to \nu_\mu, t_1) &= |\langle \nu_\mu | \nu_e(t_1) \rangle|^2 \\ &= \sin^2 2\theta \, \sin^2 \left(\frac{\Delta m^2}{4p} t_1 \right) \quad (2.9) \\ &= \sin^2 2\theta \, \sin^2 \left(1.27 \frac{L_1}{E} \Delta m^2 \right), \quad (2.10) \end{aligned}$$

where we again assume that $p \gg m, E \simeq p(1 + \frac{1}{2}(\frac{m}{p})^2)$ and that the different neutrino mass eigenstates have the same momenta, but since they have different masses, their energies are different.[3] So the probability of oscillation depends on the *difference* in masses between the two mass eigenstates and on a mixing angle $\sin^2 2\theta$.[4] We can now also define an oscillation length, L_{osc},

$$L_{\text{osc}} = 2\pi \frac{2p}{|m_2^2 - m_1^2|} = 2\pi \frac{2p}{\Delta m_{21}^2} \quad (2.11)$$

which represents distance over which a neutrino completes one complete oscillation cycle. We see that oscillation length L_{osc} depends on the ratio of neutrino momentum p and Δm^2. Depending on the values of these two parameters, given a pair $\sin^2 2\theta, \Delta m^2$, different experiments will probe different values of L/E. We will find it useful to classify experiments with this ratio in mind.

[3]In the final equation, the units of L are in meters, E in MeV, and $\Delta \mathrm{m}^2$ in eV^2.

[4]Strictly speaking, neutrino oscillation experiments do not provide information about the neutrino masses themselves. Instead, the measurement gives us information about

$$\Delta m_{21}^2 \equiv m_2^2 - m_1^2$$

that is, the difference in the squares of the masses of the neutrino mass eigenstates.

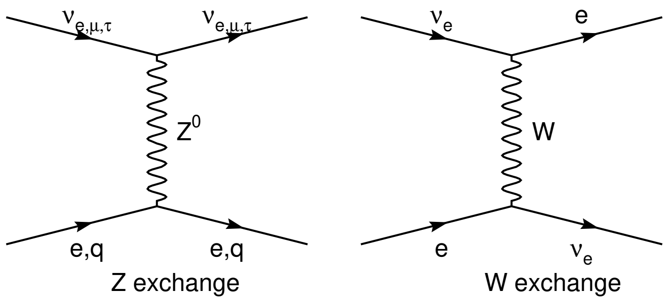

Figure 2.1: Neutrino interactions in matter

2.2.2 Matter-enhanced oscillations

In addition to the vacuum oscillation behavior described above, there exists a region in the $\sin^2 2\theta, \Delta m^2$ parameter space where a resonant enhancement of the oscillation can take place due to the presence of matter. This resonant enhancement, first described in [15] and [16], is called the MSW effect, after its discoverers, S. Mikheev, A. Smirnov, and L. Wolfenstein. It is of particular importance for solar neutrino oscillation experiments.

In matter, neutrinos undergo interactions as shown in Fig. 2.1. Since normal matter consists of only the first family of quarks and leptons, the electron neutrinos can interact more readily with matter than can muon or tau neutrinos. All three neutrino types can undergo neutral current Z^0 exchange, while only the electron neutrinos undergo additional charged current W exchange. This difference leads to an interference term in the forward scattering amplitude, which can lead to a resonant enhancement of the oscillation probability.

In the case of matter oscillations, the pure vacuum eigenstates ν_1, ν_2 are not eigenstates of the Hamiltonian governing the motion of particles in transit through matter. We can now have additional transitions between ν_e and ν_μ mediated by this new Hamiltonian. The matter eigenstates, ν_{1m} and ν_{2m}, are modified by the addition of an effective interaction which connects ν_e to ν_e.

$$H_{\text{eff}} = \sqrt{2} G_F N_e |\nu_e\rangle\langle\nu_e|$$

The old and new mixing angles are related by the ratio of the strength of this interaction, $A \equiv 2\sqrt{2} G_F N_e p$, and Δm^2.

$$\sin^2 2\theta_m = \frac{\sin^2 2\theta}{(\cos 2\theta - A/\Delta m^2)^2 + \sin^2 2\theta} \tag{2.12}$$

where θ is the vacuum mixing angle and θ_m is the effective mass mixing angle between the flavor eigenstate and the mass-matter eigenstate. For $A_R \equiv \cos 2\theta\ \Delta m^2$, we have a resonant condition where $\sin^2 2\theta_m \simeq 1$ regardless of the value of the vacuum mixing angle θ and hence mixing is maximal. Fig. 2.2 shows the behavior of the matter mixing angle as a function of $A/\Delta m^2$ for several values of θ. We see that while the width of the resonance region changes, even for small values of θ, maximal mixing ($\sin^2 2\theta_m = 1$) occurs.

2.2.3 MSW resonance

Fig. 2.2 shows mixing can occur for any value of the vacuum mixing angle.[5]

$$\tan 2\theta_m = \frac{\sin 2\theta}{\cos 2\theta - \frac{A}{A_R}} \tag{2.13}$$

[5]This discussion of MSW resonance is largely based on [12, 13].

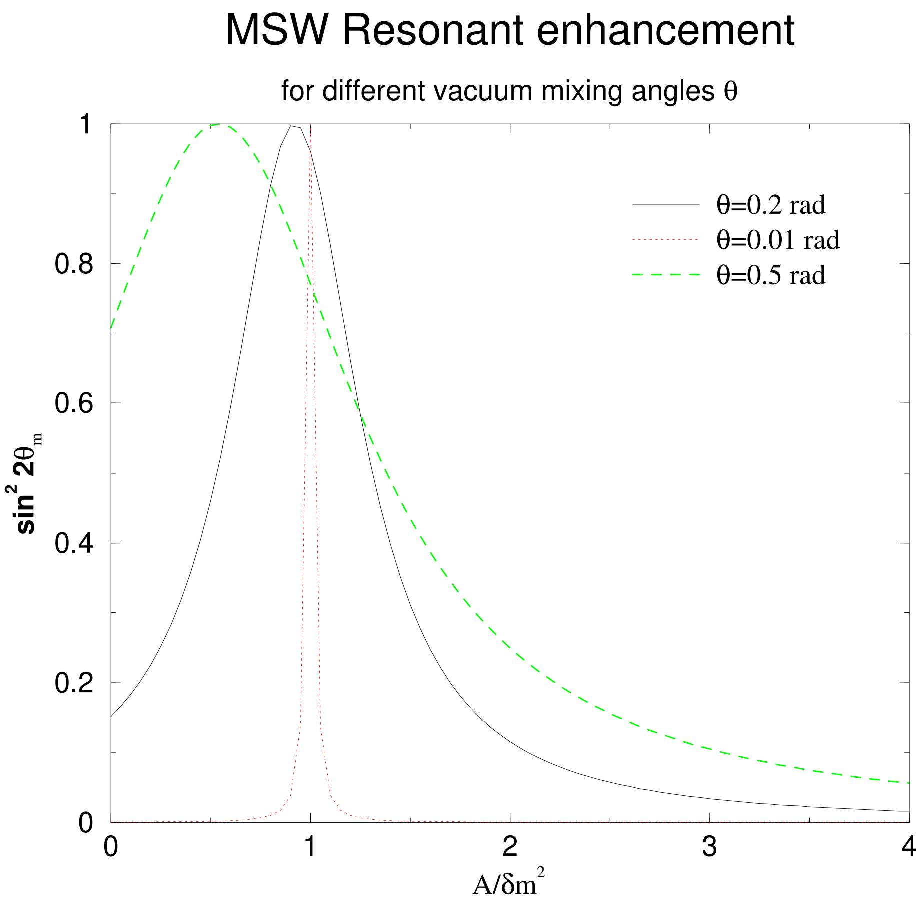

Figure 2.2: MSW resonant enhancement of neutrino mixing the presence of matter. Resonance occurs for $A = A_R \equiv \Delta m^2 \cos\theta$.

Eq. 2.13 shows the behavior of the mixing angle θ_m, and Fig. 2.3 illustrates it graphically. We see that the state changes from being exclusively one flavor, $\theta \simeq 0$, at small values of $A/\Delta m^2$, to exclusively a different flavor, $\theta \simeq 90°$, on the other side of the resonance condition at $A/\Delta m^2 = 1$.

We now discuss how this can be applied to solar neutrinos. In the case of constant density, the resonance would only occur for fine-tuned values of the density. However, solar neutrinos pass through the sun, which has regions of changing density and hence resonance can occur for many values of the mixing parameters.

We assume that neutrinos in the sun are born in a region of high density such that $A \gg A_R$ where A_R is the resonance density. As the neutrino is created in the weak interaction, it is in the weak eigenstate ν_e. We see from Fig. 2.3 that in regions of high matter density ($A \propto N_e$), the mixing angle between ν_e and ν_{1m} is almost $90°$, i.e., at high A, $\nu_e = \nu_{2m}$. Passing from high to low matter densities, A passes through the resonance region, where $\theta_m = 45°$. Finally, upon exiting the sun, the mixing angle returns to the vacuum value θ.

There are two different conditions, called the adiabatic and non-adiabatic condition, which change the behavior in the resonance region. Adiabatic refers to the rate of change of A or of the electron density of the sun. In the adiabatic regime, the changes occur very slowly, and the path length in the resonance region is many times longer than the oscillation length. States that are born as a particular matter-enhanced wave function eigenstate are likely to remain in this eigenstate throughout the resonance region and leave the sun in this state. In the case of small mixing angles and at large densities, the angle between the flavor eigenstate and the mass eigenstate is $90°$ rotated with respect to the angle at which the particle eventually emerges from the sun. This means that in this case, the probability that a neutrino

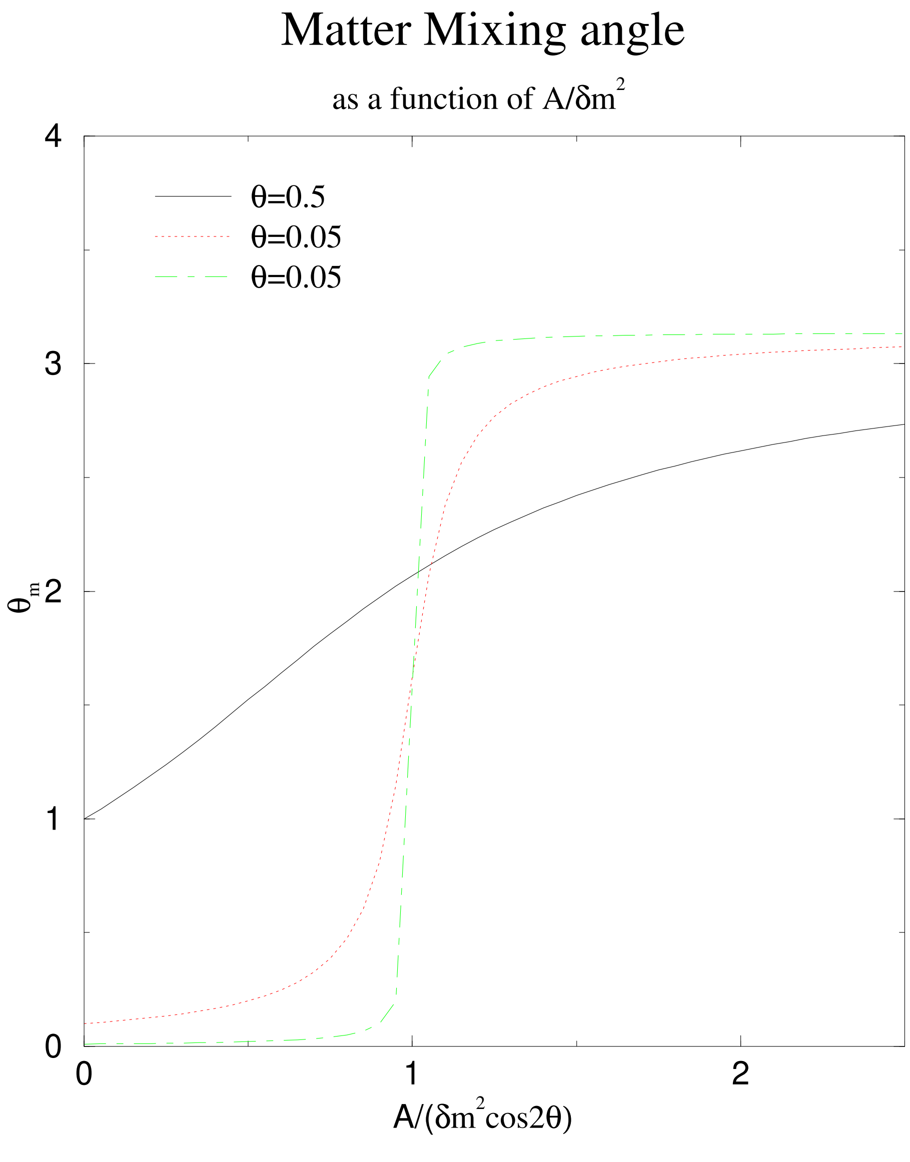

Figure 2.3: Behavior of θ_m as a function of $A_R \equiv \frac{A}{\Delta m^2 \cos 2\theta}$.

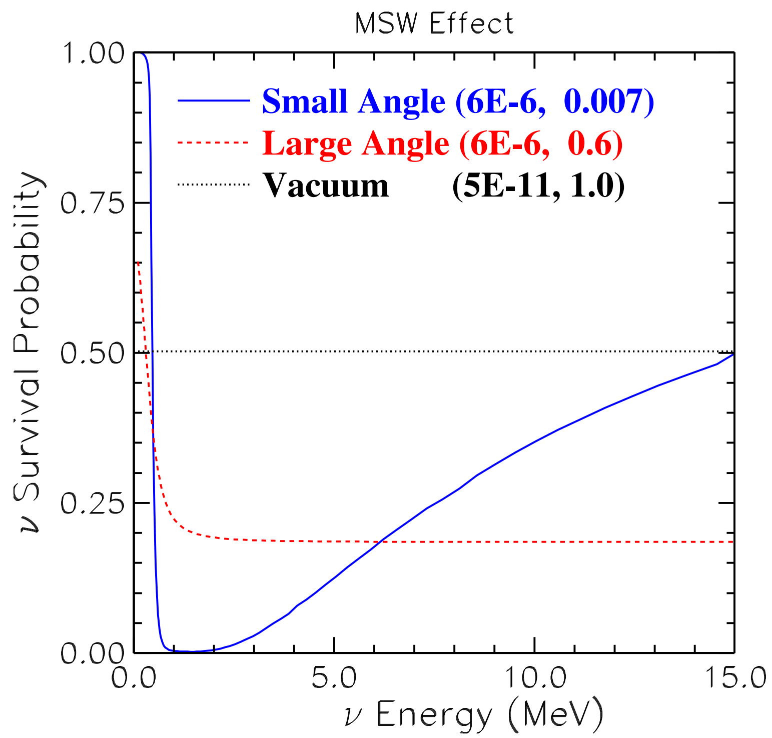

Figure 2.4: ν_e survival probability as a function of energy for different neutrino oscillation schemes. The curve shown for the vacuum solution is an average of a rapidly varying probability.

that is born as ν_e will emerge from the sun as ν_e can be almost zero.

In the non-adiabatic case, there is a larger probability of a level-jump in the region where the two eigenstates are closest. This means that there is a larger survival probability for electron neutrinos in the non-adiabatic case.

The results from previous experiments have traditionally produced three allowed regions in the $\Delta m^2, \sin^2 2\theta$ plane. These allowed regions are typically referred to as the 'large mixing angle (LMA),' 'small mixing angle (SMA),' and 'just-so' solutions. The LMA and SMA are in the region of the mixing parameters where resonance occurs, while the 'just-so' solution is a vacuum oscillation solution. Fig. 2.4 shows the neutrino survival probability for these different solutions. The 'just-so' solution, labelled 'Vacuum' in Fig. 2.4, includes the effects of position uncertainty averaged over the neutrino creation region in the sun, which causes the rapid oscillations to average to 1/2 for most parts of the parameter space. The survival probability in the SMA solution has a strong dependence on the energy, which would be visible in a detector as a distortion of the shape of the energy spectrum with respect to the prediction of the SSM. The variations as a function of energy could also explain the different suppressions seen by the different experiments in Fig. 1.2.

MSW resonances mean that for the SNP, the neutrino oscillation hypothesis no longer needs to be fine-tuned in order to fit the experimental data. This makes neutrino oscillations aesthetically more appealing than they would be in the case where only vacuum oscillations were possible.

Fig. 2.5 shows a recent overlay of the $\sin^2 2\theta, \Delta m^2$ parameter space for neutrino oscillation. The solar neutrino allowed regions are shown.

2.3 Standard Solar model

The standard solar model [1] attempts to describe a main sequence star during its stable hydrogen burning phase. The sun is such a star. It derives its power from a process of hydrogen burning. The main process of hydrogen fusion can be summarized as

$$2\mathrm{e}^- + 4\mathrm{p} \rightarrow {}^4\mathrm{He} + 2\nu_\mathrm{e} + 26.7\ \mathrm{MeV}. \tag{2.14}$$

In this process, protons are fused to produce the heavier helium, emitting two neutrinos and ~ 27 MeV of energy in the form of photons.

Equation 2.14 relates the total amount of energy emitted by the sun in form of

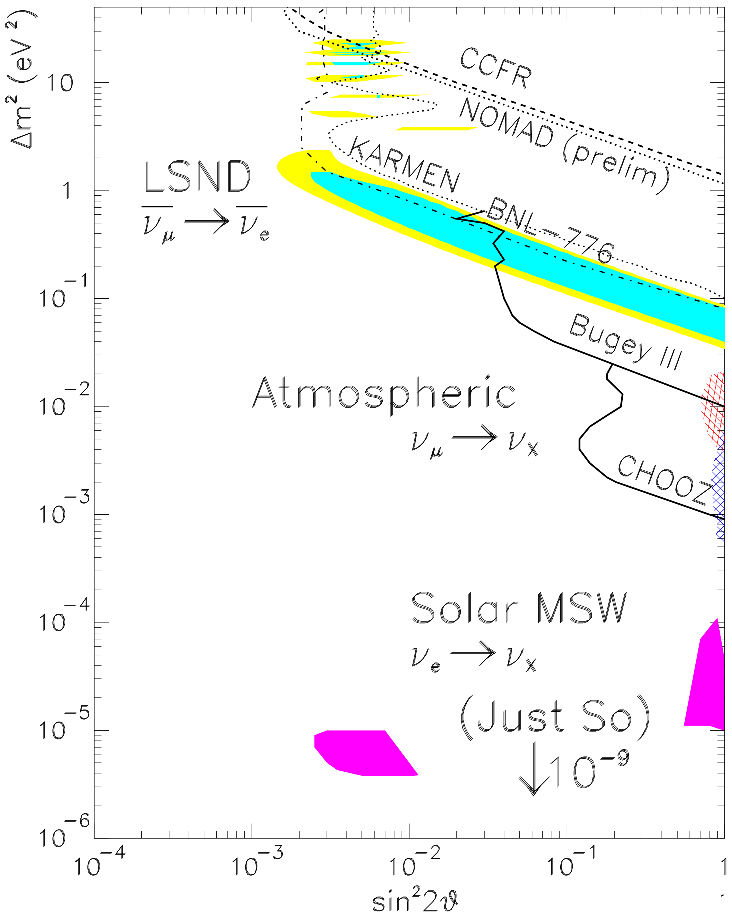

Figure 2.5: $\Delta m^2, \sin^2 2\theta$ overlay of all neutrino experiments. Figure from [17].

thermal energy to the total neutrino flux – for every two neutrinos, 26.7 MeV of energy are released. The total measured solar flux on earth therefore indicates the flux of solar neutrinos at earth is $\Phi_\nu \simeq 6 \times 10^{10}\, \nu_e$ cm^{-2} s^{-1} [14, p. 134].

The initial motivation for solar neutrino physics was to use the neutrinos as a tool for probing the solar interior. Photons emitted from the sun last interacted in the outer layers of the sun, and therefore are not good messengers from the solar core. Neutrinos, on the other hand, are so weakly interacting that they escape from their production region at the center without modification and therefore could provide clues to the processes in the solar interior.

The solar model rests on uncontroversial assumptions. The sun is assumed to be in hydrostatic equilibrium, where the gravitational pressure from the outside counters the nuclear radiation pressure from the inside. Nuclear fusion reactions are its only source of energy. The sun is also assumed to be in thermal equilibrium. The following measured solar properties are inputs into the solar model.

$$L_\odot = 3.86 \times 10^{33}\,\mathrm{ergs/sec} \tag{2.15}$$

$$M_\odot = 1.99 \times 10^{33}\,\mathrm{g} \tag{2.16}$$

$$R_\odot = 6.96 \times 10^{5}\,\mathrm{Km} \tag{2.17}$$

$$T_\odot^{\mathrm{surface}} = 5.78 \times 10^{3}\,\mathrm{K}. \tag{2.18}$$

In the SSM, all energy transport in the sun occurs in the form of convection or in the form of radiation, and all nuclear abundances are solely due to the nuclear fusion reactions mentioned above. The fusion reactions for the *pp* chain are shown in Fig. 2.6. Neutrino-emitting reactions have the names by which the neutrinos are referred to

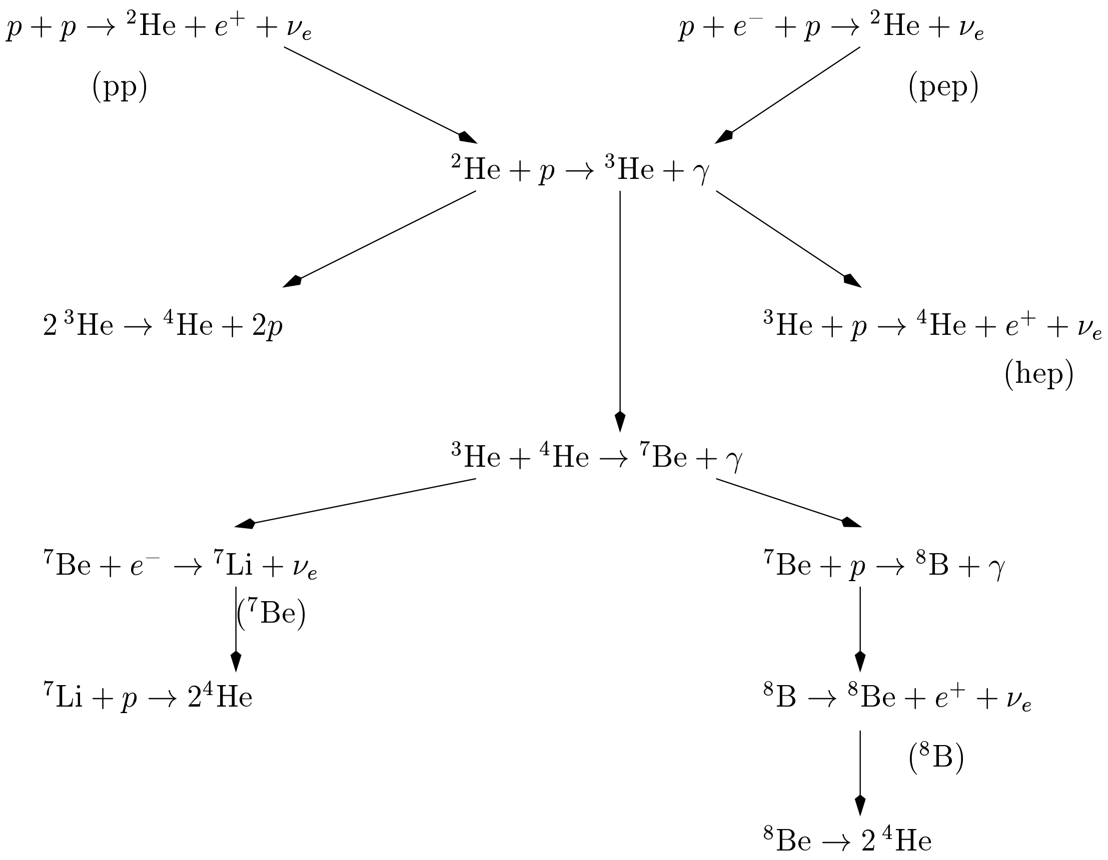

Figure 2.6: Nuclear reactions of the *pp* chain. Adapted from [13]. The ν-producing reactions are labelled with the names of the associated neutrinos. See Fig. 1.1 for the energy spectra of the neutrinos. The CNO cycle, which also produces neutrinos, is not shown.

indicated in parenthesis. An additional chain, the CNO cycle, is responsible for a small fraction of the energy emitted by the sun, and is not shown. The spectra of the neutrinos emitted in the *pp* cycle was shown in Fig. 2.6. The *pp* neutrinos have an energy range of $E_\nu \leq 18\mathrm{MeV}$.

The SSM is an attempt to describe the workings of main sequence stars such as the sun. Using nuclear physics inputs and the measured luminosity, radius, and mass of the sun, the SSM predicts the details of the nuclear fusions that give the sun its power. These predictions are born out by the helioseismological measurements, whose measured solar density profiles agree with the prediction to about 1%. The measured discrepancy between the SSM and experiment remains in the neutrino sector, where both the ^{8}B and *pp* neutrinos have been measured and are found to be in disagreement with the predictions of the SSM.

In addition to the SSM, many non-standard solar models have been proposed to address the SNP. However, it has been shown that the changes to the solar model cannot account for all the experimental evidence in the SNP [7, 8].

2.4 In conclusion

This section briefly introduced the standard model of particle physics and neutrino oscillations. A brief description of the standard solar model followed. With these two topics, we have fully introduced the SNP and motivated why it is interesting for particle physicists.

Chapter 3

SNO Signals and Backgrounds

We will outline a method for testing the neutrino oscillation hypothesis as a solution to the solar neutrino problem, and show how SNO implements this method. The test is designed to determine, with a minimum of dependence on the theoretical input, whether neutrinos changing flavor from one species to another explains the experimentally observed deficit in neutrinos from the sun.

We also will discuss SNO's signals in some detail, how they are seen in the SNO detector, and outline strategies for characterizing each of them.

3.1 Strategy to Test the Oscillation Hypothesis

Previous neutrino experiments have measured a deficit of electron neutrinos coming from the sun. The comparison between the experimental measurement and the theoretical prediction lead to the belief that neutrinos had changed flavor to another type of neutrino that the experimental apparatus was not able to detect.

If a detector could measure the flux of all neutrinos, independent of flavor, and also the flux of electron neutrinos, that is, those emitted in the nuclear reactions that power the sun, the ratio of these two numbers could show conclusively that the solar neutrinos have changed from one flavor to another. If

$$\frac{\Phi_{\nu_e}}{\Phi_{\nu_x}} \neq 1, \tag{3.1}$$

neutrino oscillation explains the solar neutrino problem.

Neutrinos only interact weakly, and the energies of the solar neutrinos are relatively low ($E_\nu <$ 20 MeV). These two properties conspire to make neutrinos particularly hard to detect. The secondary particles which one detects are of such a low energy that many radioactive signals occur in the same signal region. In order to maximize one's chance for success, such a measurement must be made in an environment free from backgrounds. This includes cosmic ray background and radioactive backgrounds from the materials that make up the detector.

Cosmic rays would increase dead time (or dead regions) in the detector, reducing the effective fiducial volume. In order to reduce the impact of this background, one can build an adequately shielded detector by placing the detector in a deep mine. The overburden of rock such a mine provides only allows the most energetic muons to penetrate to the depth of the experiment and decreases the cosmic ray flux at the detector.

Backgrounds from normal radioactivity mimic the solar neutrino signals. Due to the low interaction cross section of the neutrinos, the rate of physics events is sufficiently low that even a small contamination can skew the results. Moreover, these backgrounds are hard to get rid of. Most materials one constructs detectors out of contain some amount of radioactivity. By a judicious choice of materials, one can reduce but not eliminate this background.

3.2 How SNO Implements this Strategy

The Sudbury Neutrino Observatory is an attempt to implement this strategy in a real-world detector. By using heavy water, SNO is able to realize the test Eq. 3.1 proposes. SNO's location and construction process address the background concerns.

SNO can detect solar neutrinos via three separate reactions. Two of these three reactions are unique to SNO, and allow SNO to implement the plan outlined above. We re-cap the interactions from Sec. 1.2.

$$\nu_e + d \rightarrow e^- + p + p \tag{3.2}$$

$$\nu_x + d \rightarrow n + p + \nu_x \tag{3.3}$$

$$\nu_x + e^- \rightarrow \nu_x + e^- \tag{3.4}$$

where $\nu_x = \nu_{e,\mu,\tau}$. Reaction (3.2) is called the *charged current* (CC) reaction. In it, the recoil electron is observed via its Čerenkov radiation. This reaction measures the electron neutrino flux. In the *neutral current* (NC) reaction (3.3), all neutrino flavors e, μ, τ interact with equal strength. This reaction is a measure of the total flux of active neutrinos from the sun. By comparing the ratio $\frac{\Phi_{CC}}{\Phi_{NC}}$ to the ratio described in Eq. 3.1,

$$\frac{\Phi_{CC}}{\Phi_{NC}} \equiv \frac{\Phi_{\nu_e}}{\Phi_{\nu_x}}, \tag{3.5}$$

we can make a model-independent test of whether neutrinos from the sun oscillate. Both reactions in Eq. 3.2 and Eq. 3.3 are made possible by the use of heavy water as the target.

The third reaction shown in Eq. 3.4, the *elastic scattering* reaction, is the more usual reaction used in light water Čerenkov solar neutrino detectors. At first glance, it too appears to be a measure of the flux of all neutrinos rather than the flux of electron neutrinos, and we wonder why these experiments would measure a deficit of neutrinos, as they are sensitive to all neutrino flavors. The cross section for electron neutrinos is approximately six times larger than the cross section for muon and tau neutrinos. This makes Eq. 3.4 a bad measure for the flux of all neutrinos.

3.3 Backgrounds

In order to escape from the cosmic ray background, the SNO detector is located 6800 feet underground. This corresponds to roughly 6000 meters water equivalent. At this depth, only the most energetic cosmic rays penetrate to the detector, reducing the muon flux to about 3 per hour.

During construction, all detector materials were assayed and chosen to minimize their intrinsic radioactivity. Contamination with ^{232}Th and ^{238}U produce long-lived daughters that produce backgrounds for the neutral current detection. These elements have short-lived radioactive daughters that emit γ-rays with an energy of 2.4 (2.6) MeV for ^{214}Bi (^{208}Tl). These γ-rays have an energy above the deuteron photodisintegration threshold (2.2 MeV). This reaction mimics the neutral current signal. The γ-rays are a more difficult problem, since the Compton scattering length for γ-rays in water is about 30 cm. So while neutrons emitted in the light water are quenched, γ-rays can and do travel into the heavy water and interact there. To minimize this background, all materials in the SNO detector must be kept as radioactively clean as possible.

The final defense against radioactivity in detector components is a physical barrier between detector materials and the heavy water itself. The heavy water is surrounded by 1.7 kT of ultra-pure light water, with a 2.4 m minimum separation between the photo-tubes and the heavy water. An additional 5 kT of heavy water surrounds the phototubes, providing a barrier for radioactivity from the rock wall itself.

Cleanliness is especially important at SNO since the mining environment where the lab sits is very dirty, and relatively high in radon. To address this problem, the underground lab is run as a clean room, and all materials[1] that come into the lab must be cleaned of all mine dust and dirt.

3.4 Detecting the neutrinos

Ultimately, neutrino charged-current, neutral-current and elastic-scattering interactions produce highly relativistic electrons in SNO (see Eqs. 3.4, 3.2, and 3.3). When charged particles travel faster than the speed of light in their medium, the radiated waves from the medium can constructively interfere to produce a coherent signal far from the source. This signal is called Čerenkov radiation and is the basis for neutrino detection in SNO.[2] In the CC and ES reactions, the neutrino interaction directly produces an electron, whereas in the NC reaction, the free neutron's capture on a nucleus produces signal which is in turn detected.[3] These reactions, their properties, and strategies for measuring them are discussed below.

[1] including humans!

[2] See [18] for the seminal textbook on Čerenkov radiation.

[3] The neutron detection method described refers to two of three planned detection mechanisms. A third involves direct detection of the neutron in ^{3}He proportional counters, as will be discussed later in this section.

3.4.1 CC and ES

It makes sense to discuss the *ES* and *CC* reactions simultaneously, as both produce detectable electrons. On an event-by-event level, the electrons from CC and ES reactions are indistinguishable. They can be separated statistically, as their angular, radial and energy distributions differ.

The electrons produced in the two reactions differ in their angular distributions with respect to the neutrino direction (and, equivalently, the instantaneous vector pointing to the sun). ES electrons are highly peaked in the forward direction, while the CC neutrinos have a slight forward-backward asymmetry in their angular distribution [1, p197].

$$P(\cos\theta) = 1 - \frac{\beta}{3}\cos\theta \tag{3.6}$$

where $\beta = v/c$. These different angular distributions allow one to measure the relative and absolute number of CC and ES neutrinos given a $\cos\theta_{\odot}$ distribution.

CC and *ES* electrons also have different energy distributions. Fig. 3.1 shows the energy spectrum of electrons that have undergone *CC* and *ES* reactions induced by mono-energetic neutrinos of energy $E_\nu = 10.0$ MeV. We can see that the *CC* reaction has a narrower distribution than the *ES* reaction. The *CC* reaction therefore allows a better estimate of the incident neutrino energy. The broad *ES* response makes this task much more difficult for the *ES* neutrinos. Fig. 3.2 shows the convolution of the ^{8}B spectral shape with the *CC* and *ES* response. The *CC* spectrum has a higher mean, which is crucial for neutrino experiments as more of the events will be at higher energies and away from radioactive backgrounds. Additionally, as discussed in Sec. 2.2.3, the MSW solution predicts distortions in the shape of the ^{8}B spectrum that are more readily measurable with a *CC*-like spectral shape rather than an *ES*-like

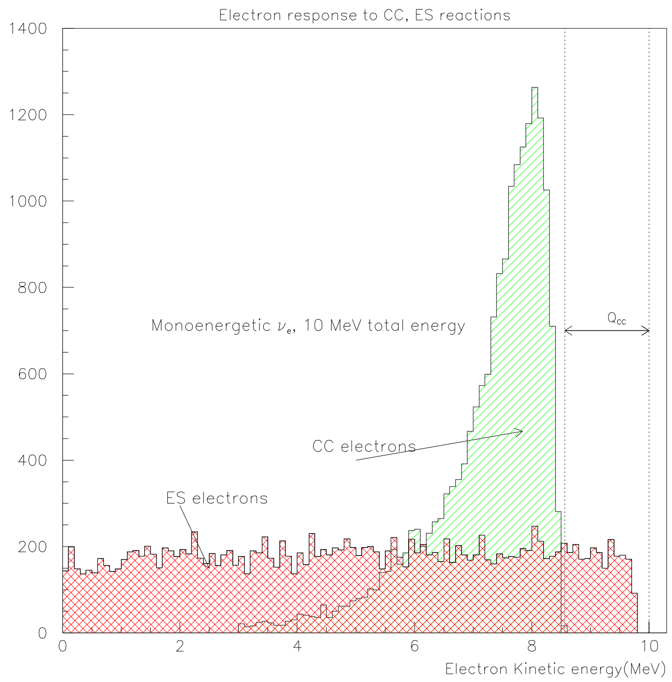

Figure 3.1: Spectrum for Monte Carlo CC and ES electrons, as predicted by SNOMAN 3_0192. The Q value for the CC reaction is 1.44 MeV. The total electron energy is plotted on the abscissa. For neutrino of the same energy, the CC electron energy spectrum is much more sharply peaked than the ES reaction. The incident neutrino has an energy of 10 MeV.

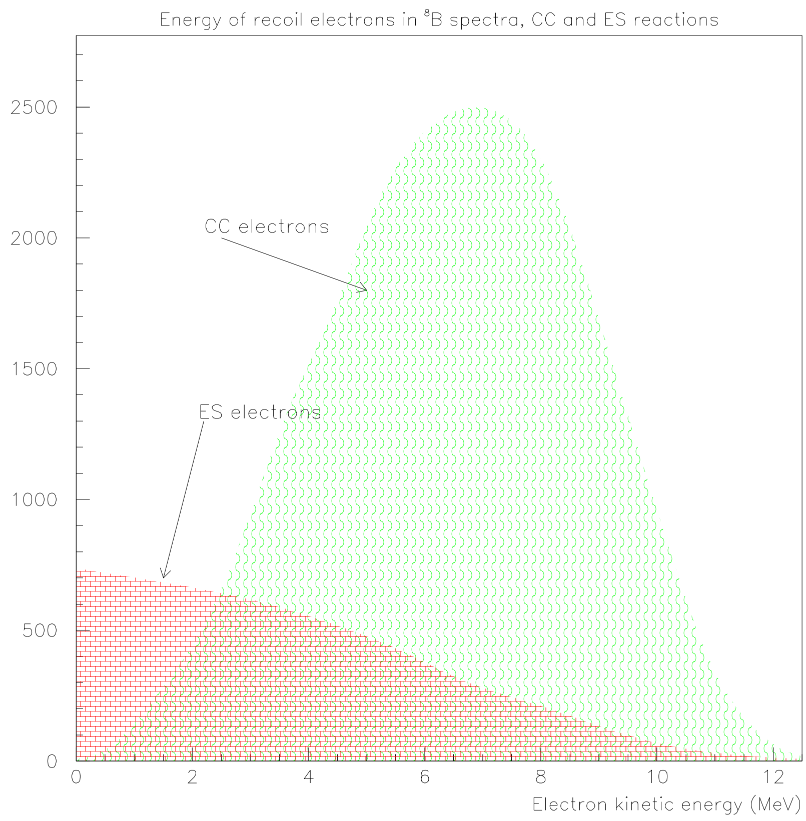

Figure 3.2: Convolution of the CC and ES response shapes with the ^{8}B spectral shape.

shape.

Another tool in distinguishing the ES and CC reactions lies in their radial distributions. Since the CC reaction only occurs in the presence of heavy water, the number of CC neutrinos must decrease sharply at the edge of the heavy water volume. The ES reaction, however, will occur in both the light and heavy water. Given a good understanding of the acceptance of the detector in this region, the radial distribution of events will allow an additional check on the relative contributions of CC and ES electrons. However, the varying optical properties near the light water-heavy water interface and the backgrounds in the light water from the PSUP and PMTs are likely to make this a difficult measurement.

The radial distribution and $\cos\theta_\odot$ distributions of CC, NC, ES and background signals can be used to estimate the relative contributions of these signals. While the energy spectra of CC, NC and ES events are totally different, measuring the spectral shape is one of the goals of the SNO experiment and therefore cannot be used *a priori* to determine CC and ES relative distributions.

3.4.2 NC

The neutral current reaction is presented in Eq. 3.3. As none of the escaping particles are charged, no Čerenkov radiation is emitted. However, the produced free neutron can be detected. SNO has three methods for detecting these neutrons. In two of the methods, the neutron thermalizes in the heavy water volume, and the capture nucleus emits a γ-ray during its deexcitation. The third method involves capture on ^{3}He and detection of the $^3\text{He}(n,t)p$ reaction using ^{3}He proportional counters as an auxiliary Neutral Current Detector (NCD). The three methods are very different and

		Run Phase	
Isotope	Pure D_2O	Salt	NCD
D	27%	3 %	12%
$H(D_2O)$	15%	2 %	7%
$H(H_2O)$	7%	2 %	4%
H(AV)	34%	9 %	21%
^{16}O	5 %	0 %	2 %
^{17}O	11 %	1 %	5 %
Cl		82%	
^{3}He			45%
Ni			2%
Signal	27 %	85 %	57 %

Table 3.1: Capture percentages of neutrons by isotope, for pure D_2O for salt, and for NCDs. Capture percentages less than two percent are not shown. 'Signal' refers to the sum of all interactions that are detectable.

have different backgrounds and systematic effects associated with them. The SNO collaboration will use each method in turn in order to measure the NC signal. In Tab. 3.1, the neutron capture efficiency of SNO in various detector configurations is summarized.

Phase One: Neutron capture in pure heavy water

In the first phase, SNO is filled with D_2O of an isotopic purity of better than 99.92%. In this configuration, the free neutron from the NC reaction can be captured on one of three targets: H, D, and O.

If the neutron is captured by a deuteron, SNO can detect it via the following reaction.

$$n + d \rightarrow t + \gamma, \qquad E_\gamma = 6.25 \text{ MeV}, \ \sigma_c = 5 \times 10^{-4} \text{barns} \tag{3.7}$$

The γ-ray will most likely Compton-scatter, transferring most of its energy to an electron which emits Čerenkov photons and is detected. The γ-ray, however, has no angular correlation to the incoming neutrino direction, so one cannot infer the direction of the neutrino from this process. In addition, the reaction does not contain any energy information except that the neutrino was above the NC reaction threshold, $E_{th} = 2.226$ MeV.

If the neutron wanders into a hydrogen nucleus, SNO cannot detect it.[4]

$$n + H \to d + \gamma, \qquad E_\gamma = 2.2 \text{ MeV}, \ \sigma_c = 0.3 \text{ barns} \tag{3.8}$$

The γ-ray in this case has an energy too low to result in a detected neutral current event. Unfortunately, the capture cross section for n on H is approximately 635 times larger than the capture cross section on d [19]. Therefore, H_2O contamination in the heavy water immediately affects the NC detection efficiency, and any degradation of the D_2O leads to a decrease in NC detection efficiency listed in Tab. 3.1.

The neutron can also capture on an oxygen nucleus. As with the capture on H, the capture leads to an undetectable γ-ray.

$$n + {}^{16}\text{O} \ \rightarrow \ {}^{17}\text{O} + \gamma, \ \sigma_c = 1.8 \times 10^{-4} \text{ barns} \tag{3.9}$$

$$n + {}^{17}\text{O} \ \rightarrow \ {}^{14}\text{C} + \alpha, \ \sigma(n, \alpha) = 0.235 \text{ barns} \tag{3.10}$$

The total cross section for the latter process is competitive with the capture on H.

In the pure D_2O phase, several strategies will have to be used to understand the relative number of NC, CC and ES neutrinos. As discussed above, the CC and

[4]To be accurate, the reaction might trigger the detector, as our current hardware threshold is $E_T \simeq 2$ MeV. However, the signal-to-noise ratio in this energy region is such that the data is swamped by background.

ES reactions have very different angular distributions relative to the sun's position. While the *CC* and *ES* events cannot be separated on an event-by-event basis, they can be distinguished from one another and from the *NC* reaction statistically by this angular distribution. However, no such correlation exists with *NC* events. The flat angular distribution of the neutral currents cannot be distinguished from background events that occur in the D_2O and that leak into the data sample. The reconstructed direction of background events is also likely to be uncorrelated with the solar direction. Some attempts have been made to exploit the differences between electron-induced events and γ-induced events to distinguish individual *NC* events from *ES* and *CC* events, though these techniques appear to be more successful on a statistical level rather than on an event level.

It would also be possible to understand the neutral current signal strength by examining the radial distribution of events near the edge of the acrylic vessel. The *NC* rate is expected to drop precipitously near the AV. Neutrons produced near the AV but in the heavy water will often escape to the light water due to the long mean free path in heavy water and short mean free path in the acrylic vessel and light water. The region near the AV is therefore expected to be a "*NC*-free" region.[5] However, as stated above, correctly performing this measurement requires a good understanding of the optics and acceptance of the detector near the edges of the AV.

Given the relative cross-sections of neutron capture on light water, heavy water, and oxygen, the capture efficiency of neutrons in the pure heavy water configuration

[5]The depletion of *NC* events near the acrylic vessel is akin to using a neutron *poison* in this region. An option for SNO to obtain a clean measure of the *CC* energy spectrum would be to introduce a neutron-quenching material such that neutrons are silently captured. This would allow a clean measurement of any distortions in the *CC* energy spectrum due to MSW effects. The NCD ^{3}He detectors provide some of the similar benefits as a nice side effect, but not as efficiently as a real poison.

is 27%. SNO will improve its NC capture efficiency by two means in the subsequent two phases of data taking.

Phase Two: Neutron capture with salt

After the pure D_2O phase, SNO will increase its capture efficiency of NC events by introducing a small amount of $MgCl_2$ or $NaCl$ salt into the heavy water. The detection then proceeds as shown below.

$$n + {}^{35}\mathrm{Cl} \rightarrow {}^{36}\mathrm{Cl} + \gamma, \qquad \sum E_\gamma = 8.6 \text{ MeV}, \; \sigma_c = 43 \text{ barns} \tag{3.11}$$

The chlorine's neutron cross section in reaction Eq. 3.11 is 143 (86×10^3) times the capture cross section on H(d) (see Eq. 3.8,3.7). One or more γ-rays are emitted in the capture, for a total energy of 8.6 MeV. As in the $n + d$ reaction, the γ-rays Compton scatter and are detected by SNO. The total neutron capture efficiency is calculated to increase to 85% in the salt configuration.

Note that the γ-ray energy in this configuration is 2.5 MeV higher than in the pure D_2O configuration. While the CC signal stays constant, the NC signal shifts. By comparing the electron energy spectra measured with pure D_2O and with salt, the contribution in the energy spectrum due to the NC signal can be determined.

Phase Three: Neutron capture with ^{3}He detectors (NCD)

The third phase for detecting neutrons involves installing a separate detection subsystem consisting of ^{3}He proportional counters in the D_2O. These detectors are designed to allow an event-by-event separation between neutral-current and charged-current or elastic-scattering signals. While ES and CC events would still be detected as before via Čerenkov light, the NCDs provide a completely separate data path for the detection of neutrons. The detection reaction is listed below.

$$^{3}\mathrm{He} + n \rightarrow p + {}^{3}\mathrm{H} + 764\ \mathrm{KeV},\ \ \sigma(n,p) = 5.33 \times 10^{3}\ \mathrm{barns} \tag{3.12}$$

Due to the long free path lengths of neutrons in heavy water, and the large cross section for the reaction in Eq. 3.12, a relatively sparse array of detectors is sufficient for a high neutron capture efficiency. The sparseness of the array minimizes the optical effects on the propagation of Čerenkov light from CC and ES reactions.

As with all other SNO detector elements, a big challenge for the collaboration is to assure that the NCDs do not introduce an unacceptable amount of radioactivity into the detector. The NCDs sit in the heavy water, making their need for cleanliness even more acute. In order to study the induced backgrounds of the NCDs, a small uninstrumented test detector has been built from several lengths of tubing similar to the material of the NCDs themselves. The collaboration will take data with this "CHIME" device[6] suspended inside the heavy water. This data will be analyzed in an attempt to extract an upper limit on the NCD-induced background [20].

From the point of view of the PMTs, the neutral current detectors act as a "poison" for the neutral current signal. As discussed in the previous section, the shape of the CC signal is an important means of determining the type of MSW solution to the solar neutrino problem. In the case of either pure D_2O (phase 1) or salt (phase 2), the NC signal complicates this analysis as neutrons are detected by Compton-scattered electrons, that are indistinguishable from CC events on an event-by-event basis.

As can be seen in Tab. 3.1, neutrons still capture on d in Phase Three, so the

[6]The CHIME acronym stands for "Construction Hardware In-situ Monitoring Experiment."

PMT signal is not neutron-free. The ratio of cross sections between the ^{3}He in the NCDs and ^{2}H in the heavy water is well known. By measuring the neutron capture rate on NCDs, we will be able to accurately calculate how many neutrons capture in the heavy water and subtract this amount from the PMT signal to get a neutron-free signal. This should improve the measurement of any MSW-induced CC energy spectrum distortion.[7]

3.4.3 Summary

In this section, we presented a strategy for testing the neutrino oscillation hypothesis to the solar neutrino problem. We showed how SNO realizes this strategy in a physical detector. Finally, the chapter concluded with an introduction of the main SNO signals and some of the strategies that SNO will use to measure them.

[7]If a precise measurement of the CC spectrum cannot be obtained otherwise, SNO could insert a real poison with a higher neutron capture efficiency.

Chapter 4

The SNO detector

This chapter provides a description of the SNO detector. As mentioned in Section 1.2, SNO is a second generation solar neutrino observatory. Like Kamiokande and SuperKamiokande, SNO is a water Čerenkov detector. In contrast, SNO's target is ultra-pure heavy water, which has significant advantages, as has been described in Ch. 3. A cut-away diagram of the detector can be seen in Fig. 4.1.

In this chapter, we will first describe the location and major components of the detector. Next, we will briefly discuss the calibration system and its goals. Afterwards, we will introduce the SNO analysis software suite and some of its functions. Finally, we will present an overview of the SNO physics signals.

4.1 Overview

The SNO experiment is located near Sudbury, Ontario, Canada. The detector cavity sits 2039 m below the earth's surface in INCO Limited's Creighton mine. The effective shielding provided by the Norite rock above is equivalent to roughly 6000 m of

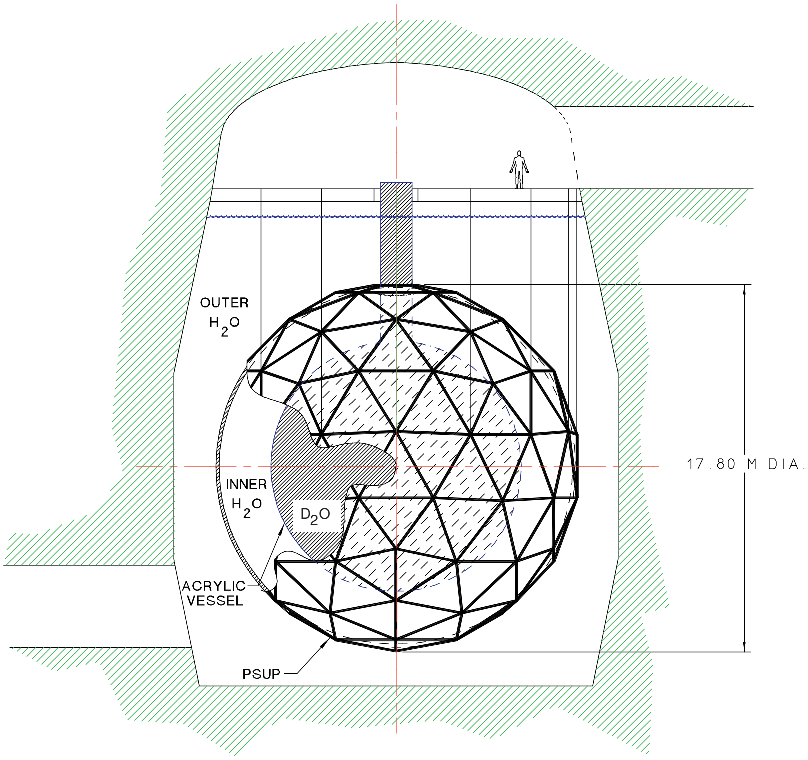

Figure 4.1: Cut-away of the SNO detector. Originally published in [21].

water shielding. At that depth, approximately three muons pass through the detector volume each hour.

The cavity is 34 m tall and 22 m in diameter. Encircling the detector is a set of coils, embedded in the cavity's rock wall which bathe the detector in a magnetic field. This field compensates for the vertical component of the magnetic field of the earth in order to increase the collection efficiency of the photo-tubes. At the detector, the Earth's magnetic field strength is roughly 55 μT and points approximately 15 degrees off the vertical axis. The compensation coils are designed to cancel the field to a value of 19μT, at which point the expected reduction in PMT gain from the zero-field value is 2.5%. The measured increase in photon detection efficiency with and without the compensation coils is about 10%.

The heavy water is contained in a spherical vessel 12 m in diameter (the "acrylic vessel", or AV).[1] The entire vessel is suspended by Vectran ropes from a construction platform which is known as the deck. Looking in at the vessel are 9438 eight-inch photo-tubes. The photo-tubes are mounted on a geodesic sphere of 17.8 m in diameter, called "Phototube Support" or PSUP. It, too, is suspended from the deck.

The area between the AV and the PSUP is filled with light water (labeled "inner H_2O" in the figure), as is the area between the outside of the PSUP and the wall of the barrel-shaped detector cavity ("outer H_2O").

The instrumentation for the phototubes sits on a deck above the detector.

[1]See Fig. 4.1 for a view of the AV and other detector parts that are described in this section.

4.2 The acrylic vessel

A spherical vessel with an elongated neck, the AV functions primarily to contain the D_2O. The vessel consists of 122 individual panels, each of which needed to be small enough to fit into the elevator (the "cage") that transports material from the surface to the 6800 foot level of the mine. The individual panels are approximately five centimeters thick. All materials used in the vessel were assayed to assure their low radioactivity, and the acrylic itself was chosen because its light transmission characteristics closely match Čerenkov light and PMT frequency responses. The AV must be able to prevent D_2O-H_2O mixing and to maintain its structural integrity while submerged in ultra-pure water for the experiment's ten year lifetime.

4.3 The photo-tubes

SNO's approximately 10,000 photo-tubes (PMTs) are eight-inch Hamamatsu R1408 photo-tubes. As with the rest of the detector, PMT materials were chosen to minimize the radioactivity introduced into the detector. The glass envelopes of the tubes were made for SNO by Schott Glaswerke of especially low radioactivity glass and produced in a special clean furnace. Internally, the tubes have twelve stages, and are run at a gain of slightly greater than 10^7. The average dark current noise rate per tube as measured in the detector is approximately 500 Hz, well below the design specification of < 8 kHz. In addition to the inward looking tubes, 91 outward looking tubes ("OWL" tubes) are mounted on the outside of the PSUP. These PMTs provide a means of detecting particles entering or leaving the inner detection volume. In order to extend the sensitivity of the detector to very bright events inside the central volume,

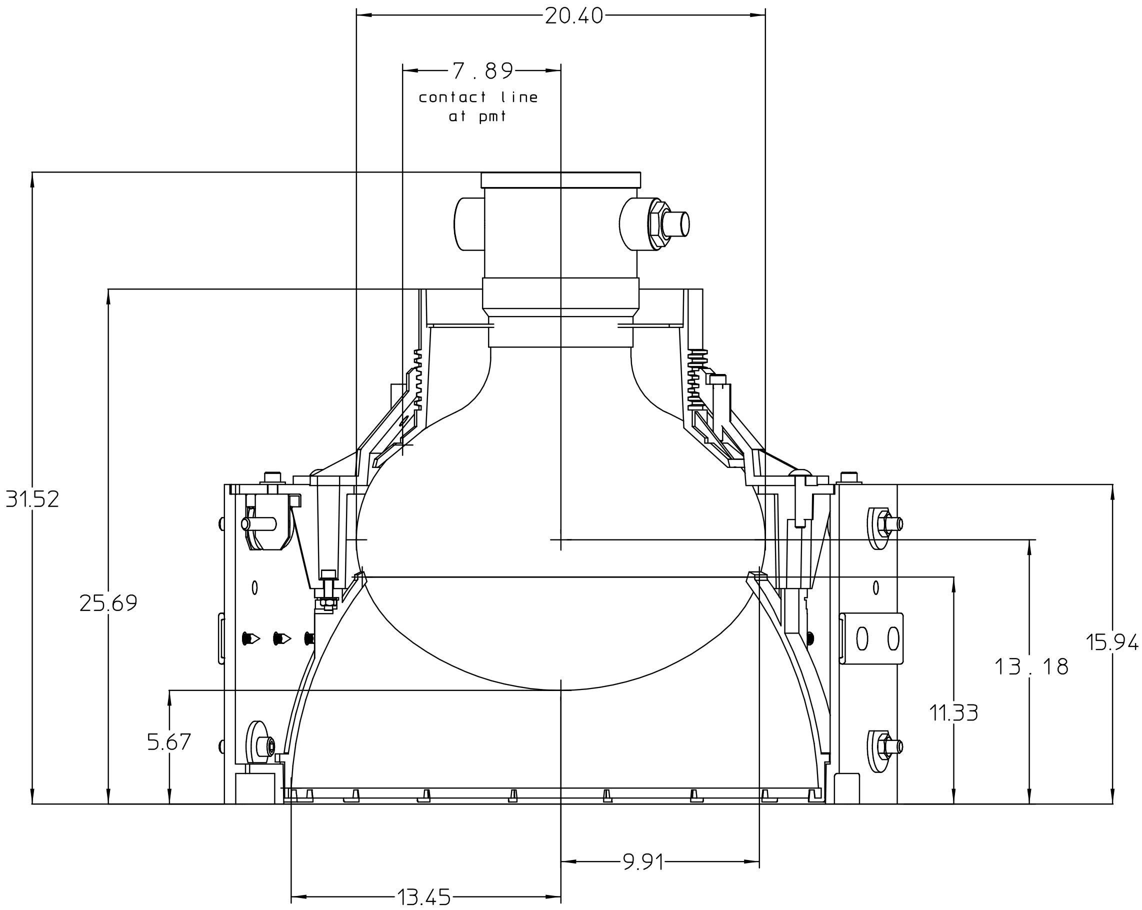

Figure 4.2: SNO photo tube and light concentrator. Figure originally published in [21]. Dimensions are in centimeters. Also shown is the photo-tube base and the HV/signal cable connector.

49 of the inward-looking tubes have a second cable and electronics channel attached to a dynode ("low gain" tubes). The dynamic range of these channels is extended by a factor of 125 from 2000 photo-electrons for normal tubes to 250,000 photo-electrons for these channels. A schematic drawing of the photo-tubes can be seen in Fig. 4.2.

In addition to the photo-tube, the concentrators are shown. Each inward-looking tube is embedded in a concentrator. Without the concentrators, the total photo-cathode coverage of the sensitive volume inside the D_2O is 39%. The concentrators raise this coverage to 54%. As a side effect, the concentrators act as blinders. Their use greatly reduces the number of tubes which can view any individual event at radius larger than $7\frac{1}{2}$ m and also reduces the number of tubes which see radioactive decays in neighboring PMTs (PMT $\beta - \gamma$s).

The transit time of photo-electrons in the tubes is 30 ns, with an average transit time spread of less than $\sigma_t = 1.6$ ns . The narrow spread in transit times allows a precise measurement of the time at which the tube fires and hence good position reconstruction.

A single custom RG59/U-like waterproof cable connects each SNO PMT to the electronics supplying the tubes with high voltage as well as providing the signal path. The cables are connected to the PMT bases via a custom connector.

4.4 Data Acquisition hardware and software

The data acquisition and electronics are designed to accommodate the short high-rate bursts of neutrino data that would occur due to a local supernova, as well as the more routine steady-state low rate data taking of neutrinos due to the sun and radioactive backgrounds. A block diagram of the SNO detector data acquisition electronics can be seen in Fig. 4.3. The electronics consists of two main data streams: the trigger and PMT streams. In order to be able to handle bursts, the electronics is built with several levels of hardware buffering. A detailed description of the flow of data in SNO, along with a typical SNO event, can be found in Ch. 5 and in [22].

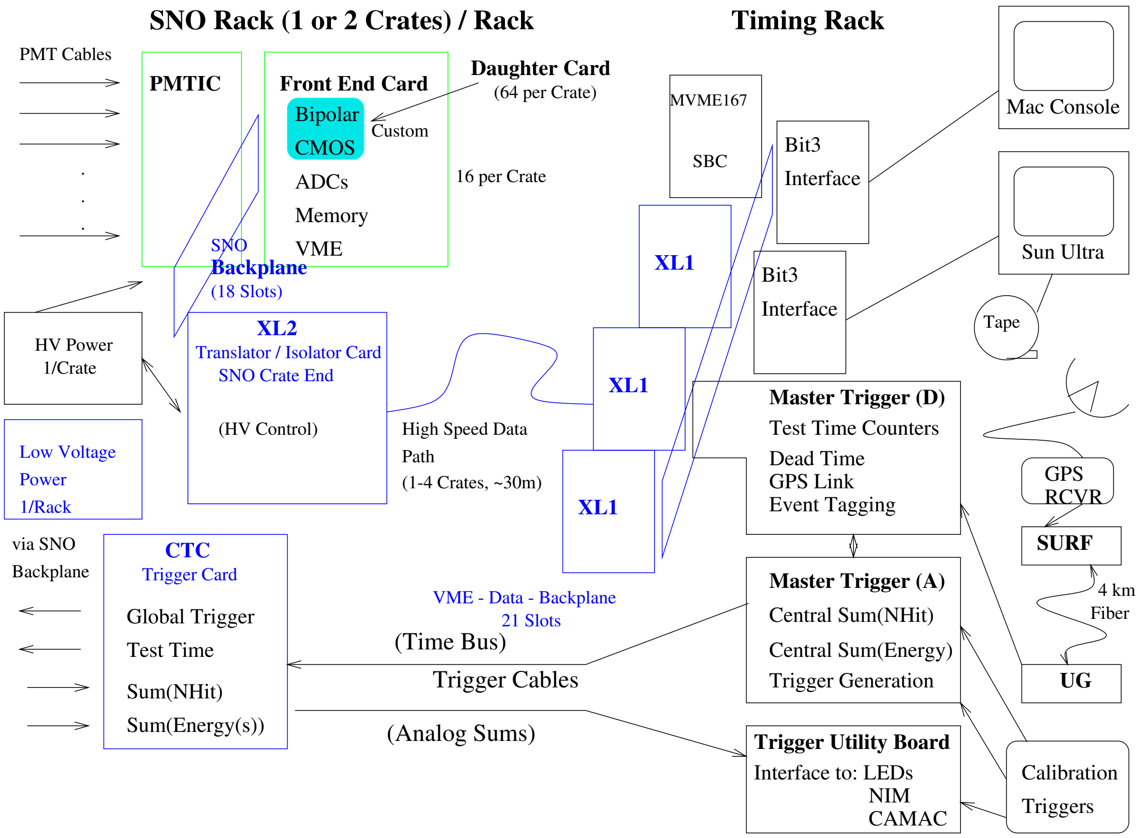

Figure 4.3: Block diagram of the SNO electronics. Data flows from the PMT cables in the upper left corner of the diagram to the tape drive in the upper right corner. Figure courtesy of Rick Van Berg.

4.4.1 Trigger data path

The trigger system must determine when an event worth storing has occurred in the detector, and issue a signal to the electronics to initiate the process of storing the data. The main detector trigger in SNO is the so-called `NHIT100` trigger. This trigger is a simple $N-$fold coincidence trigger. Each PMT channel issues a 100 ns analog trigger pulse when its discriminator fires.[2] These pulses are summed in stages at the board-, crate-, and system-levels. If the summed signal exceeds a pre-determined threshold, a detector wide *global trigger* (GT) is sent to all channels. Each channel that has a discriminator firing within the last 400 ns (the timeout period) stores its data.[3] On all channels, the GT increments a global trigger counter. This counter is used to maintain synchronism between channels.

In addition to the main N_{Hit} trigger, six other trigger types exist. In one trigger, a tighter time coincidence trigger (20 ns rather than 100 ns) with variable delays can be used to selectively image different parts of the detector. By varying the delays between different positions, a narrow coincidence for light emanating from a specific off-center part of the detector can be used to trigger the detector selectively at lower thresholds than the regular 100 ns trigger allows. The normal SNO triggers are synchronized to a 50 MHz clock in order to provide accurate inter-event timing, and tagged with a 10 MHz clock which is fed by a GPS system.

An energy-sum trigger sums a copy of the analog output of all tubes in the experiment. In addition to being useful for triggering events with single high-charge tubes, this signal is one of the single best detector diagnostics and is monitored by

[2]The actual maximum width of the width is 93 ns rather than the designed 128 ns, so the actual pulse width of the `NHIT100` trigger is 93 ns.

[3]If no trigger is received within the timeout period, the buffered data on the individual channel is discarded.

shift operators continuously. The detector can also be triggered on an energy sum or coincidence trigger of OWL tubes. A zero-bias trigger (the so-called *pulse-GT* trigger) runs continuously to record baseline detector behavior in the data stream. In this mode, triggers are set at fixed time intervals, regardless of number of tubes hit. This data allows one to measure the average tube rate in the detector in a non-invasive way.

The trigger system is flexible and allows for a wide variety of different triggering schemes. A large number of external trigger inputs are available for uses such as triggering on calibration sources. A special asynchronous trigger is provided for the timing calibrations. The usual 50 MHz synchronization is incompatible with the PCA timing calibration described in Sec. 4.6.1, and is circumvented by using this `EXTASYNC` trigger during these calibrations.

One of the more unusual external triggers permanently connected to the detector is an experimental hydrophone (aka HIB) trigger, which triggers the detector on local seismic events. Shift operators report correlations between the audible blasts that occur due to INCO mining activities and certain instrumental backgrounds. The HIB trigger triggers on seismic activities and is used to correlate detector behavior with INCO blasting and earthquakes.[4]

4.4.2 PMT data path

Each channel of electronics instruments a single PMT. For each such channel, four quantities are measured: the time of arrival of the PMT signal relative to the detector global trigger, and three samples of the charge integral of the signal. The charge

[4] See [23] for a complete description of the SNO trigger system.

integral is measured at two times (in order to distinguish prompt from late light) and, additionally, at twelve times reduced gain to extend the dynamic range of the charge measurement. The time and charge measurements are implemented in custom ASICs, as are the pre-amp and discriminator and an analog storage array.[5] A state machine reads out the front end through several pipelines, starting with a sixteen deep analog storage array for each channel. Thirty-two channels make up a single front end card. Each of these cards in turn can store 3×10^6 bundles of digitized PMT information locally in 4 MB of standard RAM. The data stored in RAM is read out by the DAQ computers.[6]

4.4.3 Data acquisition software

The data acquisition software consists of a series of custom computer codes which run on a Motorola MVME 167 single board computer (SBC), a Sun Ultra-1 Workstation, and an Apple MacIntosh. The data is collected from the electronics by the SBC and collated into events on the Sun by an event builder. The Mac provides the console and interface for all detector accesses and provides diagnostics and monitoring capabilities.

4.5 Water systems

SNO has two separate water systems for the light and heavy water. Each system is designed to provide ultra-pure water for use in the detector.

[5]See [24, 25] for a complete description of the three SNO custom ASICs mentioned above.
[6]The front end electronics will be described in detail in Ch. 5.

4.5.1 Light Water System

The light water system cleans, purifies, de-oxygenates, cools, and finally regasses water from INCO's water system. The water is introduced in the inner volume (see Fig. 4.1) and forced through a mostly leak-tight barrier on the PSUP into the outer water region, where it is extracted and enters the purification process.

The first stage of this process is a multimedia filter which removes large and fine particles. A charcoal filter follows, which reduces the amount of organic contamination. The water is softened next and sent to a reverse osmosis (RO) unit, which purifies the water of inorganic and organic contaminants. After the RO, the water is sent through an additional batch of filters and mercury lamps to remove any remaining organic components. Next, a Process Degasser (PD) reduces the amount of O_2 and Rn, followed by a fine 0.1-μm filter and a UV lamp to kill bacteria. Before introduction in the cavity, the water is re-gassed with N_2 and chilled to 10°C. A ten ton holding tank is included to allow the system to maintain a constant water level in the detector volume.

4.5.2 Heavy Water System

The heavy water system is similarly organized. The water was delivered to SNO in trucks from Ontario Hydro's Bruce heavy water plant. It then made its way underground in special rail cars.

Since water fill of the AV, the water group has been recirculating the heavy water periodically in order to purify it and measure its radioactivity. The main concern *vis a vis* radioactivity is to reduce the neutron production rates due to photo-disintegration. The goal is to reach levels of 3×10^{-15} g/g Th and 4.5×10^{-14} g/g U in the D_2O . The water group has several complementary methods for performing these assays, and the water is approaching these target levels of cleanliness.

In addition to purification, the heavy water system is also responsible for introducing and removing $MgCl_2$ salt for the second phase of neutron detection (see Sec. 3.4.2).

4.6 SNO calibrations

The process of taking the times and charges of photo-tube hits and relating them to electron positions and energies, and, ultimately, to neutrino properties depends heavily on the SNO calibration program. The calibration program requires measurement of the characteristics of the electronics, the response of the electronics to PMT signals, the optical properties of the acrylic vessel and the light and heavy water, the energy response, and calibration of the detector geometry. When these properties of the detector are understood, SNO will interpret the physical meaning of its data with confidence.

The calibrations for SNO can be divided into roughly five groups: electronic, PMT, optical, energy, and geometry.

4.6.1 Electronics calibration

The electronics described in Sec. 4.4 and Ch. 5 has 19 crates of 512 channels = 9728 electronics channels. Each of these channels must be individually calibrated. Every channel can instrument a PMT. Since there are only 9438 PMTs, spare electronics channels exist. In addition to the electronics, the PMTs themselves must also be calibrated. We first describe the electronics calibration (ECA) and then describe the

PMT calibration (PCA).

Electronics calibration: ECA

The Electronics Calibration (ECA) extracts information that is needed to describe the behavior of an electronics channel. It determines the zero offsets, or pedestals, for each of the three charges and one time measured on a channel. For each PMT hit, three charge measurements and one time measurement are performed. For the pipelining described above, the analog charge and time measurements are stored in a sixteen cell deep local analog storage array, and two time-to-amplitude converters (TACs) are used to reduce the channel-based dead time. Each of these (3 charges + one time) $\times$ 16 cells = 64 cells must have its zero offset measured. In addition, the time slope of the two TACs must be measured.

To allow calibration, the discriminators on each channel have a calibration mode called "pedestal enable." When a channel is put in this mode, its discriminator will respond as if a signal were present at its input when it receives a "pedestal" signal. In this way, the response of the channel to zero charge can be measured.

By varying the arrival time of the pedestal signal and the detector trigger (GT), the timing response of the channel can be mapped out. The trigger system can provide sub-nanosecond control of the timing of the pedestal signal relative to the global trigger. In this way, the ECA is used to calibrate each channel's TAC. In data analysis, the ECA calibration constants are the first step in the process of converting digitized values into physical quantities, such as ADC counts into nano-seconds.

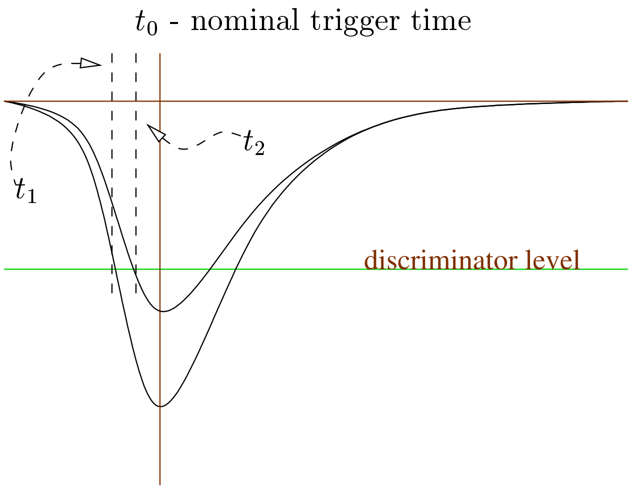

Figure 4.4: Time Walk of PMT signals

Phototube calibration: PCA

The Phototube calibration is performed to describe the behavior of individual phototubes and their response to single photon light. For each channel, the discriminator walk is characterized. Walk is the slew whereby signals of varying amplitudes appear to have arrived at different times due to the differences in the pulse shapes. Fig. 4.4 shows two PMT signals of different amplitudes. As drawn, the larger of the two signals will register at time t_1, while the smaller one will register at time t_2, though both nominally occurred at the same time t_0. By examining the behavior of a known source of single photons, we can measure the magnitude of this charge-based time variation for each channel. Once this behavior is understood, the charge-dependent behavior can be accounted for. Finally, the stray capacitance of the discriminator itself can cause an additional time adjustment. For signals just above threshold, this time correction can be as large as 5 ns.

A different walk correction must be applied at high charges. Each photon that hits the PMT generates a signal that will be measured at a time that is normally distributed around the PMT transit time with a Gaussian width that is called the PMT transit time spread. In the case of multiple photons hitting the photo-cathode, the discriminator measures the signal that arrives earliest. This process leads to a slow creep where the measured signal time moves to earlier times at higher charges. This high charge walk is also corrected in the PCA.[7]

In addition to the walk correction, the PCA measures two quantities. First, the relative timing offsets between channels due to the length of the cable connecting the channel to the front end electronics is measured. This relative offset between channels is removed in the PCA calibration.[8] Second, by comparing the signal size between channels, the relative gain of each PMT is measured.

The "laser-ball" source used for this measurement is described in Sec. 4.6.2 and elsewhere [21, 26]. The main source requirement for this measurement is to be able to emit light of a time duration short compared to the intrinsic time spread of the photo tube, of varying intensities, and at a well-defined position. The laser-ball fills these requirements. An alternate sonoluminescent source has also been used to study the timing response of the SNO PMTs.[9]

[7] Due to the somewhat poor charge resolution of the photo-tubes, there exists a charge regime where it is unclear whether to apply the low or high charge walk correction. This decision must now take additional factors such as the total charge in the event or the distribution of hit photo-tubes relative to this one into account.

[8] Additional delays due to the travel time in the of the photo-tube signal in the electronics are also removed.

[9] See [27, Section 3.4].

4.6.2 Optical calibration

All properties that affect the travel of light in the detector are corralled under the rubric "optical calibration." The optical calibration is used to determine the attenuation, scattering, and absorption lengths of heavy and light water in the detector, as well as the detailed optical properties of the acrylic vessel panels. As these optical properties vary substantially with wavelength and Čerenkov light has a broad spectrum, the optical calibration must be performed at several wavelengths to map out the behavior of Čerenkov light in SNO.

The laser-ball system mentioned in Sec. 4.6.1 is intended to fill this function. The laser-ball is a fed by a 337 nm laser, which in turn can be used to pump dye lasers from 360 – 700nm, thus emitting nearly monochromatic light of different wavelengths. The laser light pulse travels down an optic fiber bundle to a light diffusing ball to produce nearly isotropic light in the SNO detector. The time spread of the 337 nm light pulse is 600 ps [28]. The optical calibration follows the strategy of studying the hit distributions in the detector as a function of time and position. To determine the attenuation length, laser-ball data is taken at different positions. Given a known laser intensity, the change in the number of hits in a narrow in–time window as a function of distance allows us to determine the extinction length. Scattering and reflections cause structure in the time distributions of the hits. This structure can be used, for instance, to understand the amount of Rayleigh scattering in the water. The optical properties of the acrylic panels were measured before their installation underground and will be verified in the *in situ* measurement.

At the time of this writing, the optical calibration is still a work in progress.[10]

[10] A detailed proposal (mostly historic at this point) is [29].

Source name	detection reaction	Energy in SNO
^{16}N	$^{16}\text{N} \rightarrow {}^{16}\text{O} + \beta + \gamma$	$E_\gamma = 6.1$ MeV
pt	$t + p \rightarrow {}^{3}\text{He} + \gamma + n$	$E_\gamma = 19.8$ MeV + neutrons
^{8}Li	$^{8}\text{Li} \rightarrow \beta + {}^{8}\text{Be}^{*}, {}^{8}\text{Be}^{*} \rightarrow 2\alpha$	E_e, β-spectrum, ~ 13 MeV endpoint

Table 4.1: Types of energy scaling sources in SNO for electrons

4.6.3 Energy response

After the optical and electronic properties of the apparatus have been measured and the Monte Carlo has been calibrated with the results of these measurements, the response of the detector can be tested using each of several radioactive sources. A comparison between the response of the calibrated Monte Carlo and the source data will determine how well we are able to model the detector. This test allows an independent verification of the optical parameters and the SNO Monte Carlo at several energy points. The response of the detector to neutron-induced γ-rays and *vs* its response to electrons can also be tested this way. Tab. 4.1 lists some of the sources in SNO's arsenal.

4.7 SNOMAN: the SNO Monte Carlo package

All analysis in this thesis was performed with the official SNO analysis package, SNOMAN (SNO **M**onte Carlo and **An**alysis package), mostly with development releases in the 3.0185–3.0192 range. SNOMAN is a fully featured, integrated analysis package written in the `FORTRAN 77` programming language. SNOMAN simulates e, μ, n, and γ physics and optical photon propagation in the detector. The full geometry and optical properties of the detector are modeled using bench measurements of detector component properties initially and *in situ* calibrations when available. SNOMAN also simulates the electronics and DAQ, including, for instance, crosstalk.

SNOMAN relies on several external software packages. Most importantly, the transport and interactions of γ and e^-, are simulated using the EGS4 package [30]. Extensive use is made of the CERNLIB software packages. SNOMAN overcomes `FORTRAN`'s lack of built-in memory management tools with the CERNlib ZEBRA package. The CERNlib HBOOK package allows storage and manipulation of histograms and n-tuples for data analysis and display.

The management of these constants and all time-varying data is provided by SNODB, an interface to HEPDB, a CERNLIB-based `FORTRAN` database system developed at Penn for SNO. The state of the detector, as recorded by the DAQ software, is available to the data analyzer if she needs it via a set of banks by the name of "DQXX," which are stored in the data base. The collaboration-standard data reduction information is communicated to the user via a set of DAMN (**D**ata **A**nalysis **M**ask **N**umber) banks, which are also distributed via SNODB.

SNOMAN is fully documented, with a web-based reference guide (the "companion") and user's and programmer's manual [31, 32].

4.8 Summary

This chapter introduced the SNO detector, describing the physical detector parts. Both the calibration system and the analysis software were discussed. Finally, the physics signals SNO was designed to measure were outlined.

As mentioned in the introduction, SNO's ability to measure both ν_e and ν_x neutrinos independently are its greatest strength. The extreme care the collaboration has taken during construction will allow analysis thresholds well below 5 MeV. The next chapters will discuss a preliminary measurement of the CC integrated flux, which will be the first measurement of the ν_e -only flux from the sun. Before discussing the CC measurement, an outline of the data flow in the detector will be presented.

Chapter 5

Data flow

This chapter describes the hardware and software data acquisition system at SNO. The emphasis is on the physics goals that the design tries to achieve, and the design choices that are made to reach these goals. Some of the work I have performed on this system will be described here, and in more detail in the Appendices A–C.

5.1 Physics goals motivate electronics design

The design of the electronics and DAQ systems at SNO is determined by the physics goals of the Observatory. SNO is designed to take solar neutrino data continuously and to be capable of handling a burst of data in case of a nearby supernova. These goals put two requirements on the DAQ hardware and software.

- **Steady-state**: Continuous throughput of N_{Hit} data at a trigger rate of 100 Hz, due to solar neutrinos and detector backgrounds.

- **Burst**: Burst throughput at high event rates of at least 1 kHz for short periods, due to supernova neutrinos.

PMT data		
3 LW	96 bits	12 bytes
MTC Event data		
6 LW	192 bits	24 bytes
$N_{\mathrm{Hit}} \sim 30$ at 0.1 KHz		
$\rightarrow$ 384 bytes/event $= \frac{3}{8}$kB/event		
$\rightarrow 3.75 \frac{\mathrm{kB}}{\mathrm{s}}$		
$\rightarrow \sim 3 \frac{\mathrm{GB}}{\mathrm{day}}$		

Table 5.1: Size of the SNO data stream used in the electronics design. The actual data rates are roughly 1/5 of this estimate.

No aspects of the DAQ or electronics should limit the ability to do physics analysis - rather, any such limits should be due to intrinsic limitations of the detector itself, such as radioactive backgrounds or diminishing signal-to-noise ratio. Present indications are that these goals have been met and that SNO could take data at a much higher rate. The SNO rates are low compared to those of collider experiments. As displayed in Tab. 5.1, the expected average data rate at SNO is 37.5 kb/s. Unlike collider experiments, the SNO data rate is asynchronous, which brings some different challenges.

5.1.1 Solar neutrinos

At SNO, the expected solar neutrino rate is very low. The total event rate due to solar neutrinos in the SNO detector is roughly 40 events per day.[1] In order to maximize the solar neutrino data rate, the detector electronics must achieve 100% live time and no event-based dead time. This includes the ability to take physics data while doing electronics and physics calibrations, and to take data with subsets of the detector if

[1]This estimate is based on the unmodified SSM. See Ch. 3.

the whole detector is not available.

5.1.2 Supernova neutrinos

The picture looks a bit different for supernova neutrinos. While the understanding of supernova dynamics is not complete, most models predict a burst of neutrinos preceding the light emission by a few hours. The burst is expected to last about 10 seconds, with an initially high flux of electron neutrinos, followed by a flux of neutrinos and anti-neutrinos of all flavors after a few tens of milliseconds. The energy scale of electron neutrinos is expected to average to 15 MeV , while the other neutrinos are expected to have a somewhat higher average of 20 MeV.[2] The structure of the burst has a lot of information about supernova dynamics and mass differences between neutrino flavors in it, and it is clearly best for SNO to capture as much of this data as possible. Therefore, for supernova neutrinos, SNO must be capable of acquiring short bursts of high rate data without loss. The DAQ system achieves this goal via a series of hardware pipelines.

5.2 An event at SNO

What follows is the description of a typical SNO event, from a neutrino interaction to the final analysis on a physicist's workstation. I will present a short overview of the elements of the SNO DAQ hardware to simplify the discussion, and then follow with a blow-by-blow description of a SNO event.

[2]See [10, 33] for a discussion of supernova neutrinos at SNO.

5.2.1 SNO DAQ Hardware Overview

The hardware consists a timing rack and eleven PMT racks. The timing rack has in it

- A standard VME crate, which has in it
  - 1 Master Trigger Card, Digital (MTCD). The MTCD is the digital part of the trigger system.
  - Five VME-SNObus translator cards (XL1). This is the VME side of the interface between the VME crate and the SNObus custom backplane. Each XL1 can address up to four SNO crates.
  - One Motorola MVME 167 Single Board Computer, part of the DAQ system
  - Two Bit3 VME extender cards, to allow communication with
    - the DAQ console, an Apple MacIntosh computer
    - the event builder computer, a Sun Solaris workstation
- The SNO analog trigger crate, which has in it
  - 8 Master Trigger Cards, Analog (MTCA). The MTCA is the analog part of the trigger system.
  - One Analog Measurement Board (AMB). The Analog Measurement board measures the integral, derivative, and peak of the ESUM trigger.
- Two Trigger Utility Boards (TUB). The TUB is a general-purpose board. It provides
  - The SNO Master 100 MHz clock

- ECL-TTL translation
- BNC to SMB interconnectivity
- Trigger fanout for monitoring purposes
- A speaker connected to the main N_{Hit} trigger, for monitoring purposes
- a discriminator with an manually adjustable threshold

The SNO crates are spread across eleven racks, with one or two crates each. There are 19 SNO crates in total. Each contains

- Sixteen Front End Cards (FECs). Each FEC handles 32 channels and has on it
  - a state machine to manage the board's channels, known as the "sequencer"
  - four 12 bit analog-to-digital converters (ADCs)
  - four daughter boards (DBs). Each DB has on it
    - Eight channels of electronics, consisting of three custom ASICs.
      1. A discriminator (SNODISC)
      2. An integrator (SNOINT)
      3. A time-to-amplitude converter and analog storage array (QUSN7, usually referred to as the CMOS chip.)
  - four MBs of DRAM memory
  - A state machine to manage the signal digitization
  - A FIFO memory controller
- 16 PMT Interface Cards. Each PMTIC has up to 32 PMT cables on it, and provides high voltage to the PMTs and connects directly to the FECs above. The PMT signals are transformer coupled to the FECs via the PMTIC.

- One Crate Trigger Card (CTC), which does the crate-level summing of the trigger signals and distributes the detector-wide Global Trigger (GT) to the FECs.

- One VME-SNObus translator card (XL2). This is the SNObus side of the interface between the commercial VME crate and the SNObus custom backplane.

5.2.2 The event

As an example, let us follow a neutrino interaction in the SNO detector, though any other light source will produce the same sequence. The event starts when a neutrino interacts. The neutrino interaction creates Čerenkov radiation, which is then detected by some subset of the SNO PMTs.

At the Front End

The hit PMTs produce a pulse which travels along the cable to the front end (FE). At the FE, the SNODISC fires if the signal is above the discriminator's roughly 1/4 photo-electron (PE) threshold.[3] The SNODISC signals the start of a possible event to both the QUSN7 TAC and to the SNOINT.

The QUSN7 CMOS chip produces a `TR100` trigger signal, roughly 100 ns in width, which is sent to the trigger system. At the same time, it also starts a timing ramp timing ramp (time-to-amplitude converter, or TAC) on each channel to measure the time of the PMT's firing. The SNOINT measures the integral charge of the signal that caused the discriminator to fire. This charge measurement will be used to correct charge-dependent time walk described in Sec. 4.6.1. The SNOINT measures charge in

[3] A PE is defined as the amount of charge generated by a single photon striking the photocathode of the PMT.

three ways.[4] First, the charge is sampled at a time consistent with one photon hitting the PMT. This is the main measurement. A second measurement is made near the end of the time window, with the same sensitivity, but with a longer sample time. This long sample measurement indicates whether the tube was hit by late light, i.e., photons that hit the tube after the first charge measurement was made. These charge measurements have a dynamic range of about 1/30 to 100 PE. An additional charge measurement is made with lower resolution but much larger dynamic range, up to 1200 photo-electrons. This is the third measurement.

When the GT signal arrives at the FEC, every channel whose discriminator fired less than the timeout time T_{tot} ago, stores its values. The timers which were started when the individual PMT's discriminators fired now all stop at the same time, giving them a common reference point.[5] The three analog charges and the TAC signal are stored in an analog storage array. This array can store up to sixteen events internally before its storage space runs out. This is the first level of buffering, allowing the system to store roughly 5000 events at the very front end.[6] The analog data is tagged by a global trigger ID (GTID), which is simply the number of global triggers since the beginning of the run. Each channel with data asserts a signal that indicates that it can be read out.

The Trigger

Concurrently, the trigger system determines whether the event should be recorded. The trigger system sums trigger signals in stages, first from individual tubes, then

[4]For a description of the ASICs built by Penn for SNO, see [25].

[5]i.e., SNO timing is a *common-stop* system.

[6]N = (9500 PMTs $\times$ 16 cells/PMT)/(30N_{Hit} /event) $\simeq$ 5000 events in all cells, for optimal coverage.

from daughter boards, FECs, electronics crates, and finally for the entire system. The trigger sum is analog, since the system as designed requires that the trigger system return a signal to the FECs within a few hundred nano-seconds before the individual PMT data is discarded. A fast analog sum is better suited to this task than a digital sum, though additional precaution must be taken to insure that pickup does not destroy the trigger resolution.

At the central timing rack, the final detector-wide sum of trigger signals is performed on the Master Trigger Card, Analog (MTCA). The outcome is compared to a threshold value, and, if the signal is above threshold, the MTCA signals its digital counterpart, the Master Trigger Card, Digital (MTCD), to store the event. The event time is synchronized to a Global Positioning System (GPS) 50 MHz clock, which gives the event its universal time. The MTCD records event header information, such as the time at which the event occurred and which triggers are currently above threshold, and produces a signal called the global trigger (GT) which is distributed to each FEC. The MTCD has a mega-word buffer and can store a million events' worth of data before needing to be read out.

The Middle Ground (between the Front and Back ends)

At this point, the data must be converted to a digital form such that the DAQ computers can read it out and process it. The 32 channels of each FEC are processed in turn by one set of ADCs. The sequencer, which is described in detail in Appendix A, controls the digitization of the data from each channel and writes the data relevant to a single PMT hit, including the geographic information which identifies which PMT fired (a "PMT bundle"), into the local 4 MB memory. The memory management is performed by a commercial dual-ported memory controller and some logic that makes the memory look like a large first-in, first-out (FIFO) stack. The FIFO logic is described in detail in Appendix B. The 4 MB local memory on each FEC can store up to $\sim 3 \times 10^5$ PMT hits. This is the second level of hardware buffering.[7] The sequencer writes the data into the local DRAM, and the FIFO controller signals the availability of data on this FEC to the DAQ computers via a register on the SNOBus side of the VME-SNOBus translator, the XL2.

Though the hardware buffer on the individual FECs is smaller than the buffer on the MTCD, as FECs do not, on average, participate in each event, during neutrino data taking, the choke point is not at the FECs but at the MTCD.

Event readout and collation

At this point, the DAQ computers can read out the data. The DAQ Motorola MVMe 167 Single Board computer runs an "EVREAD" program which continuously polls the crate 'data available' register on the XL2s in each crate at one Hz to determine which FECs have data. In turn, the MTCD and each FEC are read out and the event header and PMT data bundles are transferred from the electronics' hardware memory to a buffer in the event builder computer.

The event builder and EVREAD communicate via a dual-ported memory which sits in the Bit3 card. The event builder collates the PMT bundle information with the appropriate event header from the MTCD. When all data for an event has arrived, the event builder writes the data to disk in ZDAB files. ZDAB is a ZEBRA-based format which is suitable for reading by the SNOMAN analysis package. The data is written to disk. Each disk file contains data from at most one run, and is called a

[7] I was primarily responsible for both the sequencer and the FIFO controller, as well as the digital interface from the VME backplane through the XL cards to the FE card.

"sub-run". Each sub-run is at most 200 MB large to facilitate later manipulation.[8]

To Tape

SNO uses DLT tapes as its archival medium. When the event builder closes the data disk file, it creates a marker file which indicates that the file can be moved to long term storage. The process which manages the long-term storage on tape is called 'data flow', or 'dflow.' The dflow process' goal is to copy the data to tape in two locations, underground and above ground, while always keeping at least two copies of all data at all times to protect against data loss. Equally important is the use of cryptographic-style checksumming to insure that the data is not corrupted.

Dflow's first task is to generate the cryptographic checksum, or hash, which uniquely identifies the file. The MD5 algorithm is used for this purpose. This algorithm generates a unique 128 bit 'signature' based on the contents of the data file. Next, the file is written to a tape drive which is connected to the underground DAQ Sun.[9] Once this is finished, the marker file, signature, and data file itself are copied to a second workstation which sits on surface via fiber-optic Ethernet.

The surface dflow machine generates a second MD5 hash of the ZDAB file and compares it to the hash which has been copied from underground. Small changes in the file's data generate different signatures. Any changes in the signature indicate data corruption. Upon detection of corruption, the dflow process is halted and a human needs to intervene to clean up the mess. In the usual case where the signatures agree,

[8]In SNO, a run is at most 96 hours long. Operators start new runs earlier when someone or something modifies detector conditions, such as water system changes or equipment failures, or when the calibration group inserts a source.

[9]The underground tape copy is a backup copy and is only used if the above ground copy gets damaged. The underground tapes are eventually cycled to above ground to be overwritten in the above ground process.

the data is now written to the above-ground archival tape. A secondary process copies the data to a surface analysis machine for near-line data processing.[10]

A tertiary dflow process duplicates full tapes, which are distributed to all member institutions of the collaboration. SNO uses 15 GB tapes, which are currently filled at a rate of roughly one per 3 weeks. We have written a `tape_extract` program which is available to the collaboration to get the data off the tape. The program queries a globally available data base to determine which tape the run is on, and where on the tape it sits. It then extracts the data from tape to disk, and re-calculates the hash and checks it against the one stored in the data base to insure correct extraction. `tape_extract` is a mixed `C` and Bourne-shell set of programs.

5.3 Near-line data analysis and automated processing

SNO uses near-line data analysis for data quality monitoring, processing of routine electronics and PMT calibrations, and a fast first pass data reduction. Near-line analysis is managed by a framework called `SNOMAN Application Supervisor (SAS)`. SAS can operate on sub-runs (the smallest unit of data) and runs, depending on the application.

5.3.1 List of SAS tasks

- **Near-line detector monitoring (Auto-SNOMAN or `data_inspection`):** Auto-SNOMAN makes histograms such as event rate, flasher rate, occupancy

[10] I was instrumental in the development and commissioning of the dflow process, and have been its maintainer since the commissioning phase of the detector.

maps which are diagnostic of the quality of the data. Auto-SNOMAN runs as soon as possible after the data arrives at the surface analysis machines, and the output is monitored by shift operators as part of their data quality tasks.

- **Automated calibrations**: SAS automatically performs routine electronics (ECA) and PMT (PCA) calibrations when it detects such data on disk. When the analysis tasks have completed, calibration experts receive an email notification and inspect the results of the calibrations. With the blessing of the experts, the good calibration constants can now be entered into the official SNODB data base.

- **Data cleaning, or DAMN bank generation**: For rapid data analysis, SAS generates the DAMN data mask banks described in Sec. 4.7. These banks allow the physicist to do rapid turn-around analysis.

- **Runlogger**: Similar to data inspection, the Run Logger generates data summary histograms which are stored in SNODB, the SNO data base.

SAS is the final stage in SNO's data flow. It is implemented as a set of PERL scripts. I contributed to SAS and performed a large amount of the Unix system interface programming.

5.3.2 SunDAQ: Electronics Commissioning Software

During the commissioning phase, the Penn group wrote a complete DAQ software system. This was necessary to debug and commission the electronics, and the effort produced a system that was fully functional in all respects. The SunDAQ system was also used for initial data taking with the completed detector, and is still used by the electronics experts to debug electronics problems.

5.4 Summary

In this chapter, we described the SNO detector's data acquisition system. We first explained the motivations for the design, and then showed how they were implemented. In order to fill in the details, in the rest of the chapter we gave a walk-through of an event from the PMTs firing to the data being extracted from tape for analysis. This chapter introduces some of the systems I have worked on during the SNO construction and commissioning phase. In the next chapter, we will present SNO data.

Chapter 6

First analysis of SNO data

In this chapter, we introduce the step-by-step process of determining physical quantities from the raw data. This includes data reduction to remove instrumental backgrounds, background measurements to estimate the amount of non-neutrino data in the data set, calibration to turn the raw data into PMT-level physical quantities such as charges and times, and reconstruction to turn those PMT-level quantities into event-level quantities such as the energy of the candidate electron. The physics and instrumental backgrounds will be handled first, followed by a brief description of calibration and event reconstruction, and finally a summary describing the sensitivity of this measurement.

6.1 Background Estimation

SNO's backgrounds can be divided into two categories. The first, instrumental backgrounds, refers to events that are induced by some part of the detector itself or its surroundings and are not caused by any physics event. The second class, physics backgrounds, covers events that are caused by physics processes that are not the physics process the current analysis is interested in. A physics background could be a radioactive decay of a contaminant in the detector, or it could even be a signal that is otherwise of interest to the solar neutrino program but is not being measured in this particular analysis. We will describe each type of background in turn.

6.1.1 Instrumental backgrounds

In this section, we will first describe SNO's approach to eliminating instrumental backgrounds. We then then discuss in some detail the types of background events and how they are tagged.

6.1.2 Data cleaning

The job of removing instrumental backgrounds for the collaboration is handled by a subset of the collaboration called the data cleaning group. While many members of the collaboration helped in this process, the core consisted of a small group of which I was a part. The data cleaning group had the following goals.

1. Develop a set of simple cuts to tag instrumental backgrounds.
2. Determine the amount of sacrifice, i.e., the amount of signal data that is tagged by the data cleaning cuts.
3. Determine the amount of residual contamination, i.e., the amount of background that leaks past the data cleaning cuts.
4. Publish a default mask of cuts with understood contamination and sacrifice.

Note that we decided that the data cleaning group should not *remove* events, but rather just tag them. As described in Ch. 5, the SNO data rate is sufficiently low that the complete data set is not too large to overwhelm a SNO institution. Publishing a mask rather than cutting data and producing data summary tapes allows easier systematic studies of the biases of the data cleaning. In addition, as the publication mechanism is via the SNODB, using DAMN banks also reduces the pain of redistribution if the data must be reprocessed.

How the cuts were created

The cuts were developed and tested against "golden" data sets of particular instrumental backgrounds. These data sets were either hand-assembled or assembled with preliminary versions of the cuts, and we then optimized the cuts to tag as much instrumental background as possible.

The cuts were then tested against various signal-like data sets, such as the laser ball and the ^{16}N data, to see how much signal they will cut. Each cut was then tuned to not cut more than 0.1% of the ^{16}N signal as a function of radial position in the detector. The data cleaning group decided that this 0.1% was the maximum allowed sacrifice.

Biases in position and N_{Hit} were also examined, and based on these studies, cuts were either refined or evolved into more sophisticated ones.

Finally, the data cleaning group applied the cuts to a small N_{Hit}-triggered data sample and examined the resultant reduced data set and the cut data set to get a preliminary estimate of the contamination and sacrifice.

6.1.3 Types of Instrumental backgrounds

The instrumental backgrounds in SNO can be roughly grouped into those that produce light and those that do not. The former category includes flashing PMTs, some wet end breakdown, and light produced by what we think is static electric discharges across insulators in the detector.

The latter category contains events that cause the electronics discriminators to fire though the PMT does not see light. This can be caused by breakdown in the dry or wet end PMT connector or pickup in the electronics due to activity in the deck area.

Most instrumental backgrounds do not fit into these categories very cleanly. For instance, breakdown in the wet end can cause light in the detector, but the large PMT signals that breakdown generates also causes pickup on the deck due to the 2000 V pulse on the PMT signal cable.

Flashers

The PMTs in SNO flash, that is, emit light spontaneously, as do most PMTs. Each tube will flash roughly once a week, which results in a detector-wide flasher rate of one per minute. A flashing tube can illuminate as many as 300 other PMTs. Due to the high rate of flashers, cuts to eliminate them are among the most important for SNO.

Most flasher events have a distinctive topology. In such an event, one set of PMTs on one side is hit, and a second set of tubes across the detector is hit one light-crossing-time later. The first set of tubes will have one tube with very high charge (the flasher tube) and a collection of adjacent tubes with charge consistent with zero (pickup from the flasher). The light across the tank has a characteristic elliptical pattern, and those tubes have charges that are consistent with single PE light. The number of tubes hit by flashers is in the range of $N_{\text{Hit}}\, 20 \sim 300$, which is in the same range as is expected for solar neutrinos. The light is emitted over a somewhat longer period than is characteristic for Čerenkov light [34].

Several strategies have been developed to tag flashers. The first flasher cuts were developed based on marking the high charge tubes. These *charge* cuts are effective at catching flashers, but also tend to tag good events with miscalibrated tubes as flasher events. For instance, during a particular period, 7% of the data was tagged by a charge cut as flashers due to a FEC that was intermittently producing miscalibrated data. More effective cuts were needed.

The next set of cuts tagged the pickup cluster around the flasher tube (*cluster* cuts). This cut is similarly efficient against flashers as a charge-based cut, but is less susceptible to miscalibrated tubes. However, this cut also has problems. Due to the cabling of the detector, tubes that are close to one another tend to be close to one another in electronics address space as well. Therefore, higher N_{Hit} events or events close to the PSUP are more likely to have a cluster of tubes that look like a pickup cluster from a flasher. For the ^{16}N, it was determined that $\frac{4}{10}$% of all events at $r = -500$ cm (near the bottom of the detector) that pointed outwards would be cut by this cut [35]. This radial bias eliminated the cluster cut.

By demanding the coincidence of the charge and cluster cuts, one can further reduce the amount of signal data that is tagged (**QCluster cut** [36]). Other cuts that tag flashers take into account the characteristic charge *vs* time behavior of flashers described above(**QvT cut** [37]).

A second class of flashers exists which is harder to find. A "blind flasher" is a flasher where the signal path for the flashing tube is broken or accidentally disconnected, but the HV path is intact. Without the signal path but with voltage, the tube can still flash, but the large charge tube and pickup channels will not be recorded in the data. This can also happen when the operators disable a channel but do not turn off its HV.

Several strategies have been developed to deal with this problem. An obvious method would be to use reconstruction to determine that such an event occurred on the photo-sphere, and reject events at those radii. However, the data cleaning cuts are supposed to be simple, and reconstruction is inappropriate at this stage of the analysis.

The one cut for tagging blind flashers takes advantage of the long flasher light emission time. By looking at the median time difference between pairs of adjacent tubes, this difference in emission time can be resolved and blind flashers can be flagged, albeit not with a very high efficiency (**FTS cut** [38, 39]).

Bursts and Retriggers

Certain types of background tend to occur in bursts. Pickup events, breakdown in either the dry or wet end, and neck events all tend to occur in groups during very short time scales. By tagging and cutting events if more than three events happen within 1 ms, we eliminate these backgrounds.

A second version of this cut tags events if more than two events are recorded within 5 μ s. In very large events, the detector tends to retrigger due to residual light or pickup. This cut eliminates these events.

A third burst cut eliminates events when six occur with $N_{\mathrm{Hit}} > 40$ in less than four seconds. Due to the nature of all SNO physics events, including radioactive backgrounds, this is extremely unlikely to occur due to solar neutrinos, and some instrumental backgrounds fall into this category. The low regular trigger rate for the experiment (less than 20 Hz) means that physics events very rarely come in short succession and hence the cut has a low (and calculable) sacrifice. Clearly, supernova physics has different time structures and supernova analyses will not use a straight N_{Hit} burst cut.

In all burst cuts, the first event of a burst or retrigger is *not* tagged as a burst but tagged separately as the initiator of a burst [40].

Neck events

A background SNO has to contend with are "neck events," where a short light pulse is emitted in the neck region of the acrylic vessel. Due to the opacity of the neck, the light is beamed downwards and tubes at the bottom of the vessel are illuminated. In order to help tag these otherwise reasonable looking events,the collaboration has installed two PMTs in the neck region near where the light appears. The tubes detect neck light with high efficiency. The actual cut will mark an event if both neck tubes fire, or if one of the tubes fires with large charge and early time with respect to other tubes in the event [41].

Pickup

There is a large class of events that look like pickup in the detector. Due to the low PMT thresholds, the PMT discriminators are very sensitive to electrical noise in the lab.[1] The topology of these events is distinctive. For instance, breakdown in the PMT

[1] The individual PMT discriminators trigger on signals equivalent to roughly 1/4 photo-electron, which is equivalent to a roughly 2.5 mV signal at the front end.

wet and dry connectors causes breakdown with the following topology.

- The events occur in bursts.
- The hit patterns spread from one crate, to uniformity across the detector during the burst.
- The charges measured by the PMT integrators are consistent with zero charge.

Other singleton pickup events do not occur in bursts, but do have the same spectrum of PMT charges.

Several cuts address these topologies and use them to eliminate pickup events. The burst cut clearly will cut some of the pickup events. The $\mathbf{Q}/N_{\text{Hit}}$ cut tags pickup events by their low total charge relative to the number of tubes, as measured by the integrators [42]. The **AMB** cut is similar , but uses the total event charge as measured hardware by the Analog Measurement board (AMB) rather than by the integrators [43]. Finally, the **Crate Isotropy** cut tags events where most of the data occurs in one electronics crate.

6.1.4 Other Cuts

Other cuts listed here are important for SNO, but do not affect as large a fraction of the initial data stream and hence are relegated to a separate section.

OWL cut

SNO has 91 **O**ut**w**ard **L**ooking tubes, which are tubes that sit on the PSUP but look away from the heavy water rather than at it. These tubes see events that occur in the light water outside the PSUP. If more than five of these tubes fire, the event is tagged [44].

"Junk"

There is a class of events that are called *Junk* events. Junk events are events where a single PMT occurs more than once, and 'orphanages.' As described in Sec. 5.2.2, an orphan is a single PMT bundle which does not have any associated event information. These events occur at a low rate due to a flaw in the QUSN7 CMOS chip.[2] In other, rarer cases, orphans can occur during high rate data taking when PMT bundles appear at the event builder after the rest of their associated event has been flushed from the event builder's buffer. These bundles are now orphans.

Such PMT bundles are collected into "orphanages" and written to disk when they reach a certain size. The junk cut tags these events.

Wet End breakdown

A flaw in the SNO PMT base is the underwater connector which connects the cable to the PMT base. The connector does not hold high voltage reliably in ultra-pure, degassed water. During the detector water fill, when the first tubes were immersed and turned on, we immediately found that the underwater tubes were breaking down. This was a surprise to the collaboration, as the entire PMT-cable assembly had been tested in a lab environment in an underwater tank with many year's equivalent exposure time and the testers did not report any problems. However, the test tank had used regular water, rather than the ultra-pure, degassed water that went into the SNO

[2] These are sometimes called "double-zero" orphans, as the GTID in the PMT bundle is typically of the form `0x???00`, where the `?` is replaced by a valid hex number. See App. C.

detector.

A SNO physicist realized that the degassed water would cause a pressure reduction in the open volume connectors to reach equilibrium with the gas pressure in the water water. In an effect known as Paschen's Law, the breakdown voltage in the connector is related to the particle density or pressure in the connector. At low pressures, the electrons from molecules that ionize can acquire sufficient velocity before they interact with other molecules such that during these collisions, a small cascade is produced, leading to a non-negligible amount of current flowing in the signal path. Via this effect, the degassed water caused a reduction of the breakdown voltage in the connector to the point where a large number of connectors were failing.

The water group degassed the water to impede biological growth and to remove radon, an important background to the neutral current measurement. By re-gassing the water with nitrogen after the initial degassing process, the density of gas (and the pressure) inside the connectors increased and the breakdown voltage rose to the point where breakdown of the wet end is a rare occasion. The connectors seem to be performing adequately and are expected to last the lifetime of the detector.

Wet End Breakdown typically causes events that are tagged by burst-like cuts.

6.1.5 Biases in Data cleaning

A concern is that the data cleaning cuts data in a way that depends on position or direction in the detector. Figs. 6.2-6.4 show the effect of the cuts as a function of radial position in the detector, of N_{Hit}, and of $\mathbf{u} \cdot \mathbf{r}$, where $\mathbf{u}$ is the reconstructed event direction, on ^{16}N data. The plots indicate that there is no significant bias. The inward bias in Fig. 6.4 is the largest, but is sufficiently small that it is not significant.

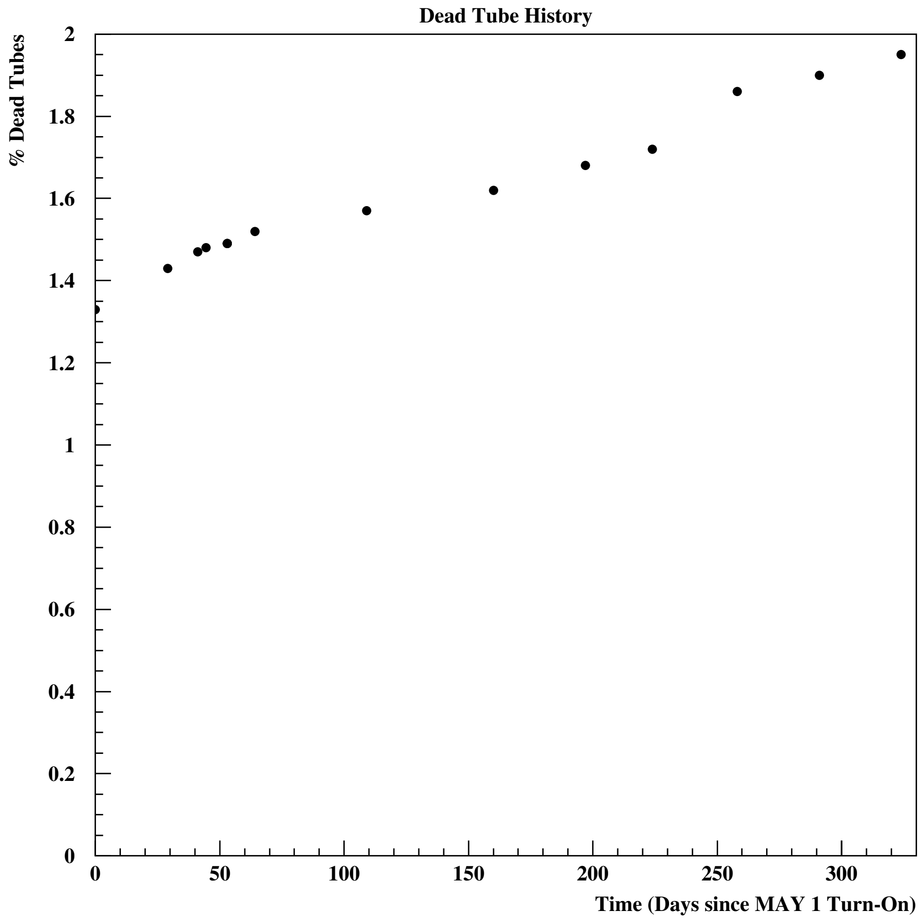

Figure 6.1: SNO PMT tube deaths, as a function of time. The approximate failure rate is 0.6% per year.

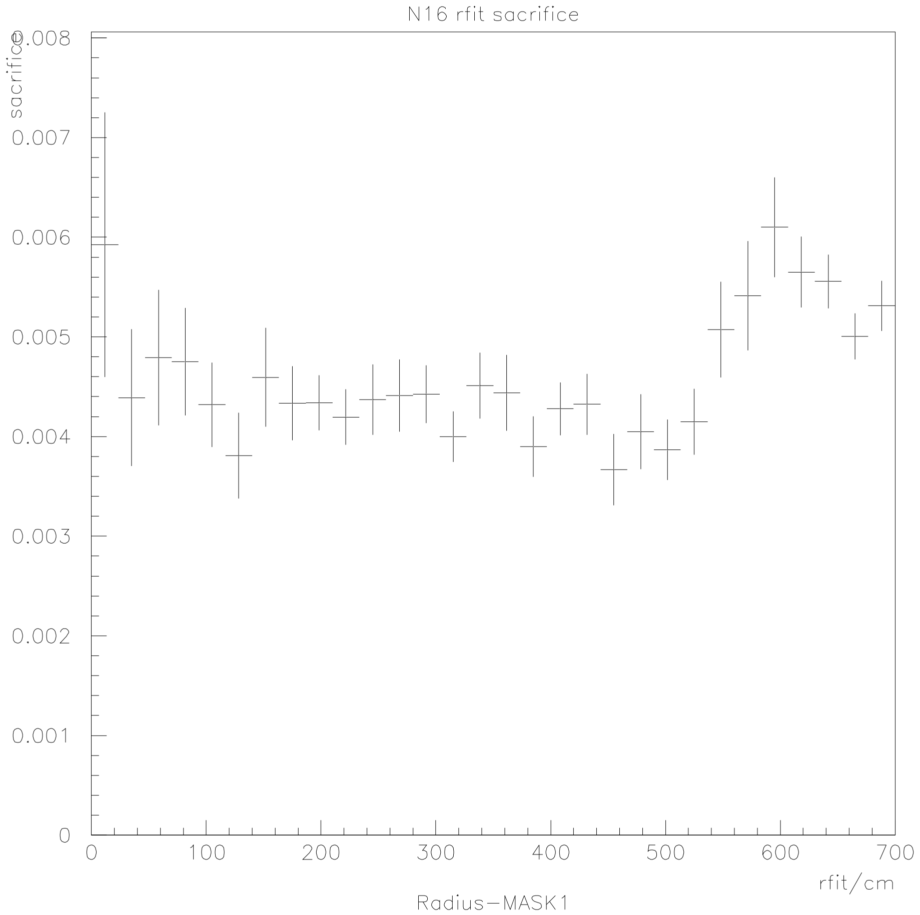

Figure 6.2: Bias in r, distance from detector center, for the data cleaning cuts.

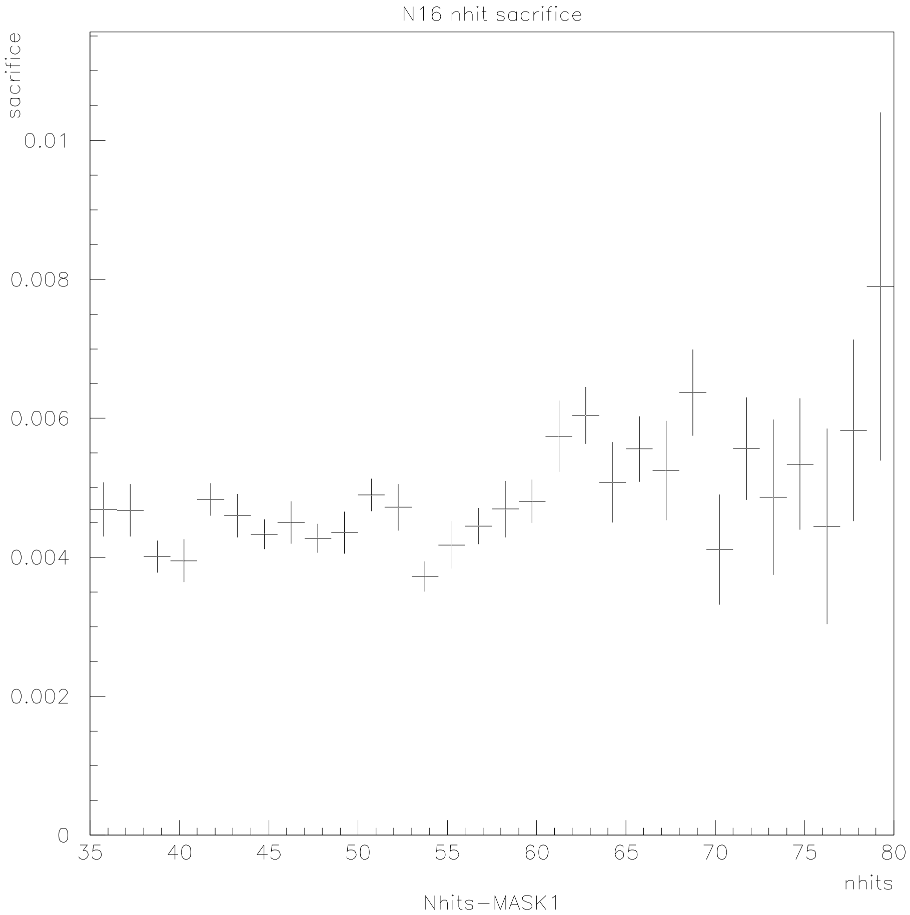

Figure 6.3: Bias with N_{Hit} for Data Cleaning Cuts.

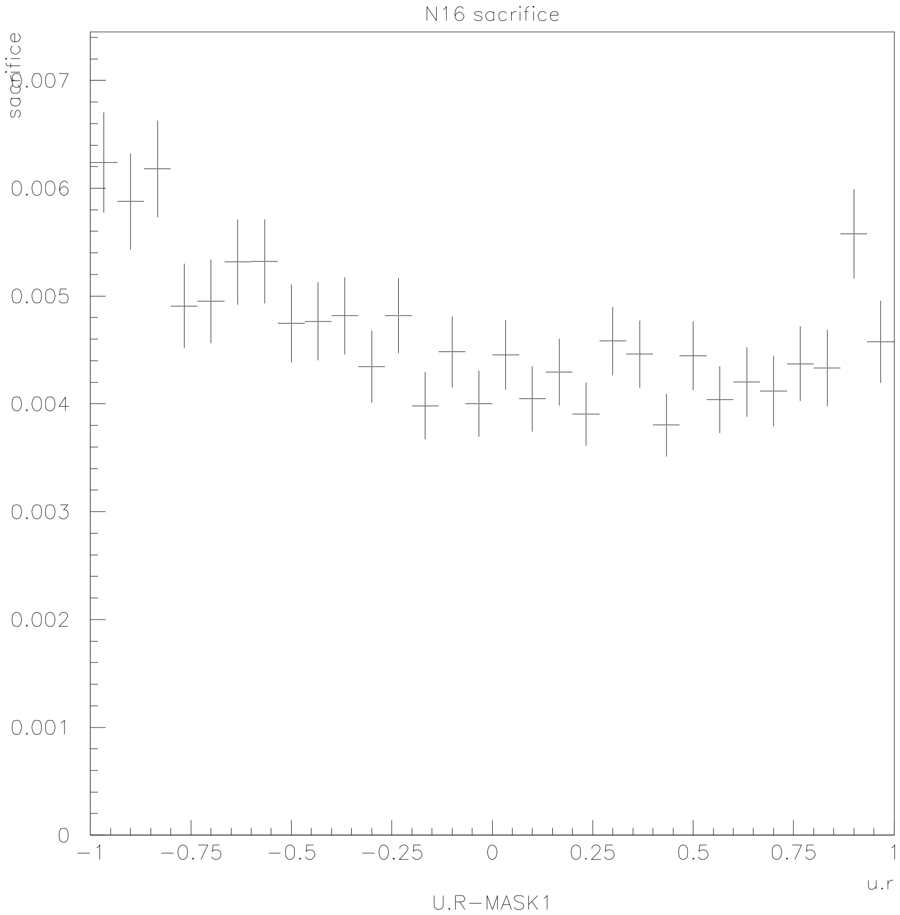

Figure 6.4: Inward-going bias in for Data Cleaning Cuts.

6.1.6 Default Mask and Data Reduction

The data cleaning group provides a default mask which specifies which cuts recommends should be applied to the data to produce a clean data set. The mask is listed below. We described most of the cuts in the previous sections. The remainder are described below.

- **Retrigger**: Events which are less than 5μs after another (but not the initial event).
- **Burst**: Bursts in which three or more events occur in 1 ms or less.
- **QvT**: Events in which the maximum QH (QL) charge is more than 1000 (180) counts away from the average and more than 60 ns earlier than the median channel time.
- **Q/NHIT**: Events in which the average charge is less than 0.25 photoelectrons, once the highest 10% of the charges are removed.
- **Crate Isotropy**: Events in which more than 70% of the hits are in one crate, and more than 80% of those are on two or fewer cards.
- **AMB**: Events in which the normalized Analog Measurement Board's Integral or Peak measurement of the ESUMHI signal are more than 3.7σ away from the average for the appropriate NHIT.
- **FTS**: Events in which the tube-to-tube pairwise time spread exceeds 6.8 ns, for tubes which are within 3 m of each other and which have time differences less than 25 ns.

- **OWL**: Events which have 5 or more OWL tubes firing.
- **Junk**: Events which are orphans or have multiple channel entries.
- **Neck**: Events in which both neck tubes are hit, or one is hit with large charge and an early time.
- **ESUM**: Events in which the ESUMHI bit is set but no others. This cut catches "shark fins," which are events where one tube fires with a large PMT pulse which has a shape similar to a shark fin.
- **Qcluster**: Events in which four or more channels in a sliding window five channels wide are hit, and at least one has a QH (QL) charge 2000 (300) counts above zero.
- **ITC**: Events in which the ratio of in-time hits in a 93 ns window to the total number of hits is less than 0.6. This cut tags pickup events.
- **FGC**: Events in which a cluster of hits in either channel or PSUP space is accompanied by hits on the PSUP whose average distance from the cluster is greater than 1200 cm. This cut tags flashers based on their geometry.
- **NHIT Burst**: Events for which six events above an NHIT of 40 occur in less than four seconds.
- **OWLEHI**: Events in which the OWLEHI trigger bit is set.

Based on this run selection mask, the total sacrifice on ^{16}N for all cuts is 0.5% [45].

Example of Data Reduction

As an example, we can look at what happens to the raw data stream when the cuts are applied. The cuts can be roughly grouped into removing the contamination we described above:

- flashers
- bursters
- neck events
- pickup

The data set we demonstrate the cuts on are one-third of the data used in this thesis. Each run is split into three equal length time slices, and for each run, a pseudo-random number is generated based on the run number to assign each slice to one of three data sets A,B, and C.[3] We apply the data cleaning cuts on Set A. Figs. 6.5–6.9 shows the N_{Hit} distribution for data Set A with the cuts applied in various stages. It is instructive to note that the last cut, requiring events to reconstruct within 540 cm of the center of the detector, eliminates the largest number of events.

[3]This data splitting technique addresses two goals.

- **Bias**: By tuning cuts on only one (albeit large) data subset and then doing the full analysis on the rest of the data while keeping cuts unchanged, the effect of bias and cut tuning is minimized.
- **Coverage**: By splitting each run in thirds, the analysis can monitor time-dependent detector behavior in the reduced data set.

This is one of the tools in SNO's initial attempt to address bias in its analysis.

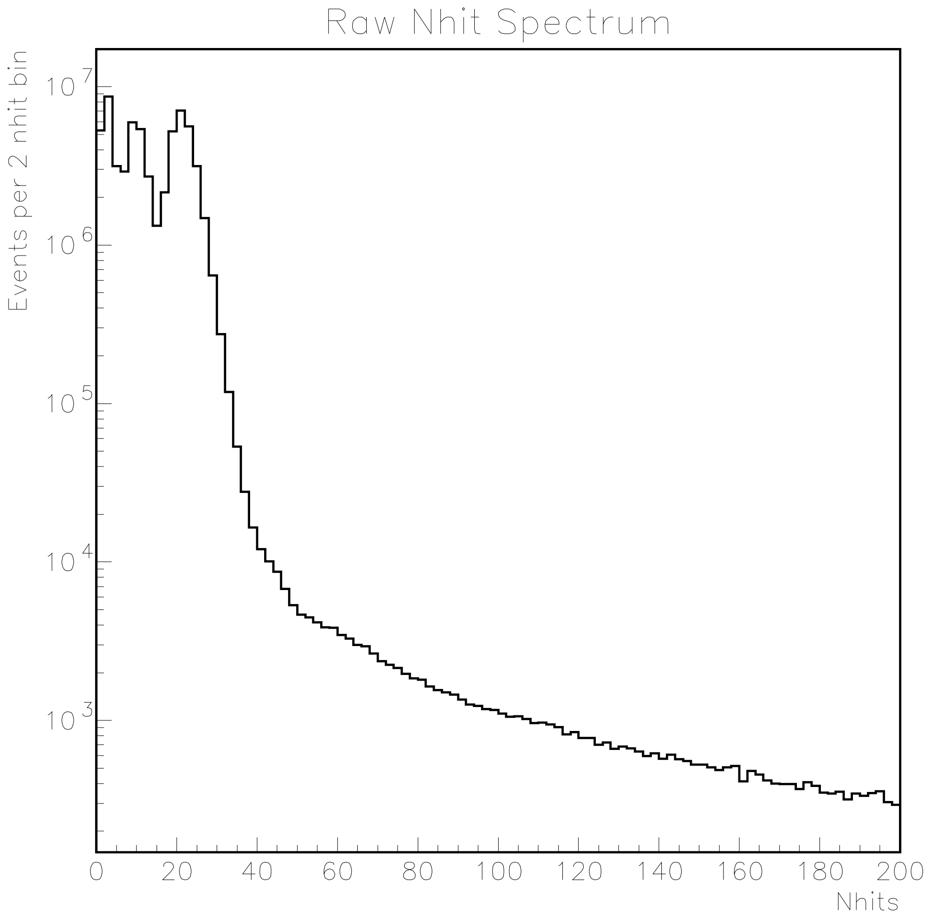

Figure 6.5: Raw N_{Hit} spectrum for all Set A data. The structure of the peaks near $N_{\mathrm{Hit}}=0$ are, from left to right, due to the PULSE_GT trigger (zero bias trigger which fires at 5 Hz), the ESUM energy sum trigger (mostly shark fins, see Sec. 6.1.6), and the `NHIT100` trigger threshold. Figures originally published in [46].

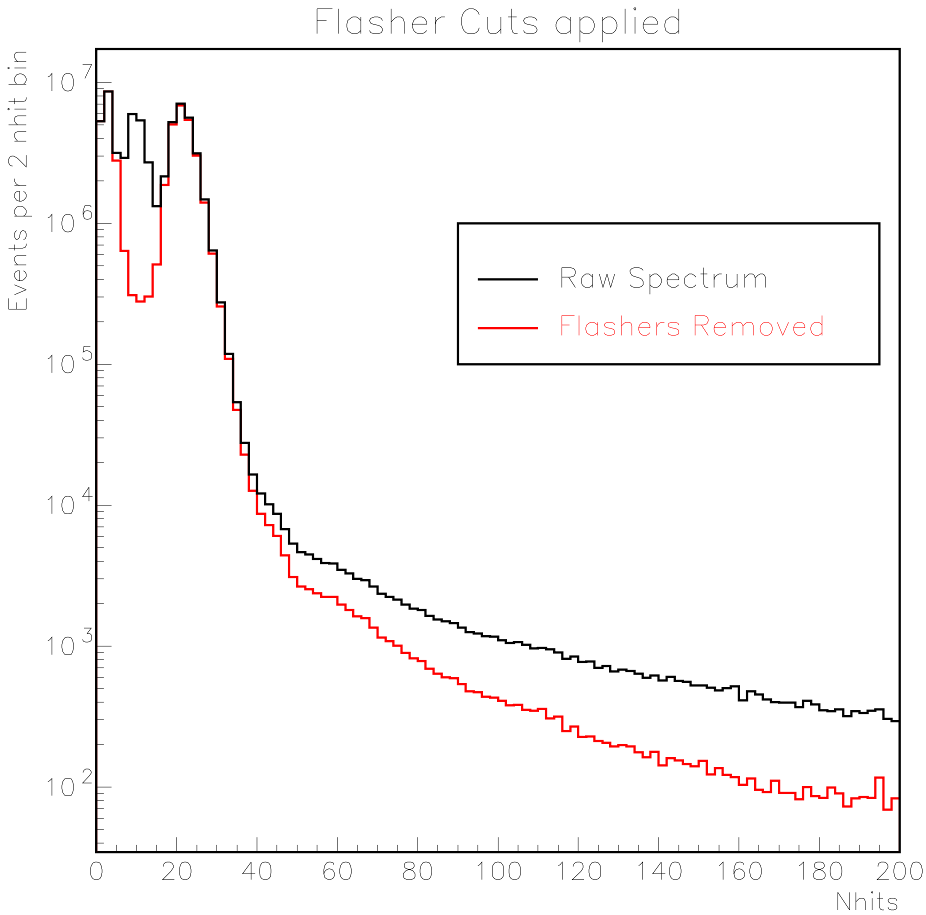

Figure 6.6: N_{Hit} spectrum for all Set A data with the Data Cleaning Group's flasher cuts applied.

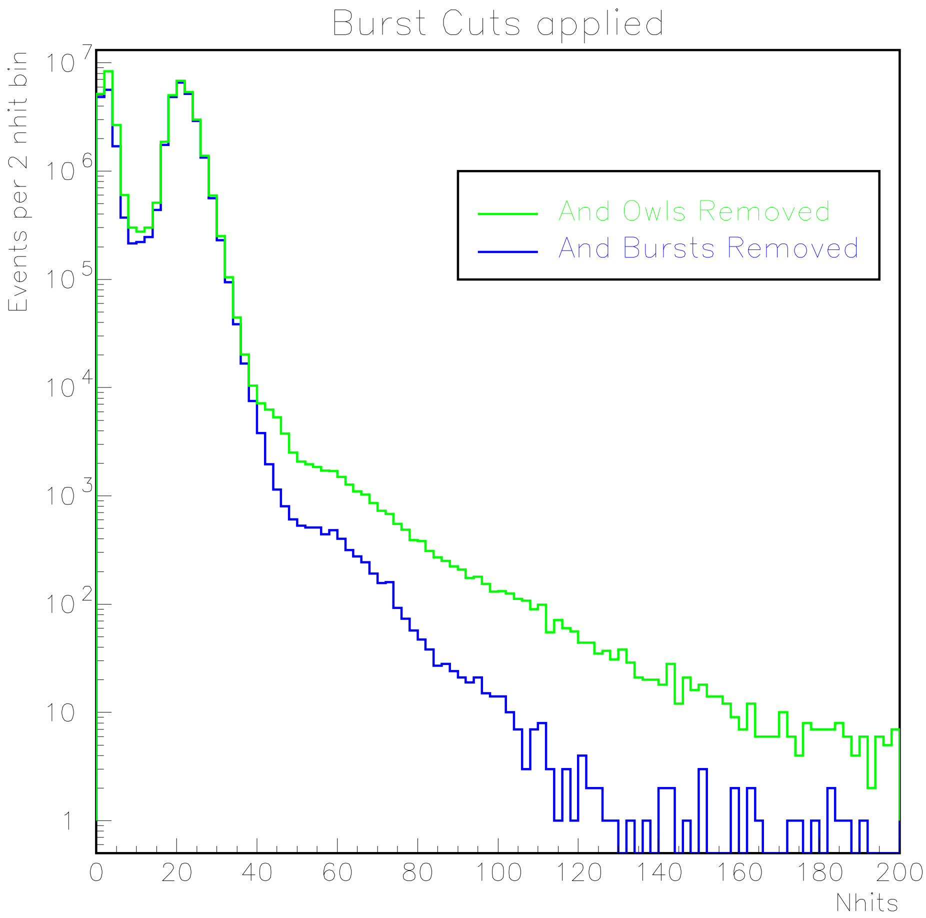

Figure 6.7: N_{Hit} spectrum for all Set A data, after flashers, bursts, and events with OWL hits have been removed.

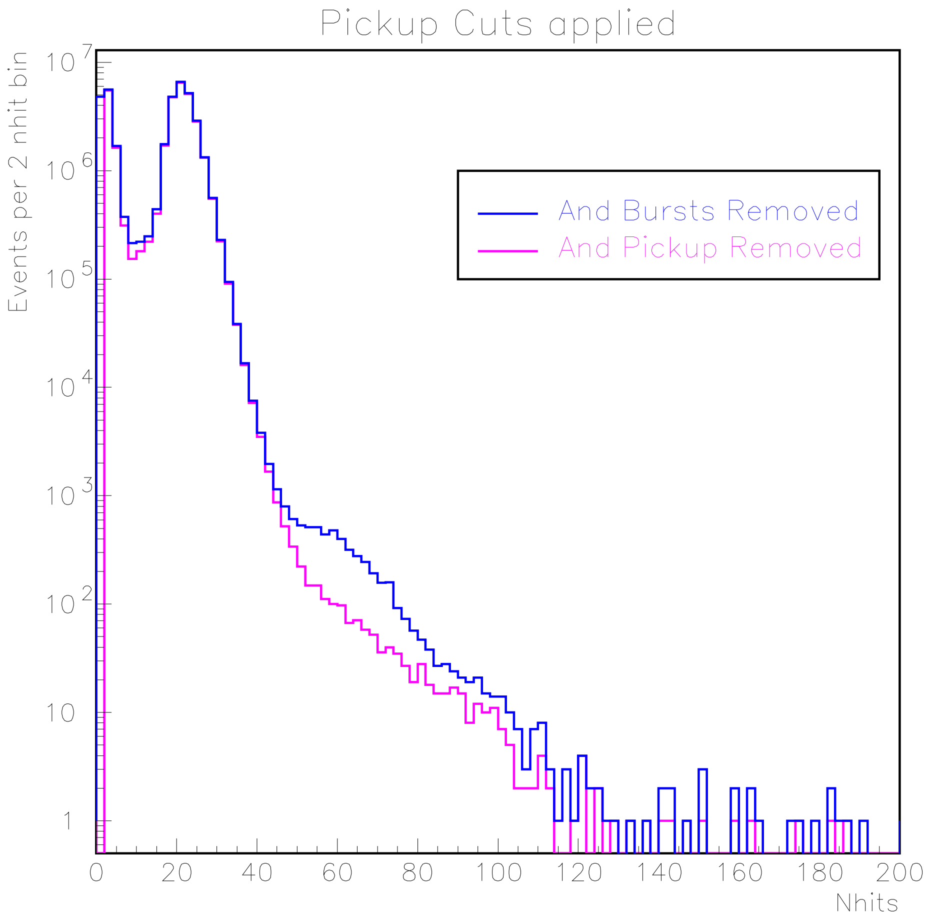

Figure 6.8: N_{Hit} spectrum as in Fig. 6.7, with pickup cuts applied.

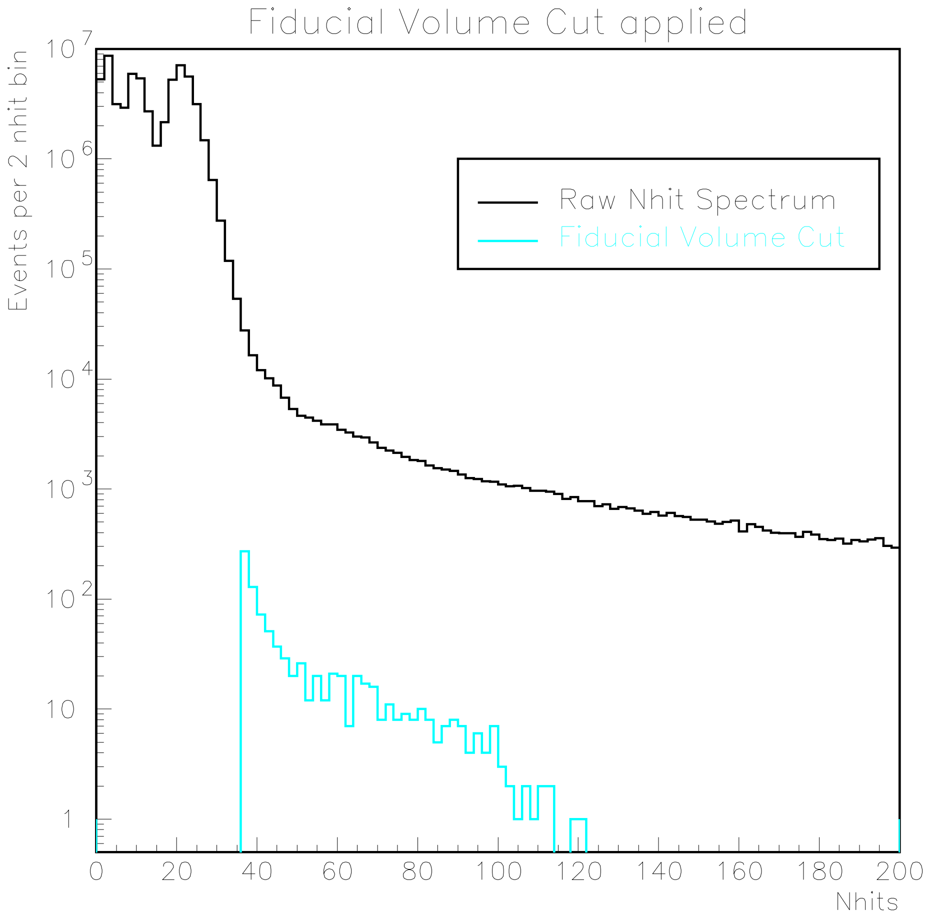

Figure 6.9: Raw N_{Hit} spectrum compared to cleaned spectrum, with 540 cm fiducial volume cut applied.

6.1.7 Instrumental Backgrounds Summary

In this section, we introduced the SNO data cleaning cuts. We showed how the cuts remove instrumental backgrounds, and how they do not remove signal-like ^{16}N events. The resultant data set is ready to be analyzed for physics backgrounds.

6.1.8 Physics Backgrounds

Physics backgrounds to this measurement include all radioactive decays in the detector, and the NC solar neutrino signal. The radioactive backgrounds include cosmic ray muons, PMT $\beta - \gamma$s, and high energy γ-rays from the rock wall, and internal radioactivity.

Muons

For solar neutrino analyses, muons are themselves not a concern. Due to the great depth of the detector, with an equivalent of 6000 m of water overhead, the muon rate is very low at roughly 3/hour. Most of these muons are through-going, and light from them hits more than 500 phototubes. Since the highest energy of the solar neutrino flux, the *hep* neutrinos, end at 18 MeV , which corresponds to less than 200 hit phototubes, the overlap region between the muons and the solar neutrinos is very small. However, muons that stop in the detector and spallation products produced by muons in the detector are a possible concern.

Stopping muons can have sufficiently low energies that they fall into the same N_{Hit} range as the solar neutrino signal. These muons decay in the detector. The decay electron is tagged via a burst cut.

Muons can produce spallation products, which may look like solar neutrinos. Spallations of n, ^{8}Li, ^{16}N, ^{8}B, ^{12}B and ^{12}N can be produced and in turn β-decay. The endpoints of these decays is well within the range of solar neutrinos (13-16 MeV), and most of the half-lives are in the range of 10-800 ms, while some are as long as 70 sec [47].

In order to remove this possible background, events that follow muons are tagged if they occur within two different time windows after the event. A short time window is used to catch most of the spallation products, and a longer decay window is used in order to mark the very long-lived spallation products. The former time limit is 20 s, while the latter is 60 s.

The muon cut is claimed to be 99.5 efficient at detecting muons, though the efficiency for stopping muons is roughly only 90% [48, 49].

High Energy γ-rays

Interactions of thermal neutrons in the PSUP steel and the cavity rock can produce γ-rays up to 10 MeV in energy [50]. These γ-rays can reach the heavy water, and react there, mimicking the solar neutrino signal. Work has been done to estimate an upper limit of this background in the acrylic vessel [46, Sec. 5.2]. In brief, in this analysis an upper limit on leaking γ-rays is made based on calibration data using the ^{16}N data at the PSUP. Fig. 6.10 gives a brief description of the analysis. With this method, the collaboration obtains an *upper estimate* on the contamination if one assumes that all radially inward events are high energy γ-ray background. A preliminary estimate of the resultant background rate for the FTE(G) grid-elastic fitter is 0.11 ± 0.04 events per day above N_{Hit} =50 [46].

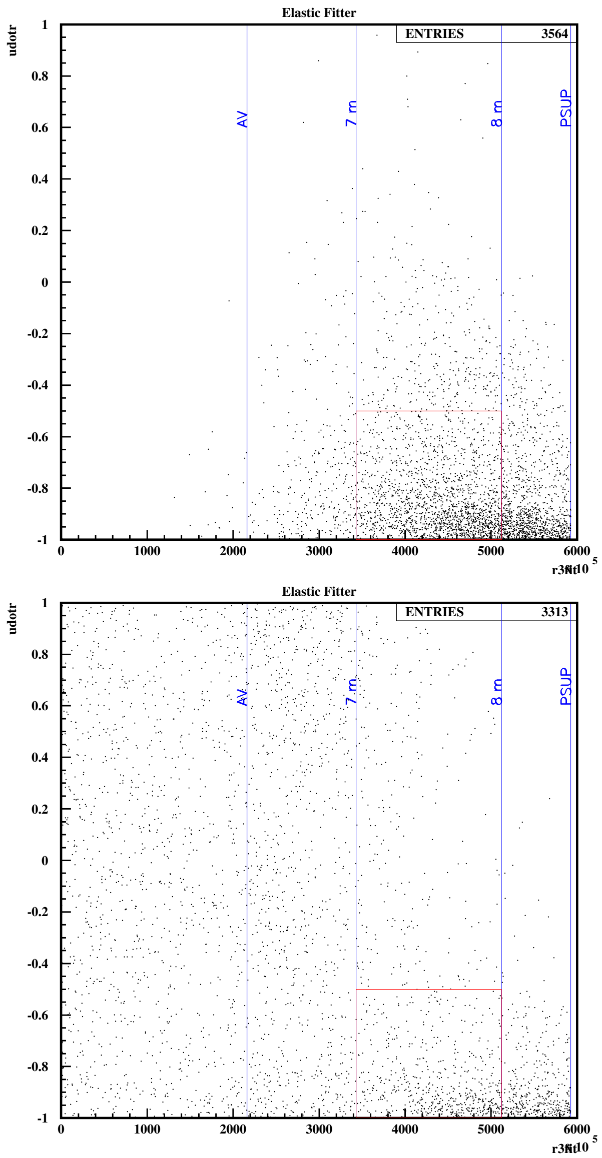

Figure 6.10: High energy γ-ray estimation using the ^{16}N source, $\hat{d} \cdot \hat{r}$ vs r^3, where $\vec{r}$ is the vector that points from the center of the detector to the reconstructed position and $\hat{d}$ is the reconstructed direction. The upper plot shows ^{16}N calibration data. Here, we measure the ratio $\mathcal{R}$ of the number of events which reconstruct in the box in the light water near the source *vs* the number that leak into the heavy water. In the lower plot, by counting the number of events in the box, we can use the ratio $\mathcal{R}$ to estimate the amount of high energy γ-ray leakage. Figures courtesy of M. S. Neubauer [51].

Solar ν signals from neutral currents and neutrons

The next two backgrounds are solar neutrino physics signals in their own right, but as we are not measuring them in this analysis, they are a background. The neutral current signal discussed in the introduction and in Sec. 3.4.2 falls in this category for a *CC* flux measurement. In addition, neutrons may be induced by radioactive decay processes. In the pure D_2O phase of the experiment, the *NC* signal consists of a 6.25 MeV γ-ray in the detector. A few comments should be made about this contamination.

- The analysis threshold for this measurement was chosen to be above the *NC* signal. The signal itself is not yet well understood, nor the backgrounds associated with it.
- Unlike the *CC* and *ES* signals, the angular distribution of the *NC* γ-ray is uncorrelated to the solar direction. This means that we cannot use the angular information of the event to measure the *NC* contamination.

We can study what neutrons look like in the SNO detector with the ^{16}N source, as it produces a very similar 6.1 MeV γ-ray. With an analysis threshold of $N_T > 60$, we expect from ^{16}N data to still have roughly 6% NC signal contamination. Going to a higher threshold would eliminate this contamination; at $N_T > 70$, we expect less than 1% neutron contamination. However, the CC signal is also rapidly fading at higher thresholds, so in the interest of sensitivity, we will accept and correct this contamination.

Solar neutrino elastic scattering interactions

The elastic scattering signal is a background for the CC flux measurement. For an individual candidate neutrino-induced electron, we cannot tell whether it was created in an ES or CC reaction. As discussed in Sec 3.4.1, the electrons that are produced in elastic scattering interactions have a direction that is strongly correlated with the direction of the incident neutrino, and hence with the instantaneous solar direction. The strength of the angular correlation makes these neutrinos easy to "see," despite the fact that there are predicted to be fewer of them than of the CC neutrinos by a factor of roughly ten. The angular distribution with respect to the sun ($\cos\theta_{\odot}$) will be our only handle for separating CC and ES events.

6.2 Energy Scale and Resolution near threshold

In order to measure the CC rate, an approximate energy scaling must be understood near the threshold of the measurement. The energy calibration is complicated, as the response of the detector is a function of position and direction in the detector, and the optical properties must be well understood to model it successfully. As a full energy calibration for SNO is still in progress, and such a calibration involves more complexity than is needed in this analysis, I will use a simple approximate method.

The total number of electrons can be described like this:

$$\frac{dN}{dE_e} = N_0 \int \frac{d\phi^{\nu}(E)}{dE_{\nu}} \frac{d\sigma(E_{\nu})}{dE_e} \, dE_{\nu} \tag{6.1}$$

The total number of detected events with some energy threshold is

$$\begin{aligned} N_{\text{total}} &= \int_{E_T}^{\infty} \frac{dN}{dE_m}\, dE_m \\ &= \int_{E_T}^{\infty} \int_{0}^{\infty} \frac{dN}{dE_e} R(E_e - E_m)\, dE_e dE_m, \end{aligned} \tag{6.2}$$

where E_m is the measured energy, E_T is the threshold energy and R is the usual resolution function. For each event, we assign an energy based on the discussion below on the correlation between N_{Hit} and energy. The resolution function is Gaussian.

$$R(E_e - E_m) = \frac{1}{\sqrt{2\pi}\sigma_{E_m}} e^{-\frac{(E_e - E_m)^2}{2\sigma_{E_m}^2}} \tag{6.3}$$

Note that due to the integration over detected and emitted energies, the resolution only matters near the energy threshold E_T. However, note that the width of the resolution is quite large, and the Monte Carlo predicts a considerable spread in values from mono-energetic electron, so the effect of the width is non-zero only far from the energy threshold.

6.2.1 Energy calibration

In order to apply Eq. 6.2, an energy calibration must exist to determine E_m for each event. The number of Čerenkov photons is related to the track length and to the energy of the produced photon [52, 53].

$$\frac{dN}{dx} = 2\pi\alpha z^2 \int_{\lambda_1}^{\lambda_2} \left(1 - \frac{1}{n^2\beta^2}\right) \frac{d\lambda}{\lambda^2} \tag{6.4}$$

For electrons, the dependence of the number of Čerenkov photons emitted per unit

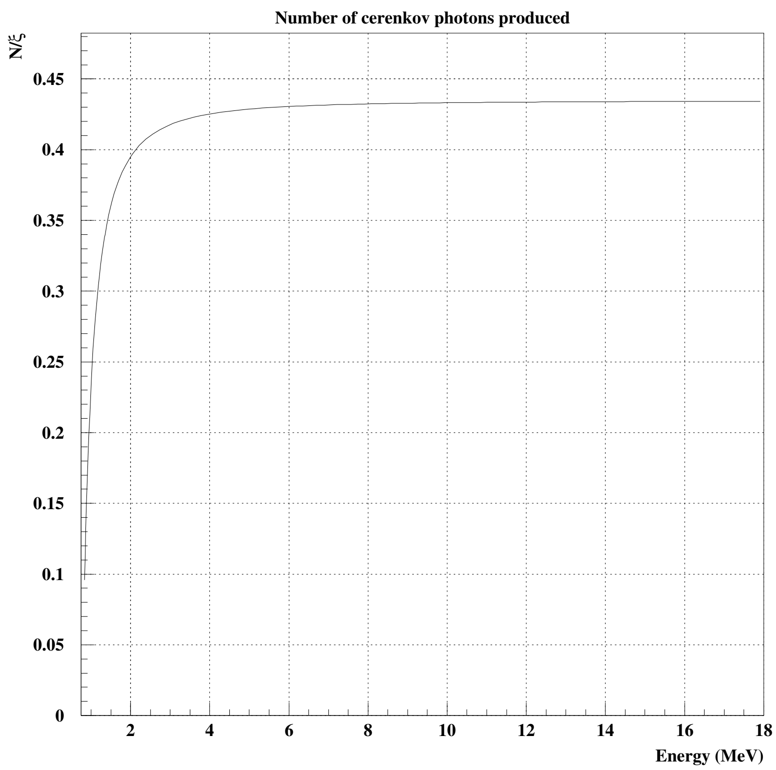

Figure 6.11: Number of Čerenkov photons emitted per unit length, as a function of energy.

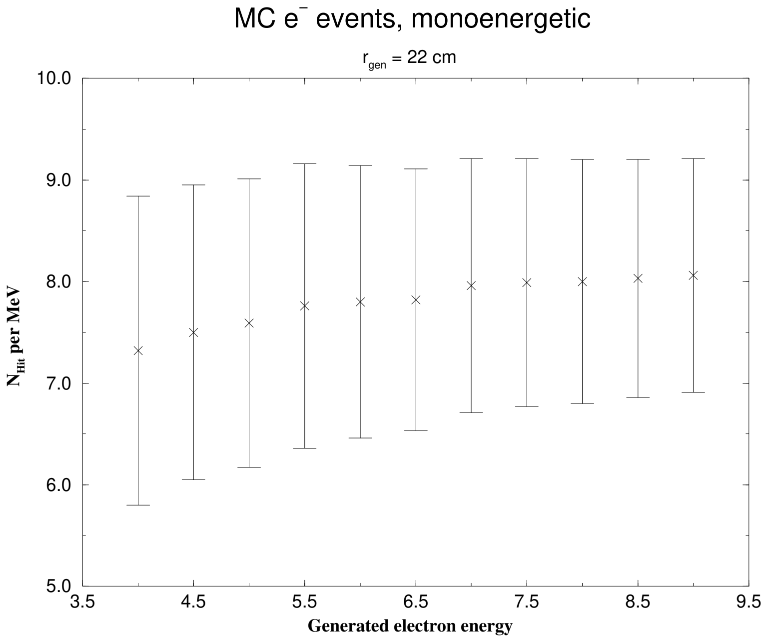

Figure 6.12: Mean number of hit photo-tubes for a simulated electron, using the SNOMAN Monte Carlo package. The particle was generated isotropically near the center at $\vec{r} = (10\text{cm}, 10\text{cm}, 10\text{cm})$.

track length on the electron energy quickly levels off at $E_e \simeq 2\text{MeV}$, as shown in Fig 6.11, and hence above those energies, the amount of light emitted by an electron is approximately linearly related to its energy. For instance, Fig. 6.12 shows a simulation of the detector response to mono-energetic electrons that are isotropic in direction at one position near the center of the detector. While the response at low energies shows a dependence on the energy, at higher energies, the number of tubes per unit energy approaches a constant value. The energy of an electron is related to the number of photo-tubes N_{Hit} that fire in a given event. I have defined a scale factor S which described this linear relationship.

$$dE = S\,dN$$

$$S\left(\vec{r},\vec{d}\right) = \frac{\mu_{N_{\text{Hits}}}}{\mu_E} \tag{6.5}$$

where $\vec{r}$ is the source position, $\vec{d}$ is the electron direction, and $\mu_{N_{\text{Hits}}}$ is the mean number of hits for a set of events whose mean energy is μ_E. By deploying a source of electrons of a known energy and direction at many points in the detector, one can map out the scaling S for SNO at one energy point and at many positions. However, as mono-energetic triggered electrons are not trivial to come by, SNO uses mono-energetic γ-ray sources for this purpose. These γ-rays Compton-scatter, transferring most of their energy to an electron in the first scatter, and provide SNO with electrons. Two main sources, which emit a 6.1 MeV and a 19.1 MeV γ, respectively, have been developed for SNO and are described elsewhere.[4] The 6.1 MeV source has been deployed.

Using the Monte Carlo to produce $e-\gamma$ correlations

The final sticking point that remains in making a rudimentary energy calibration that remains is how to relate the γ energy to the electron energy. The particles of interest in SNO, as far as detection is concerned, are electrons, while the calibration source emits γ-rays. The primary γ produces the usual collection of charged particles via Compton scattering and pair production, that in turn produce Čerenkov photons of varying energies. Once those particles dip below the threshold for producing Čerenkov radiation, however, their remaining kinetic energy is "lost", i.e., undetected. The spread in the number of particles that are produced will affect the total energy that is deposited in the detector for a given γ-ray.

This topic is well suited for a Monte Carlo study. At the energies we are interested

[4] See, for instance, [54, 55].

in for this calibration source ($E_\gamma \simeq 6$ MeV), the physics of Čerenkov radiation and electron energy loss predicts that the energy of the particle should be related to the number of Čerenkov photons produced. Therefore, one can consider using the Monte Carlo to calculate how many Čerenkov photons are produced by a mono-energetic γ-ray, and generate electrons, varying the electron energy until the mean number of Čerenkov photons produced agrees with the above.

$$\text{If } N_{e_c} = N_{\gamma_c}, \text{ then } E_e \simeq \sum_{i, E_i > E_c} E_i.$$

where N_{e_c}, N_{γ_c} is the number of Čerenkov photons produced by an electron and a γ, respectively, and E_i is the kinetic energy of the ith particle and the sum is over all energy deposited by charged particles over the Čerenkov threshold energy E_c.

The number of tubes hit in the detector is directly proportional to the number of Čerenkov photons produced in the medium by the passage of a charged particle. While the SNO Monte Carlo has not been fully calibrated, the ratio of the number of predicted counts for two types of events cancels errors in photon propagation that would enhance or suppress the total number of hit tubes in both cases. Therefore, looking at the Monte Carlo prediction for

$$S_{e-\gamma}^{MC} = \frac{N_{e^-}^{MC}}{N_{\gamma}^{MC}} \tag{6.6}$$

gives an accurate estimate of the ratio of the number of Čerenkov photons produced by γ-rays as compared to electrons for primary particles of the same energy. Given

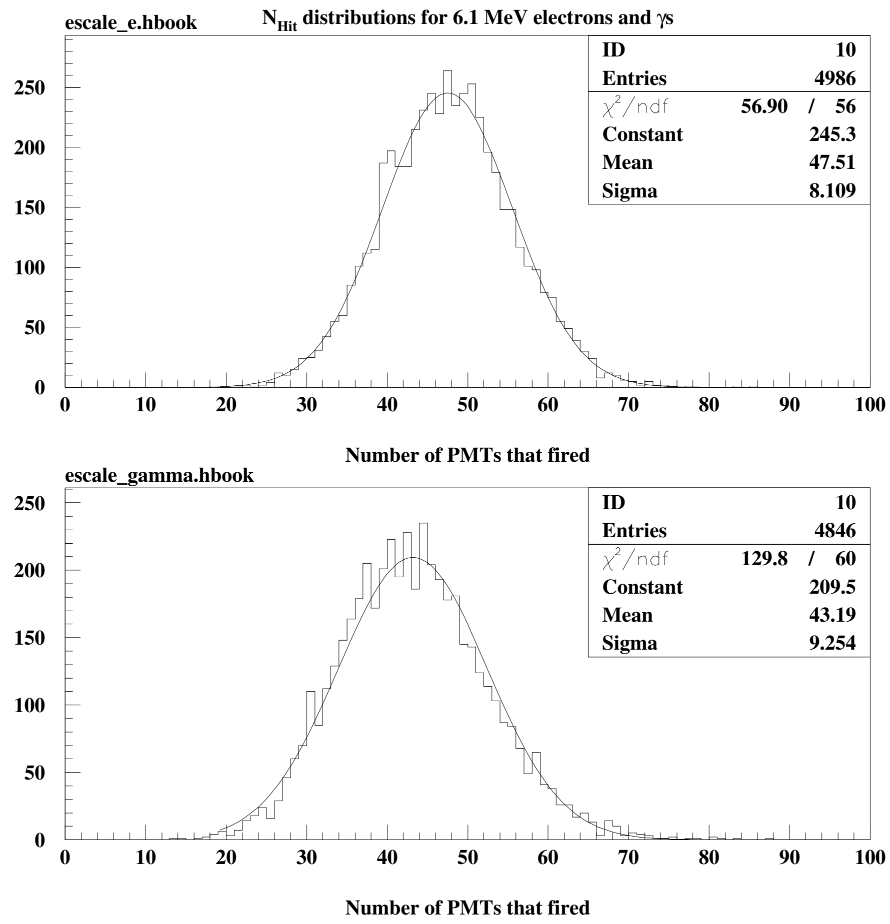

Figure 6.13: Number of hit photo-tubes for 6.1 MeV γ-rays and electrons, as predicted by SNOMAN.

the means predicted in Fig. 6.13,

$$S_{e-\gamma}^{MC} = 1.100 \pm 0.004$$

which leads to a conclusion of an effective total electron energy, that is, the energy at which the electron produces as many Čerenkov photons as a 6.1 MeV γ-ray.

$$E_e^{\text{effective}} = 5.54 \pm 0.02\,\text{MeV}\,. \tag{6.7}$$

This energy can be used to understand calibration data from the ^{16}N source and determine the scale factor S described in equation (6.5).

Alternatively, one can vary the electron energy until the mean of the number of hit PMTs of the two N_{Hit} distributions match. With this procedure, which one would expect to come to a similar result, the means match to 0.5% when $E_e = 5.6\text{MeV}\,$.

The calibration data sets

The ^{16}N used for this calibration is part of a general system of producing short-lived radio-isotopes for γ, β, and n calibrations [55]. It consists of a commercial *DT* generator, where fast neutrons are produced via the follow reaction.

$$d + t \rightarrow n + {}^{4}\text{He} \tag{6.8}$$

The neutrons have an energy up to 14.2 MeV. The neutrons then bombard CO_2, and react with the ^{16}O target to produce ^{16}N.

$$n + {}^{16}\text{O} \rightarrow p + {}^{16}\text{N} \tag{6.9}$$

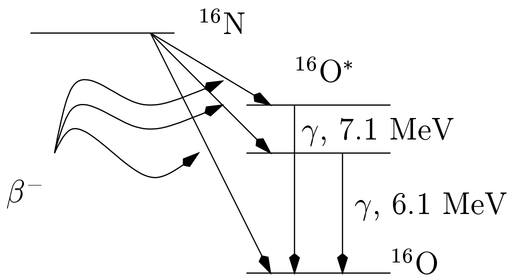

Figure 6.14: Decays of ^{16}N. See [19]. Only the two $\beta\gamma$ decays are used in SNO for calibration.

As the DT generator is far from the detector, and the half-life of ^{16}N is short ($T_{1/2} =$ 7.13 sec), a gas transport system was built to rapidly transfer the radio-isotopes to the decay chamber inside SNO's fiducial volume.[5]

The ^{16}N source is a triggered source. As shown in Fig. 6.14, the nuclear decay of ^{16}N can occur in several ways. The main decay emits γ-rays [19]:

$$^{16}\mathrm{N} \rightarrow {}^{16}\mathrm{O}^* + \beta^- + \nu_e \tag{6.10}$$

$$\hookrightarrow {}^{16}\mathrm{O} + \gamma \tag{6.11}$$

The γ-ray has an energy of 6.13 MeV (7.1 MeV) with a branching fraction of 68% (4.9%). Additionally, the β-decay can occur directly to the ground state (branching fraction of 26 %) of ^{16}O, in which case no γ-ray is emitted in the detector.

The β-ray is captured in a scintillator shell which surrounds the decay volume. A photo-tube detects the light emitted in the scintillator. The signal from this photo-tube is used as a trigger for the SNO PMT array. Additionally, the signal from this calibration photo-tube is fed into a designated channel of the SNO electronics and read out and stored in the same manner as a normal SNO PMT is handled. This

[5]In addition to ^{16}N, this calibration system is designed to produce a triggered ^{17}N-neutron source ($T_{1/2} = 4.4$ sec) and a triggered ^{8}Li β-source, where the β endpoint is roughly 13 MeV.

allows a charge and time measurement of the calibration PMT. By placing a cut on the charge measured in this electronics channel, dark noise and scintillator radioactivity contributions to the trigger rate of this tube can be effectively removed.

Several sets of scans with the ^{16}N to determine dependence on $\vec{r}$, were investigated. The data was taken in the SNO detector during the first period of production data taking. The calibration source was deployed in four sets of positions: the center, in two scans off center in the $x - z$ and $y - z$ planes, along the x axis, and in the light water between the acrylic vessel and the PSUP. In all, roughly 60 different source positions were sampled, at various radii.[6]

For each of these source positions, we selected a subset of the data. During the ^{16}N runs, the standard N_{Hit} trigger is also enabled. To find only the ^{16}N decays, the following cuts are applied:

- ^{16}N PMT channel is present in the event.
- charge in the PMT channel is $Q_{HS} > 200$ ADC counts above pedestal.
- $N_{\text{Hit}} > 20$.
- fit distance from source position: $r_{\text{fit}} < 80$ cm

For the events satisfying these criteria, the N_{Hit} distribution is collected and fit to a Gaussian distribution. The mean of this distribution is comparable to the scale factor described in Eq. 6.5. In this manner, a one point energy scale measurement is simply done at many points in the detector. The distribution of these means can be seen in figure 6.15. In the acrylic vessel ($r < 6$ m), the means vary linearly from the low point in the center.

[6]For a description of the calibration system, including the source positioning system, see [21].

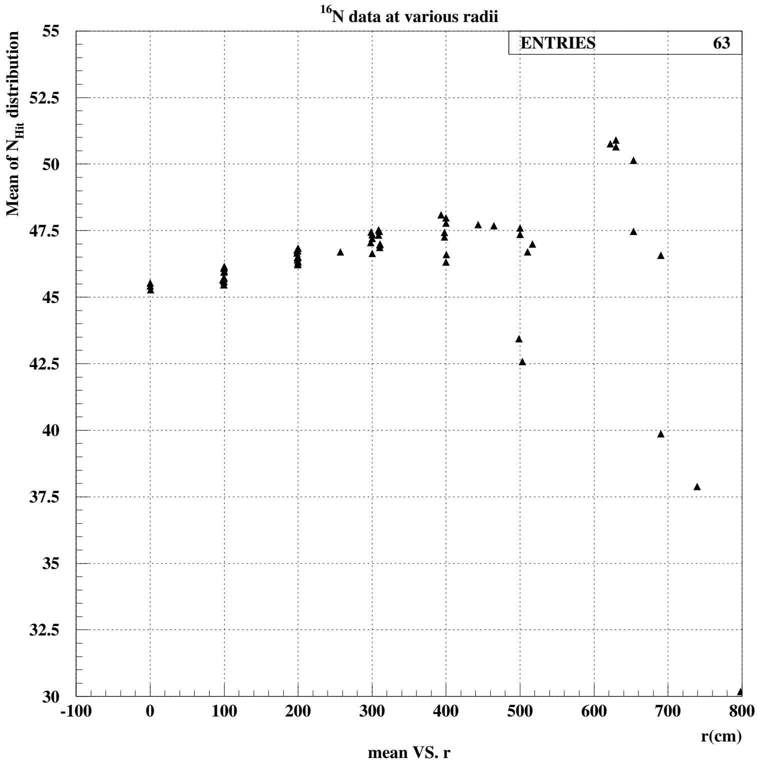

Figure 6.15: Mean response of SNO to ^{16}N γ-rays as a function of radial position. The overlapping points at the same radii are in different positions in the detector. Note the dramatic cutoff near the radius of the acrylic vessel and beyond ($r \geq 5m$). The wide spread at five meters is due to the neck region the AV, where tubes are missing and where the acrylic used is opaque. Near 7 m, the photo-tube concentrators start cutting off (see text).

In the area outside the AV, the Mean of the N_{Hit} distribution falls off sharply. The wide spread at five meters is due to the neck region the AV, where tubes are missing and where the acrylic used is opaque. At seven meters, the PMT concentrators cut off (see Sec. 4.3 for a discussion of the photo tubes and concentrators.)

Event energy determination

To determine the energy of an event, we apply the closest scale factor to the reconstructed event. If we are between points, we interpolate between the scale factors to determine the local scale factor.

We abandoned this method for this analysis as inconsistencies appeared in the data. Seemingly identical ^{16}N runs at the same location gave answers that varied by 3-5%, which is a similar number to the total radial variation that we were trying to correct. While the collaboration investigates this problem, we are forced to continue with a somewhat less ambitious plan. We now use one global scale factor from ^{16}N at the center of the detector, and accept a 3-5% systematic uncertainty on our energy threshold. The scale factor, as determined from run 10490, is

$$\overline{S} = 9.04 \pm 0.45 \frac{N_{\mathrm{Hit}}}{\mathrm{MeV}} \tag{6.12}$$

6.3 Event Reconstruction

Event reconstruction produces event times and directions from the calibrated PMT hits.

6.3.1 Position resolution

This thesis uses a combination of two fitters to reconstruct event positions and directions. Both fitters mainly use the differences in arrival times, coupled with the PMT positions, to determine where the event took place.

This list describes all information which makes up an individual event that one might consider using during fitting.

- PMT hit positions
  - relative to reconstructed position
  - relative to reconstructed direction
- PMT hit times
- PMT hit charges, i.e., estimate of the number of photons that hit the tube.

For each fit point, we compare these quantities to our expectation under our current fit hypothesis. The fit hypothesis usually assumes that the candidate event was caused by an electron. Simplifying assumptions such as no reflections and no scattering are sometimes made. We vary our fit position and direction until the quantities match up with the prediction for the current location.

Fitters used in this thesis

SNOMAN comes with several standard fitters. This thesis uses two of these fitters. We will compare our fitters with a third fitter in several locations. The third fitter is the Path Fitter (FTP), which is under development by Mark Neubauer

SNO fitters typically minimize time residuals to determine the fit position. For instance, for a simple time-based fitter, one uses a χ^2-type minimization function which minimizes time residuals under the hypothesis that every PMT is hit by a photon which travels linearly from the interaction point to the PMT.

$$\chi^2(t_f, \vec{r}_f) = \sum_i^N \left[\frac{\left(t_f - t_i - \frac{|\vec{r}_f - \vec{r}_i|}{c'} \right)}{\sigma_f} \right]^2 \tag{6.13}$$

$\vec{r}_i$ is the vector which points to the ith PMT, t_f and r_f are the current hypothesized fit position, and c' is the local velocity of light.

The main fitter we used is the elastic fitter FTE [56]. It uses both time and angular information to determine the vertex of an event, with the hypothesis that all PMT's are hit by direct Čerenkov light. The algorithm uses a weighting factor to decrease the effect of out-of-time tubes on the fit. The fitter determines the both the position and direction of an event, where the two fits are independent but the weighting for the direction fit uses information from the time fit.

The second fitter that is used is the SNOMAN grid fitter FTG [57]. The grid fitter uses a multi-step process to determine the fit position of an event. The basic algorithm is a maximum likelihood method using time of the hit tubes. In order to find an initial position, a 3-D grid of points is set up and at each point, a 1-D normalization with respect to time is performed. The fitter performs a full 4-D normalization for the actual fit vertex and time using the grid point with the best likelihood as the initial position. The grid fitting forces the fitter to explore a large parameter space to escape local minima. Angular information of the hit tubes is not used, and no explicit fit for direction is made.

We use these fitters in tandem. First, the Grid fitter determines the fit location. As the grid fitter is designed to be robust against local minima, we use it to determine a good starting location for the elastic fitter and seed the elastic fitter's initial position with the grid fitter best point for a final fit for position and direction. The elastic fitter has good directional reconstruction, so we use it to determine the final position and direction of the event. This particular combination, the FTE(G) fitter, seems to effectively leverage the strengths of each of the fitters, and is a novel combination of these two standard fitters.

MC estimate of fitter performance

In the future, the fitter performance will be determined using calibration data of electron sources. At this point, no such data or source exists. In the interim and for this thesis, we will make do using the Monte Carlo to estimate the fitter's performance. As the Monte Carlo is also not calibrated, this is far from an ideal method. Some of the problems with this method are listed below.

- **Optics**: Optical parameters in the Monte Carlo are based on estimates and bench-top measurements. In fact, the current MC optical parameters significantly disagree with the calibration group's current best estimates of these parameters. The effects of the optical model are probably strongest in the regions such as the neck or in the light water near the PSUP, and will affect the time distributions that all fitters use.

- **Čerenkov light generation**: Many of the more elaborate fitters, such as the elastic fitter, try to use the characteristic light patterns generated by Čerenkov radiation to improve the fits. If the Monte Carlo generates incorrect patterns, i.e., if its photon propagation model is incorrect, the fitters will be tuned to incorrect patterns.

- **Completeness**: The Monte Carlo will never simulate all aspects of the detector, and hence a direct measurement is always preferable to a simulation where possible.

Nonetheless, at this point, we will use the MC to estimate the position reconstruction resolution. For this purpose, we use the electrons at the threshold energy (E_T = 7 MeV) at three positions in the detector along the z-axis: $z = -200, 0, 200$ cm, and look at the distance between the reconstructed point and the generated point. Fig. 6.16 shows the ability of the fitter to reconstruct electrons. The reconstruction resolution is

$$R(r - r') = r^2 \frac{1}{\sqrt{2\pi}\sigma_r} e^{-\frac{(r-r')^2}{2\sigma_r^2}} \tag{6.14}$$

where $\sigma_r = 25\ cm$. For a rate measurement, we must know how many targets are present. The position resolution allows us to make an estimate on the error in the volume, as

$$\frac{\Delta V}{V} = 3\frac{\Delta r}{r}. \tag{6.15}$$

except in cases where we are near the edge of the physical medium.

We can make one check to assure ourselves that the MC and the data are consistent by looking at the resolution we obtain from the one Čerenkov light source that has been deployed in SNO, the ^{16}N source. While the 35 cm Compton scattering length of the γ-ray makes it hard to use this for position resolution directly, it can be useful for a one-point data-MC cross check. Fig. 6.17 shows such a comparison. We see that the MC and the data are qualitatively similar, and both measure similar Compton lengths.

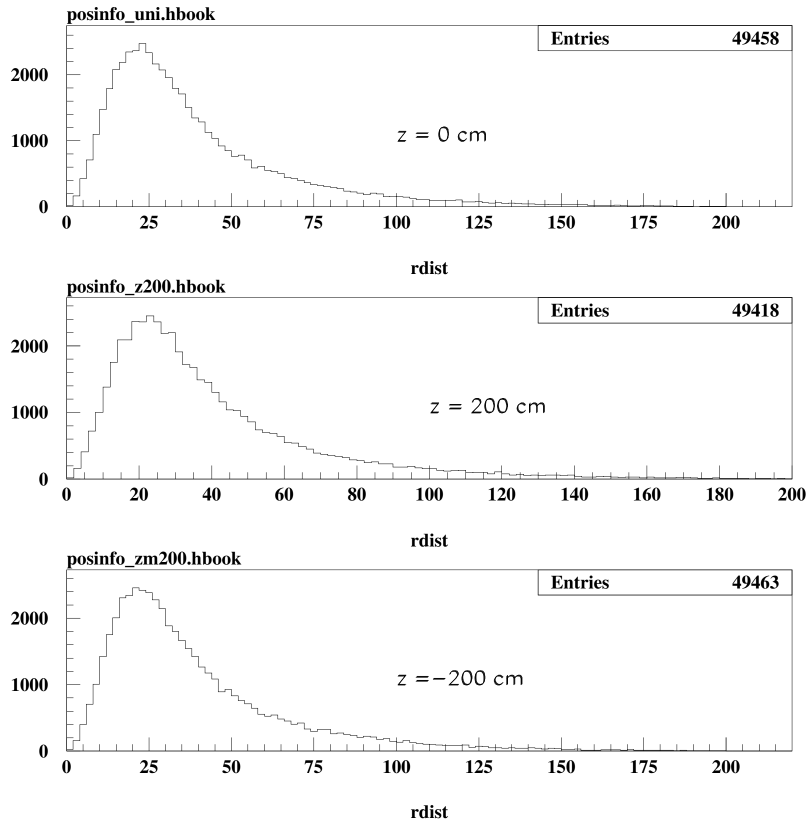

Figure 6.16: Position resolution of the FTE(G) fitter, for mono-energetic 7 MeV MC electrons at $z = 0, 200, -200$ cm.

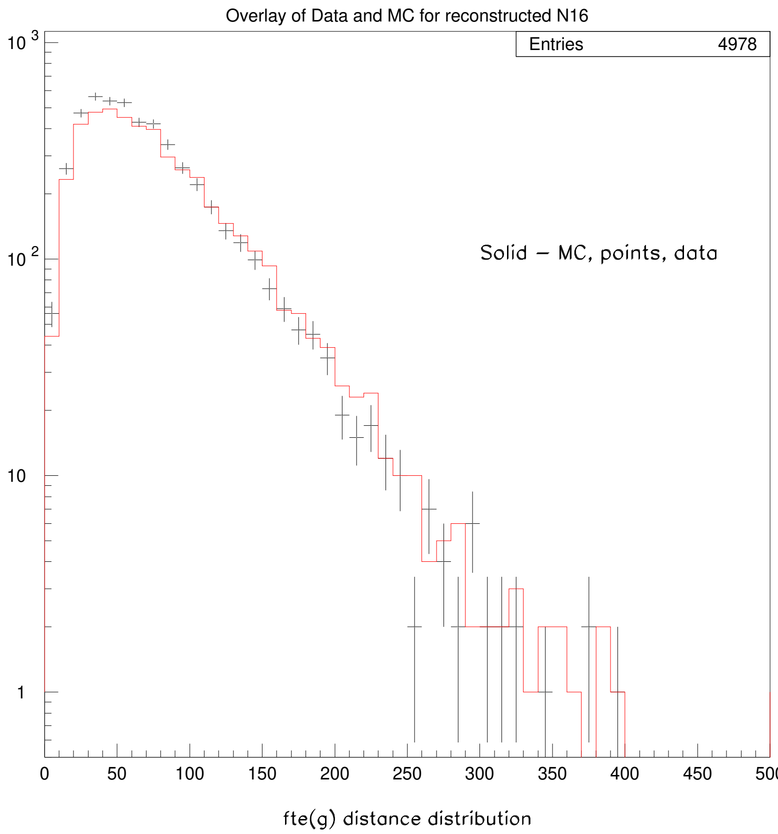

Figure 6.17: Comparison between data and MC ^{16}N position reconstruction accuracy. The histogram is the simulation, and the points are the data. We see that there is qualitative agreement between the data and the MC.

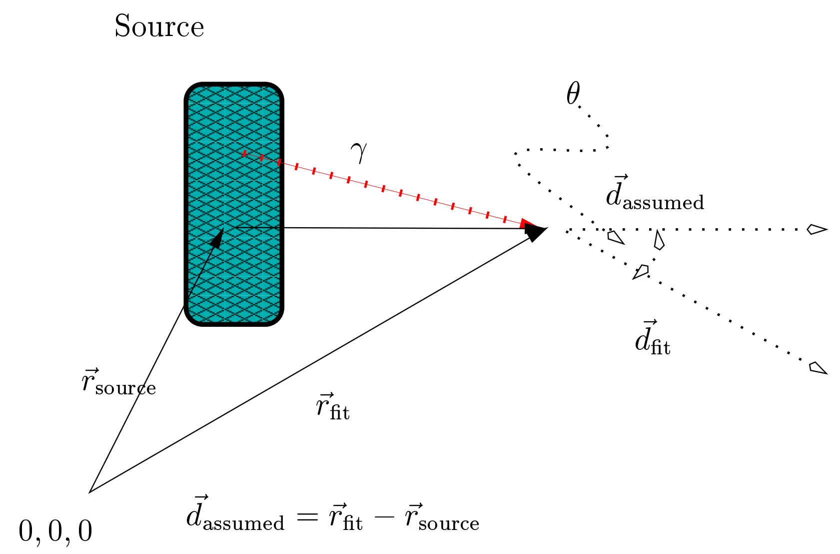

Figure 6.18: Determining the angular resolution using the ^{16}N source. The actual γ-ray track is shown by the dotted line. The γ-ray is assumed to have moved along the line of the vector connecting the center of the source position to the reconstructed vertex position.

6.3.2 Angular resolution

In order to interpret the distribution of neutrino event coordinates, we must know how well the fitter can reconstruct the direction of the incident event. The angular resolution response in SNO is measured with the ^{16}N source.[7] The method is straightforward and easy to explain. It uses a γ-ray source, whose γ-rays travel ~ 35 *cm* before interacting, typically via Compton-scattering. The Compton electron is on average co-linear with the γ-ray and hence with the ray between the source and the interaction point. In analyzing the data, the electron position and direction is reconstructed. The electron's assumed direction is along the γ direction, i.e., along the path from

[7]See, for instance, [58, 59] for extended discussions on measuring this response.

the source to the reconstructed position.

$$\hat{d}_{\text{assumed}} \parallel \vec{r}_{\text{fit}} - \vec{r}_{\text{source}} \tag{6.16}$$

This assumed direction is compared to the direction calculated by the fitter, usually by performing the dot product and plotting $\cos\theta$ between these two vectors. A schematic view can be seen in Fig. 6.18.

$$\cos\theta_{\text{fit}} = \hat{d}_{\text{assumed}} \cdot \hat{d}_{\text{fit}} \tag{6.17}$$

It is clear from the above discussion that this method introduces an unfortunate coupling between the direction fit and the angular fit. This effect can be mitigated by selecting events which reconstruct some distance from the source, say, within 100 and 200 *cm* from the source. At these locations, the smearing of the assumed direction due to reconstruction errors is reduced as the lever arm has a longer baseline.

There is an intrinsic limit to the ability of the SNO detector to measure the initial direction of the electron produced in a neutrino interaction. At solar neutrino energies, the electron undergoes multiple scattering such that its initial direction is hard to determine. However, in addition to this limit, the fitter used to reconstruct the direction of the event has its own resolution that must be considered. Ideally, the effect due to fitter resolution would be comparable to or smaller than the effect of multiple scattering. Barring this, the fitter should clearly introduce as little smearing as possible.

An overlay of the $\cos\theta_{\text{fit}}$ plots is shown in Fig. 6.19. All fitters were run on the

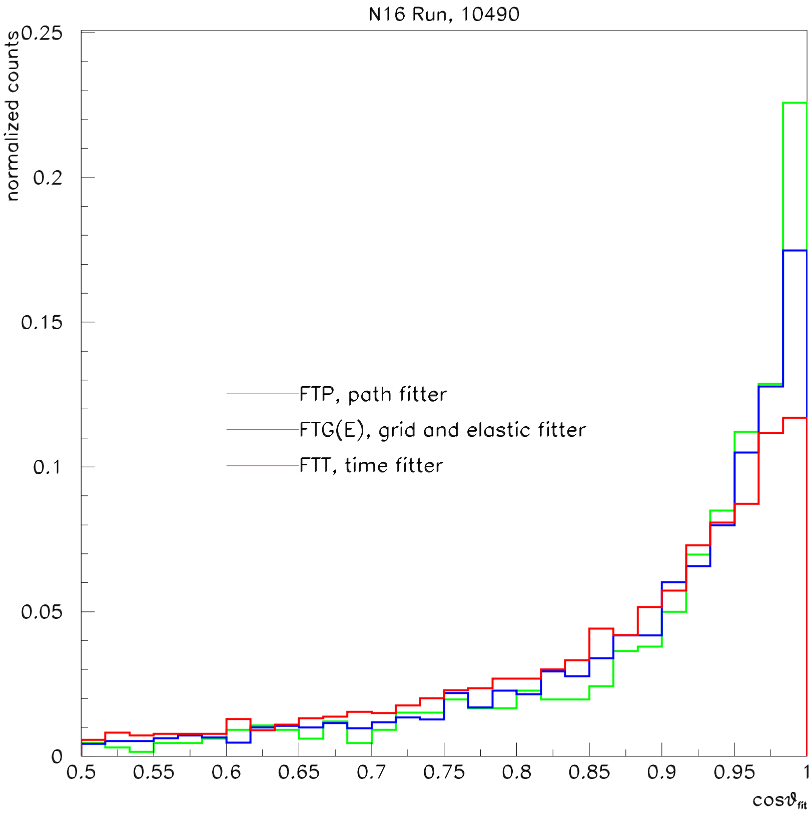

Figure 6.19: Overlay of $\cos\theta_{\rm fit}$ for three fitters, the FTP (path fitter), FTE(G) (grid/elastic combination) and the standard SNOMAN time fitter. The histogram is normalized as drawn, i.e., with the area $0.5 < \cos\theta_{\rm fit} < 1.$ normalized to one. The first 20,000 events from the data cleaning ^{16}N run 10490 at the center make up the data set.

first 20,000 events from a ^{16}N run, 10490, as selected for data cleaning.[8] Fig. 6.19 indicates that the Path Fitter has the best angular reconstruction efficiency, while the elastic/grid combination is second. This is shown by how quickly the peak falls off from the peak at zero. A third fitter, the time fitter, which does not explicitly fit for direction, has the largest tail at large $\cos\theta_{\rm fit}$. The time fitter is least steeply peaked in this area. As would be expected, explicitly fitting for the direction improves the angular resolution for Čerenkov light.

Fig. 6.20 shows the resolution of the FTE(G) fitter in more detail. The data are a ^{16}N run 10490 with the source at the center (0,0,0cm). Ignoring the backward peaking described above, the data can be fit as a function of two exponentials.

$$R(\theta, \phi) = \exp(\alpha + \beta \times \cos\theta) + \exp(\gamma + \delta \times \cos\theta) \tag{6.18}$$

where the first exponential models the sharp part of the resolution and the second the longer tail. This resolution function in $\cos\theta$, along with a flat resolution function in ϕ, allows me to model the effect of changes in resolution on a final measurement.

Fig. 6.21 is a repeat of the angular resolution measurement using the Monte Carlo. We see that the steep part of resolution function (Parameter P4 in both figures) is sharper in the simulation that in the data.

At this point, we have events whose energy we can estimate, and whose direction and positions have been reconstructed. We presented an estimate on the position and angular resolution of the sources.

[8]See Eq. 6.10 and Fig. 6.14 for the ^{16}N decay schemes. For data cleaning, all events where the β-ray fires the scintillator make up a file of 'pure' γ-ray events.

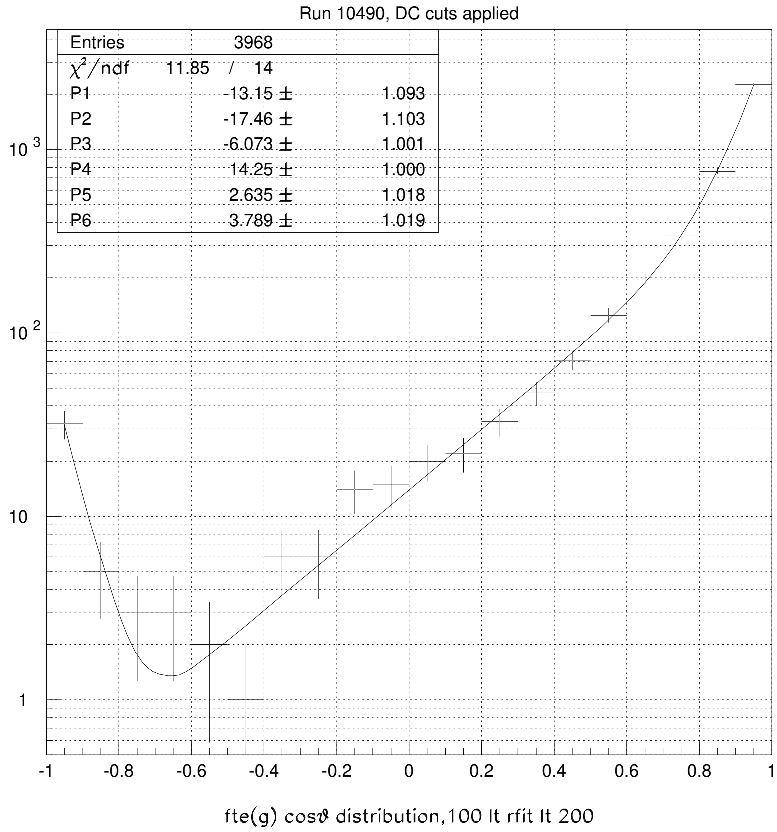

Figure 6.20: Angular resolution, as measured by the FTE(G) fitter. An additional cut 100 $cm < r_{\rm fit} < 200$ cm was applied to minimize the effects of misreconstruction. We have included a third exponential to model the backward peaking, though this has little effect on the actual resolution measurement.

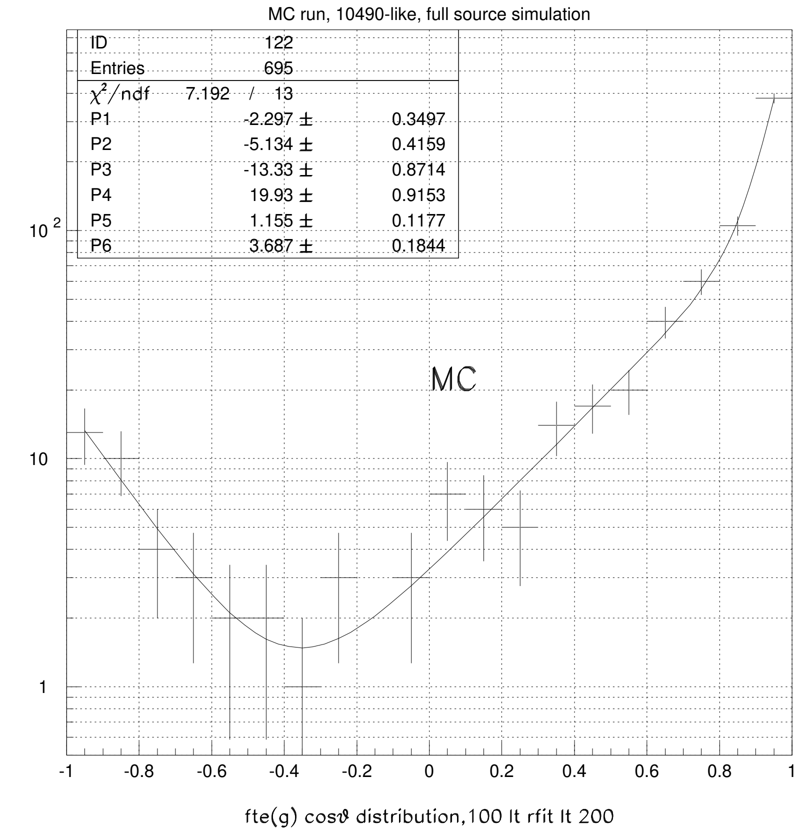

Figure 6.21: Angular resolution check with the SNOMAN MC, full source ^{16}N run.

6.4 Data Set and Run Selection

The data used in this analysis consists of a series of runs which were taken from the first phase of SNO data data taking. In SNO, a run is defined as a set of events that have the following properties.

- A consistent detector physics configuration.
- Maximum length of less than 96 hours.

A consistent detector configuration means that for the sake of analysis, nothing should change during a run that affects the physics data quality. It is the job of the run selection committee, to define how SNO is to choose runs for solar neutrino analysis. The selection depends on the data files and the shift reports.

The data files themselves contain most of the information needed to classify a run. At the start of each run, the DAQ system puts headers into the data file which describes the baseline configuration of the run. These headers define which triggers are enabled, at what levels they are set, which crates are on during this run, which sources, if any, are deployed, and the run number and starting time of the run.

In addition to the header information, we read the shift report to see whether to cut the run or not. Usually, the shift reports contain reports of power failures and other unscheduled changes to run conditions which disqualify runs. It is important to decided *a priori* what can be used to disqualify a run, and it is also important to understand how variations in reporting by different shift operators can affect whether runs get cut or not. The selection criteria listed by the run selection committee is shown in Tab. 6.1. Information from the run and trigger headers and the from DQXX

Run and Trigger Header	
Bits Set	NEUTRINO_RUN D2O_RUN COMP_COILS
Bits **not** set	TRANSITION_RUN EXPERIMENTAL_RUN SOURCE_MOVING_RUN SOURCE_RUN ECA_RUN, PCA_RUN CALIB_RUN DIAGNOSTIC_RUN
Trigger Header	
Trigger	`NHIT100MED`
DQXX Banks	
Crates	All Crates online
Channels	Fewer than 30×32 channels off-line

Table 6.1: Header information used to select runs.

banks is used.[9] In addition, the runs are rejected if any of the following conditions is met.

- Run is less than 30 minutes long.
- Lab activity in the deck causes too much pickup or light in the vessel.
- Hardware errors such as a power failure occurred during the run.
- Software errors makes tools needed for the run unavailable.

The run length is obtained from the data file itself. The decision when to cut a run based on activity in the lab is the most worrisome, as it is subjective. The information is based on the shift report, which is filled out by a continuously changing array of

[9] See Sec. 4.7 for information about the DQXX banks.

Type of Cut	Time	Fraction
Live time Summary, complete set		
Total Time	128.5 days	
Muon Follower Short Dead Time	53.6 h	1.7%
N_{Hit} Burst Dead Time	2.6 h	0.09%
Burst Dead Time	4.7 min	2.5×10^{-3}%
Retrigger Dead Time	6 s	5.4×10^{-5}%

Table 6.2: Live time and time cut by data cleaning cuts.

SNO scientists. However, at worst, this causes a decrease in livetime (and hence a decrease in sensitivity). See [60] for detailed information about run selection.

For this thesis, we analyzed data for the first fifteen production runs, tapes FD0000-FD0014. This spans the beginning of production data taking in November 2, 1999 up April 27, 2000. Tab. 6.2 shows the live time information for the selected runs, including what data is removed by various levels of data cleaning.[10] The total available live time is roughly 128 live days.

6.5 In Conclusion

In this Chapter, we showed the analysis path from the raw data stream through the data reduction path, to the event reconstruction of energy, position and direction. We are now in a position to analyze the data for physics content.

[10]The data set used here is the same as the data set used by the collaboration for analysis for the ν2000 conference. In that case, however, Set C was not used.

Chapter 7

Physics Results

This section will introduce the signal, and show that it is consistent with ^{8}B neutrinos from the sun interaction though the charged current and elastic scattering reactions. The main signal properties that we will look at are the angular distribution with respect to the current solar direction, the radial distribution of reconstructed events, and the distribution of the number of hit PMTs. Next, we will extract the rate of CC reactions due to solar neutrinos, and calculate the ratio of measured rate *vs* SSM expected rate.

7.1 Data characteristics

In this section, we will use a subset of the data which we selected for the CC data. We use only data splitting subsets A&B, as discussed in Sec. 6.1.1. In addition, not all tools that were needed for one of the 15 data tapes were in place when this analysis was done, so tape FD0014 was excluded. Tab. 7.1 shows the live time of the reduced data sample, excluding the data splitting reduction. This splitting reduces the live

Type of Cut	Time	Fraction
Live time Summary, without tape 14		
Total Time	118 days	
Muon Follower Short Dead Time	49.0 h	1.7%
N_{Hit} Burst Dead Time	2.6 h	0.09%
Burst Dead Time	4.7 min	2.8×10^{-3}%
Retrigger Dead Time	6 s	5.8×10^{-5}%

Table 7.1: Live time and time cut by data cleaning cuts, without tape FD0014.

time by 1/3.

Often, in these types of analyses, the signal is presented after a list of cuts are shown and justified. Here, we will take a slightly different approach as the cuts that were applied are motivated by physical *a priori* arguments rather than by tuning on a smaller data set.

7.1.1 Angular distribution

Fig. 7.1 shows the reconstructed $\cos\theta_\odot$ for this data set, where the angle measured is the reconstructed direction and a vector anti-parallel to the instantaneous solar direction:

$$\cos\theta_\odot = -\hat{d} \cdot \hat{r}_\odot \tag{7.1}$$

The other cuts placed on this data are straightforward, and no further tuning was applied.

- $r_{\mathrm{fit}} \leq 6m$.
- $\mid \cos\theta_\odot \mid \leq 1$, i.e., fitter converged.

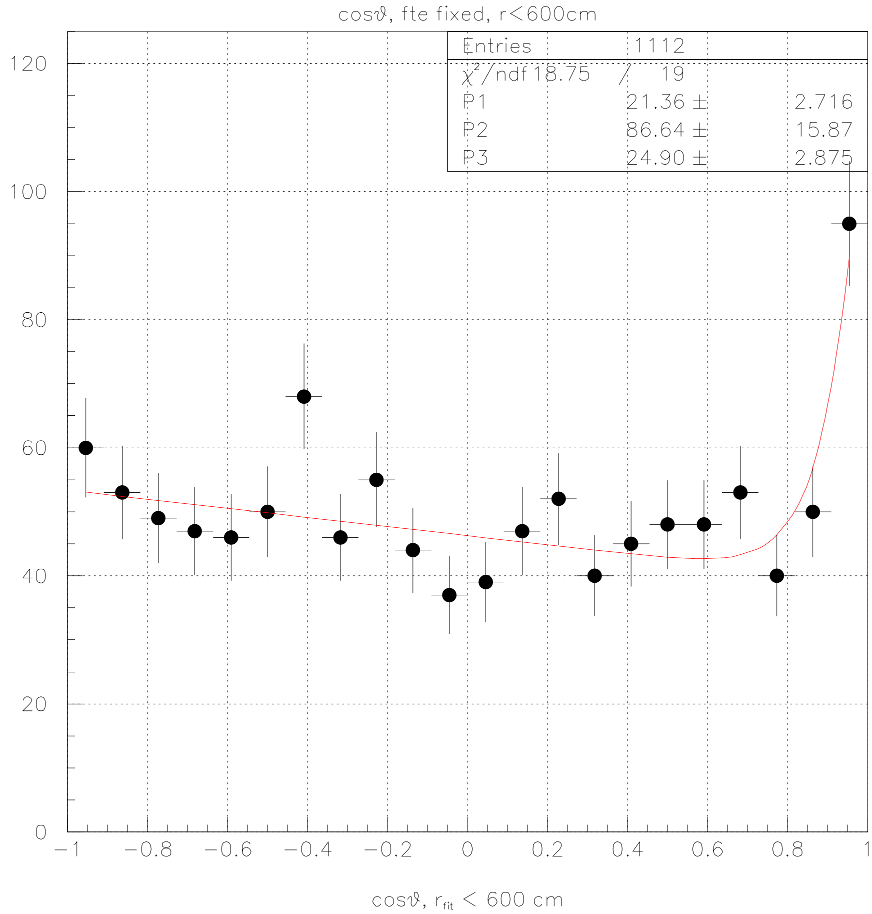

Figure 7.1: $\cos\theta_\odot$ distribution.

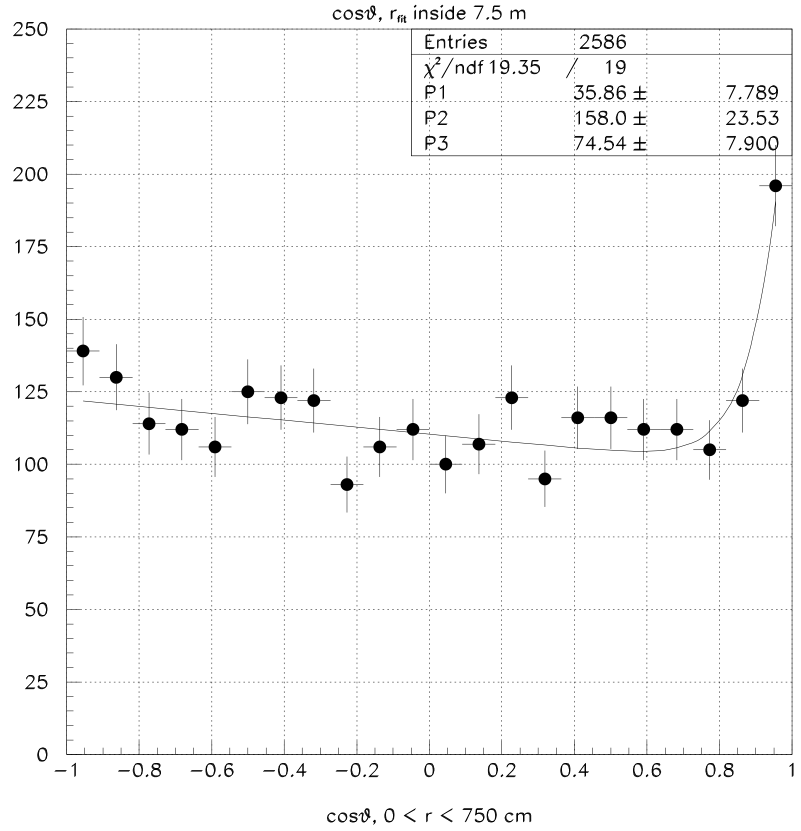

Figure 7.2: $\cos\theta_{\odot}$, with $r < 750\ cm$.

- $N_{\mathrm{Hit}} \geq 45$, which corresponds to a ball park energy threshold of 5 MeV in the heavy water (see Sec. 6.2).

These cuts are motivated by physical arguments rather than by tuning, even on a smaller data set.

Fig. 7.1 shows a convincing neutrino-electron elastic scattering. The sharpness of the peak can be attributed to the good direction reconstruction ability of the FTE(G) fitter. This peak is strong evidence that SNO does indeed see solar neutrinos. If the goal is to maximize the size and significance of the ES peak, the fiducial volume cut of 6 m can be changed to 7.5 m to include the light water reaction region for the ES reaction, up to the concentrator cutoff. Fig. 7.2 shows this data, and we see how the peak has increased. Comparing the contents of the rightmost peak in Fig. 7.1 and Fig. 7.2, we see that the contents of that bin have more than doubled. A simplistic counting of the height of the signal above an estimate of the CC and background signals shows an increase from $\sim$ 50 counts to $\sim$ 90 counts, again, consistent with the volume increase. However, as discussed in Sec. 6.2, the energy response of SNO changes for events in the light water, so the volume scaling alone should not account for differences.

The slight slope of the signal away from the ES peak in Fig. 7.1 is our first evidence that we are seeing neutrinos via the CC reaction. The drawn fit itself is a fit to the function

$$\mathcal{F}(\cos\theta) = \alpha \times (1 - \frac{1}{3}\cos\theta) + \beta \times \exp(-12 \times \cos\theta) + \gamma \tag{7.2}$$

where α, β and γ are free parameters. The parameter α characterizes the strength of the CC signal, β the strength of the ES signal, and γ the amount of flat background.[1] The flat background consists of radioactivity in the light and heavy water, and of neutron captures on deuterium, including neutrons that are created in the NC reaction. Due to the relatively weak angular dependence of the CC flux on the solar direction, the CC signal α and flat background γ are strongly correlated. The

[1] In Fig. 7.1, $\alpha \equiv P1, \beta \equiv P2$, and $\gamma \equiv P3$.

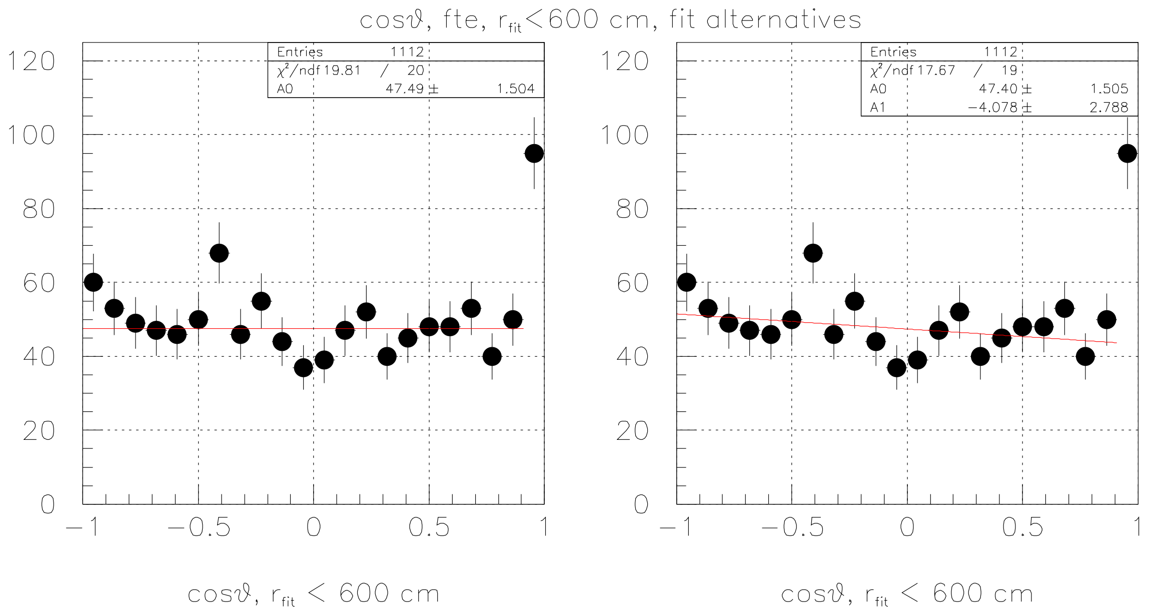

Figure 7.3: Alternative fits of a flat background and a linear function to $\cos\theta_{\odot}$ distribution. For the flat hypothesis (no CC signal), the fit has a goodness of fit of $\chi^2/n_{dof} = 19.81/20$. For the simple linear fit, the $\chi^2/n_{dof} = 17.67/19$.

correlation matrix for α, β, and γ for the fit in Fig. 7.1 shows this.[2]

$$\mathcal{M} = \begin{pmatrix} 1.000 & 0.435 & -0.985 \\ 0.435 & 1.000 & -0.471 \\ -0.985 & -0.471 & 1.000 \end{pmatrix} \tag{7.3}$$

$\mathcal{M}_{13}$ shows that α and γ are strongly correlated.[3] This is a practical example of the correlation that has been shown to make extracting a CC signal from the angular distribution alone in the presence of backgrounds difficult [61]. For this reason, for the CC rate measurement, we will go to a "background-free" region where $N_{\text{Hit}} \geq 60$. Fig. 7.3 shows an attempt to separately fit a constant and an unconstrained first

[2]The fit was performed using PAW Minuit interface.
[3]Even α and β are strongly correlated.

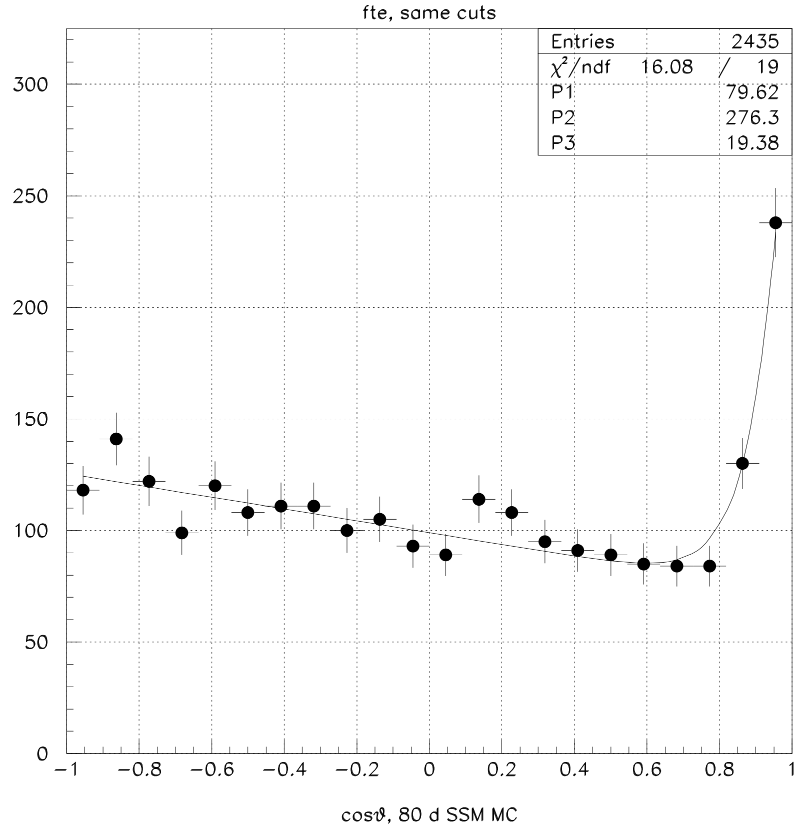

Figure 7.4: Monte Carlo Simulation of 80 days of neutrino data for the null hypothesis, standard solar model with no new neutrino physics. Note the increased total number of counts relative to Fig. 7.1 .

order polynomial to the data, excluding the ES peak in the rightmost bin. These tests determine if the data is consistent with no signal, i.e., a flat $\cos\theta_\odot$ away from the sun. For the null hypothesis (no CC signal), the fit $\chi^2/n_{\rm dof} = 19.81/20$. For the simple linear fit, $\chi^2/n_{\rm dof} = 17.67/19$. While the null hypothesis has a good χ^2, the non-flat χ^2 is "better," i.e., it is smaller. Both fits have good χ^2.

For comparison, Fig. 7.4 shows a Monte Carlo $\cos\theta_\odot$ distribution for 80 days of neutrino live time, a number that is similar to the two-thirds of the total live time

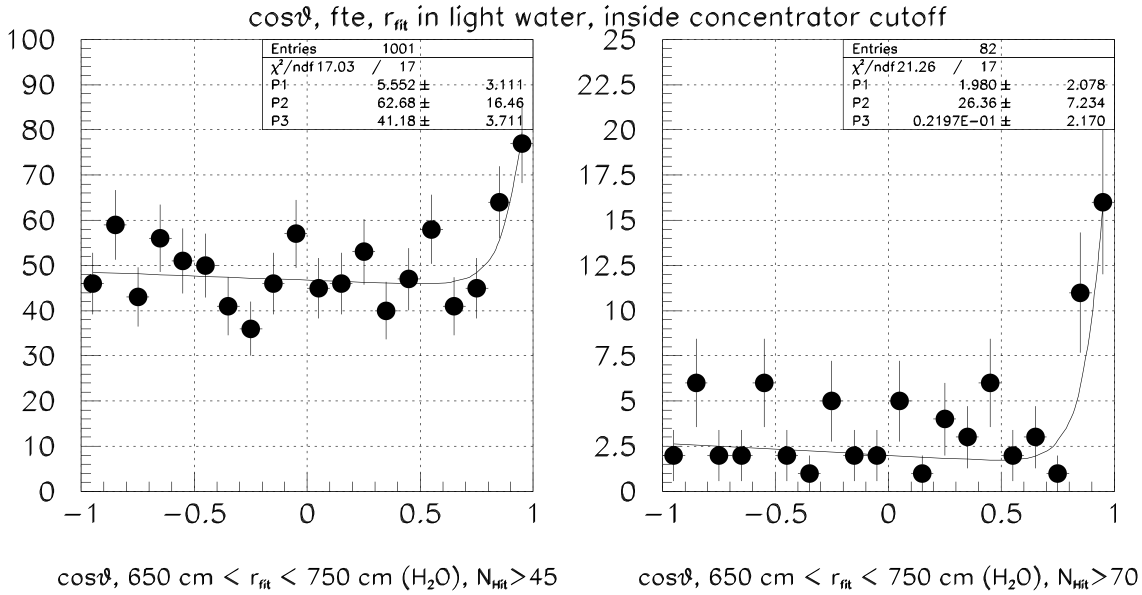

Figure 7.5: $\cos\theta_\odot$ in the light water for this data set, with two different $N_{\rm Hit}$ cuts. No believable ES peak is visible, though the peak with a higher $N_{\rm Hit} > 70$ cut is tantalizing.

listened in Tab. 6.2 for the data without tape 14. Note that the total number of counts is much higher than in Fig. 7.1.

As discussed in the section on the ES signal, one can now look at the $\cos\theta_\odot$ distribution in the light water. Fig. 7.5 shows this distribution. In the plot on the left, with a $N_{\rm Hit} \geq 45$ threshold, the ES peak is visible if one considers the two rightmost bins to contribute. With a slightly higher $N_{\rm Hit} > 70$ cut, where more background should be gone, the peaking is stronger. However, the light water data does not show the strong peaking in the first bin that the heavy water data did. One difficulty is that the data presented here is mostly in the regime for $r_f > 7$ m, where the concentrators are designed to cut off. This effect could wash out the resolution, leading to a smearing of the ES peak compared to what one sees in the heavy water data.

Possible sources of bias

There are a number of ways a signal can be induced to give a correlation with the sun. A simple way to obtain such a correlation is to have a short, high-rate burst of data at one point in the detector. This is relatively easy to notice with cuts such as the data cleaning group's burst cut or N_{Hit} burst cut. A more nefarious scenario would be to have a phenomena that occurred at one point in the detector at a certain time - such as, for instance, neck events occurring during blasting.

This scenario can be examined by looking at the data cleaning neck event data sample. By reconstructing these events with the elastic fitter and looking at $\cos\theta_{\odot}$, we can study this possibility. Fig. 7.6 shows the distribution of $\cos\theta_{\odot}$ for neck events from the data cleaning sample, DC_neck_10779_10801_10922_11269.zdab. A peaking anti-correlated with the solar direction is evident in this plot, which is a signal similar to that of the CC reaction.

How this correlation between the sun and the neck events occurred is shown in the event time distributions shown in Fig. 7.7. Neck events in this sample appear to occur preferentially in the early evening Sudbury time. At this time, the sun is in the lower hemisphere of the earth in Sudbury, while the neck events point downwards. This is all that is needed to get a strong correlation between a background and the solar direction. If neck events occurred uniformly in time, this distribution would be smeared out.[4]

We have now shown how a correlation *could* arise. Do the neck events in fact give such an effect? The neck data sample might just have sampled the evening hours preferentially, presenting a selection bias. In order to determine if this is the case, we

[4]Since the sun doesn't travel across all angles in the sky, the distribution would not be completely flat.

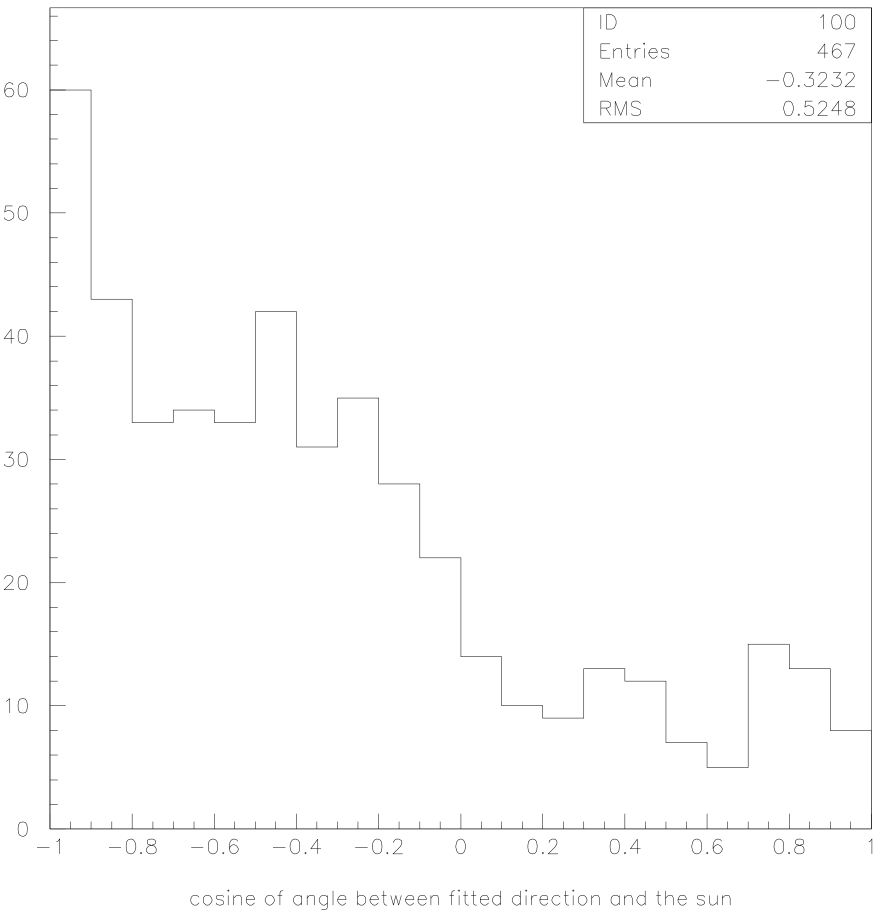

Figure 7.6: Distribution of $\cos\theta_{\odot}$ for neck events from the data cleaning sample, DC_neck_10779_10801_10922_11269.zdab. In this plot, the elastic scattering peak from a solar neutrino signal would occur at the right.

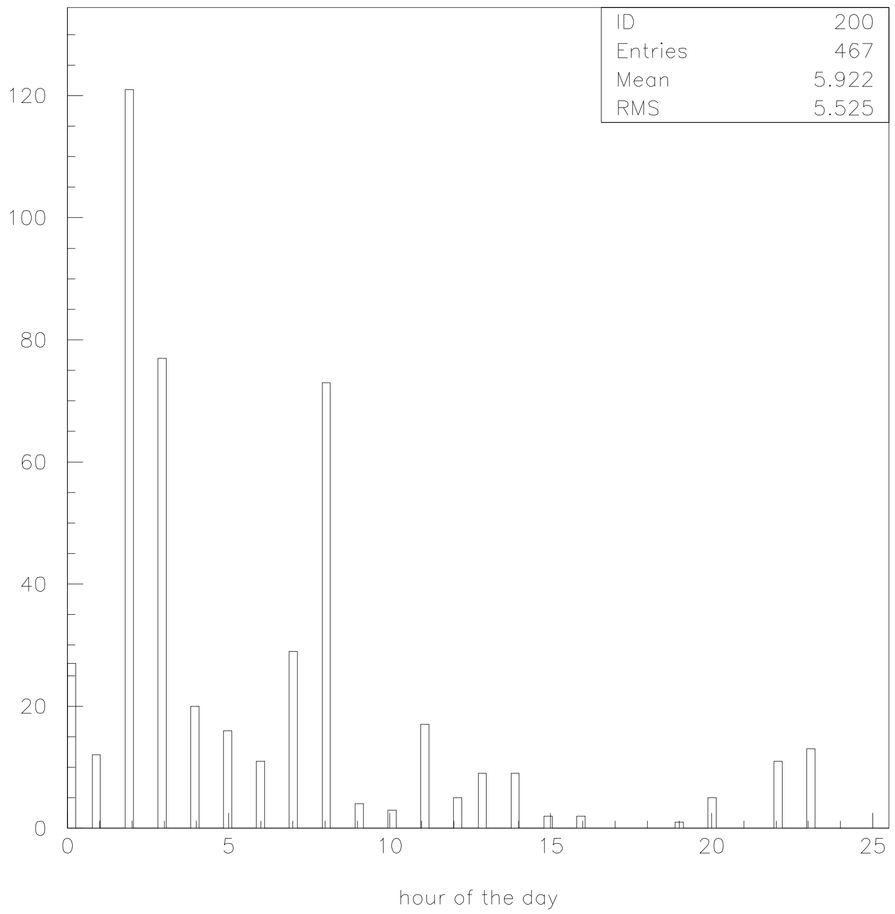

Figure 7.7: Event times in hours for the neck data sample shown in Fig. 7.6. Times listed are in UTC, so the big peak at 3 corresponds to 10 PM Sudbury time.

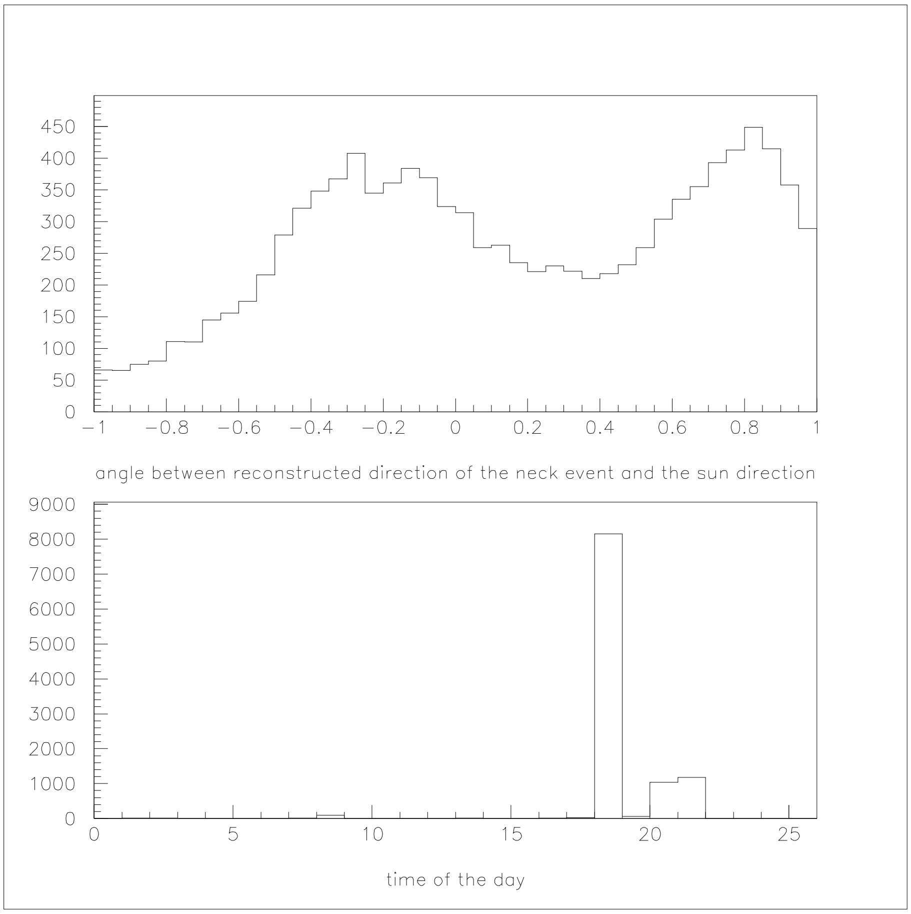

Figure 7.8: Distribution of $\cos\theta_{\odot}$ for neck events (top plot) and their times in hours (bottom plot). Note that this plot is qualitatively different from Fig. 7.6.

can look at the times of all neck events tagged by the data cleaning neck cut for a large set of data, in this case, all data on the first production data tape, FD0000.[5] Fig. 7.8 shows the distribution of $\cos\theta_{\odot}$ for all neck events on the tape FD0000 and the times at which they occurred, in hours. This plot is qualitatively different from Figs. 7.6-7.7. The time peak is at a different time, and the peaks themselves appear to have

[5]This run spans the time before the timing was changed on the neck tubes, and hence the efficiency of the neck cut is somewhat reduced. However, this inefficiency should be a global effect and not affect this test.

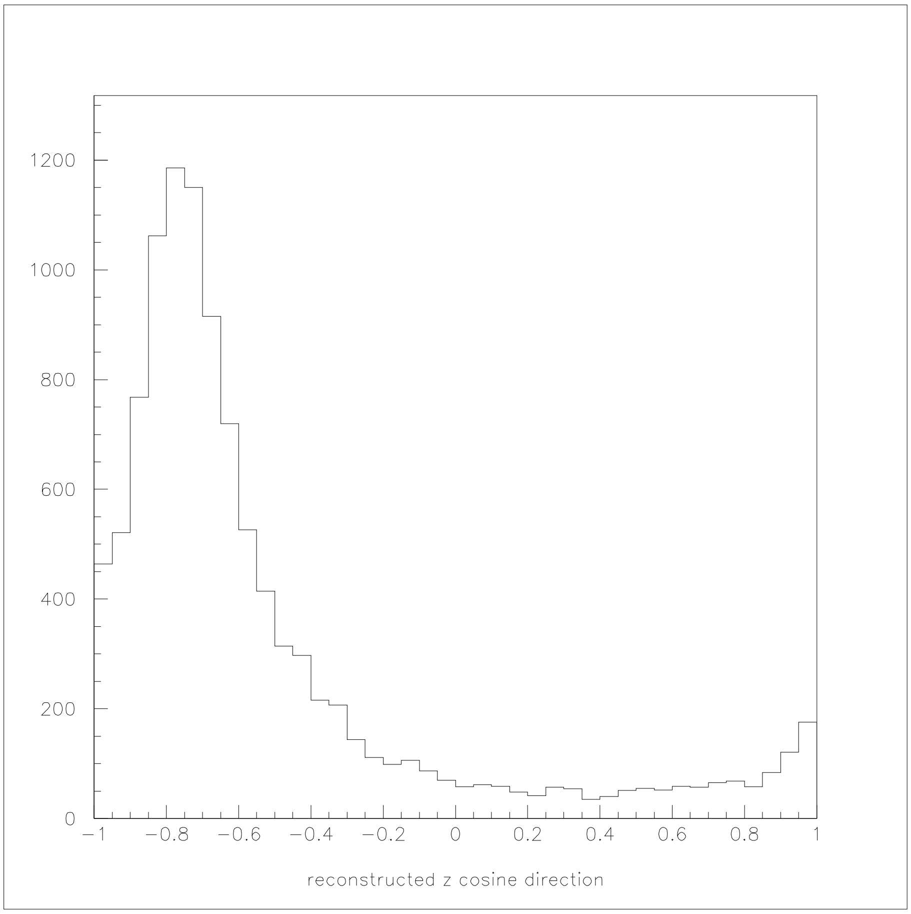

Figure 7.9: Z direction cosine of reconstructed neck events from tape FD0000.

occurred in bursts (this is not apparent from this plot.)

Though we have shown that neck events *can* cause a bias in the $\cos\theta_{\odot}$ distribution, it is not clear that neck events *do* cause this effect. If there were such an effect in the data, it would be even more obvious in a fixed coordinate in the detector. Fig. 7.9 shows the z direction cosine, $\cos z$, of events tagged by the neck cut. Note that the peaking in this distribution at $\cos z = 0$ is more pronounced than the equivalent peaking in Fig. 7.8. Therefore, an equivalent plot for the neutrino data sample should expose detector-relative fixed points in the data. Fig. 7.10 shows this distribution for

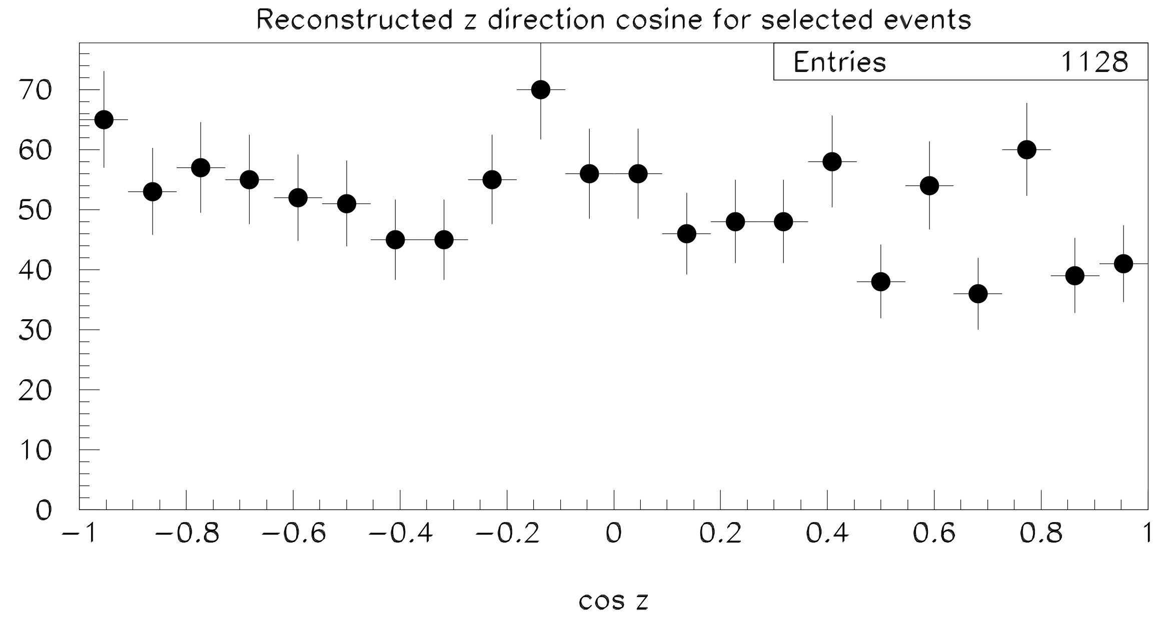

Figure 7.10: $\cos z$ for neutrino data sets A & B.

the neutrino data set. There does not appear to be any significant structure in this plot, indicating that we do not have neck event contamination. Due to the AV neck, there are acceptance effects near $\cos z \simeq 1$, which would lead to an expected decrease in this range. There is a hint of this in the data.

Alternatively, you could speculate that the fitter itself is somehow biasing your sample to create, if not the ES peak, at least the slight CC slope. As a check, we look at the $\cos\theta_{\odot}$ plot for the Path Fitter, as discussed in Sec 6.3.1. The Path fitter shares no common code with the elastic or grid fitters, and while it too uses photo-tube timing for the fit, its fitting method is quite independent of the other fitters and it should not be susceptible to the same biases. Additionally, another Penn graduate student, Mark Neubauer, performed this analysis. While this hardly constitutes two independent analyses, it is still a minimal independence cross-check. Fig. 7.11 shows this angular correlation of the neutrino with the sun as measured using the Path

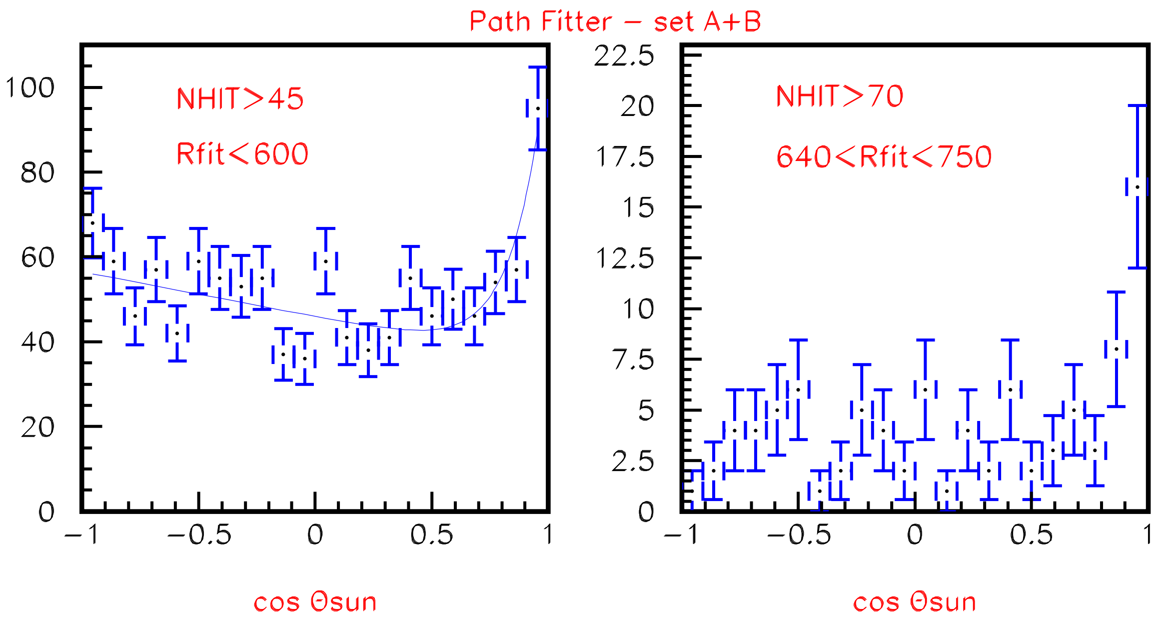

Figure 7.11: $\cos\theta_{\odot}$ distribution generated using the FTP Path fitter. Both the heavy and light water plots are shown. This check, with an independent fitter, demonstrates consistency with the $\cos\theta_{\odot}$ shown in Fig. 7.1 and Fig. 7.5. Figure courtesy of Mark Neubauer [51].

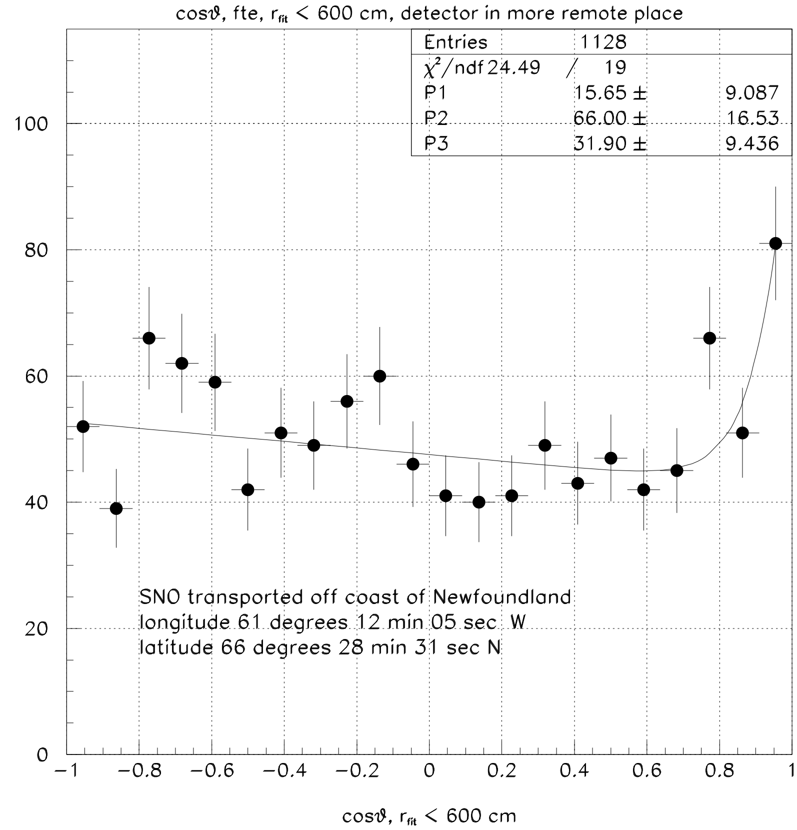

Figure 7.12: $\cos\theta_\odot$ for changed detector coordinates to off coast of Newfoundland. A peak at $\cos\theta_\odot = 1$ is still visible, though smeared. We applied the same cuts as in Fig. 7.1.

Fitter. The analyzer applied the same radial and volume cuts as described above.

This plot is qualitatively the same as Fig. 7.1, and is evidence that the signal behavior observed is not due to a pathology in the FTE(G) fitter.

Another systematic check we perform is testing to see if we know where the detector is, and if we can still produce a signal in the case of getting the detector location wrong. Figs. 7.12 and 7.13 show the result of changing the location of the detector as defined in the SNOMAN source code `solar.for` to two different locations, and

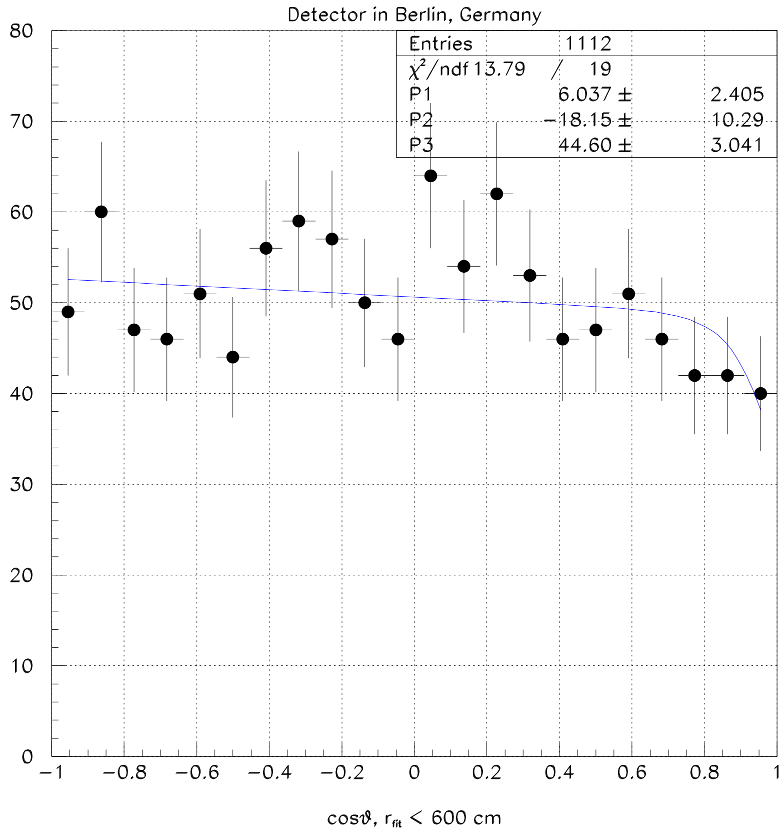

Figure 7.13: $\cos\theta_{\odot}$ for changed detector coordinates to Berlin, Germany. No *ES* peak is visible.

repeating the analysis without making any other changes. In the first plot, the given coordinates were changed by 20° for both longitude and latitude. This places the detector off the coast of Newfoundland. In this case, the $\cos\theta_\odot$ plot still shows a peaking at $\cos\theta_\odot = 1$, but it is reduced and smeared. The second plot places the detector squarely in the center of Berlin, Germany. Here, the $\cos\theta_\odot$ plot shows no visible ES peak. We conclude that the analysis itself does not produce peaks in $\cos\theta_\odot$, and that the forward peak is due to physics processes inside the detector.

7.1.2 r^3 distribution

The next figure, Fig. 7.14, shows the distribution of radii of reconstructed events using the usual FTE(G) fitter, volume-weighted, in the light and heavy water. The distribution shows a variety of features, some of which are intrinsic to the geometry of the detector and some which are attributable to the detector itself.

The first feature appears in the heavy water. Here, the r^3 distribution is far from flat, as one might expect for a clean neutrino signal. This effect is partially due to how neutrons behave in SNO. Due to the long mean free path for neutrons in heavy water, many neutrons will disappear from the heavy water by being captured in the light water or in the acrylic vessel. Fig. 7.15 demonstrates this effect with a simulation of neutrons generated uniformly in position and isotropically in direction in the heavy water. The radial distribution shows a strong fall-off in r^3.

The neutrino-induced CC electrons themselves contribute to the slope due to the boundary at the AV. As there are no deuteron at $r > 600$ cm, the rate of CC electrons drops to zero here. However, the finite reconstruction resolution causes leakage of events into the light water. This leakage is shown in Fig. 7.16.

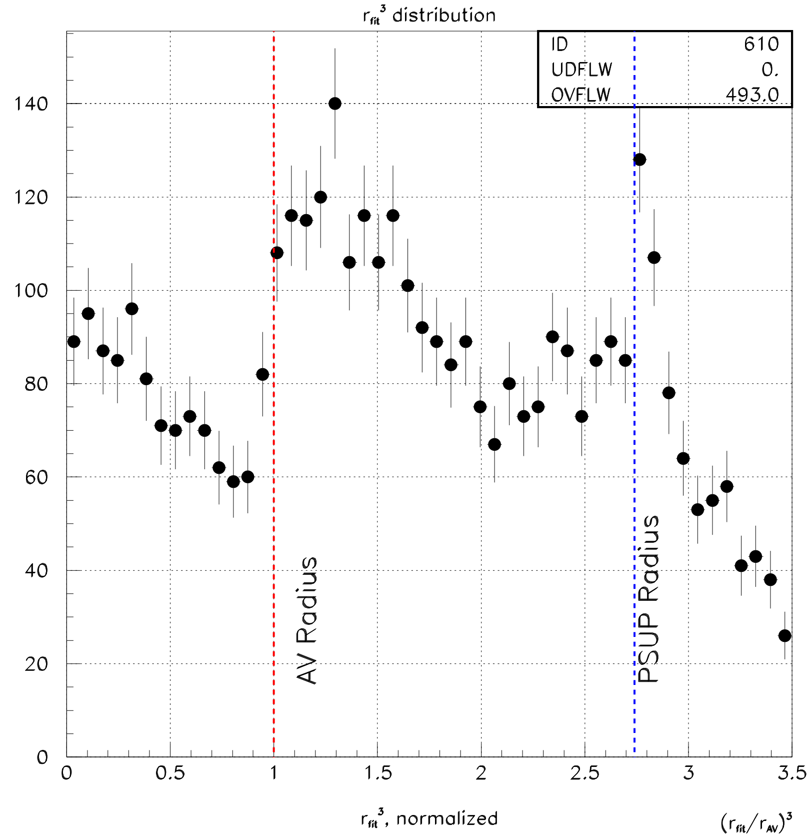

Figure 7.14: Distribution of reconstructed radial position in units $(\frac{r_{\text{fit}}}{r_{\text{AV}}})^3$ for events in the cleaned data set. An N_{Hit} cut of 45 tubes is included in this plot.

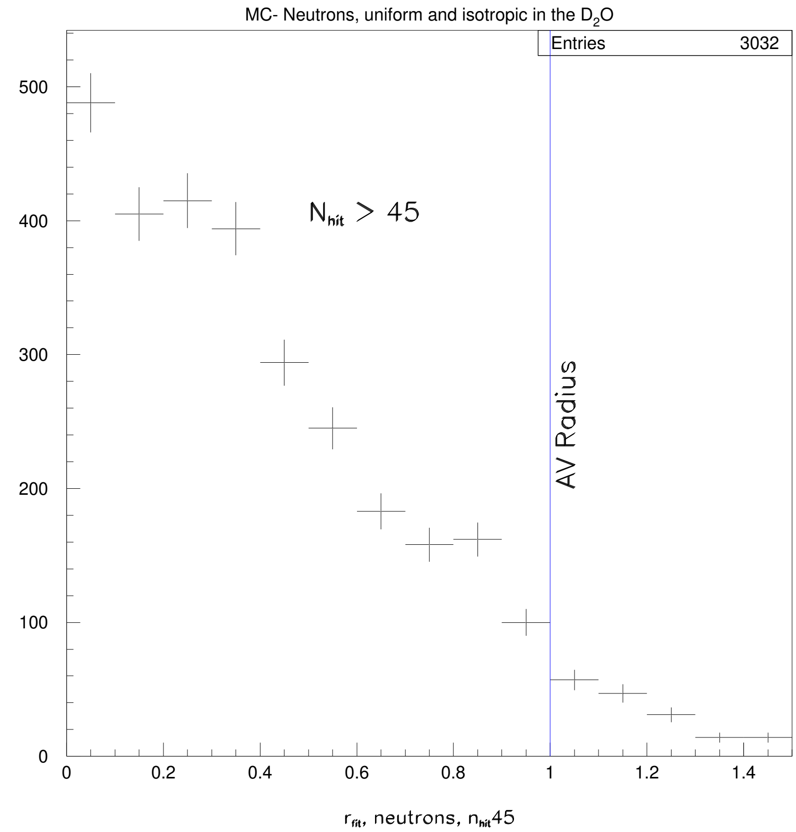

Figure 7.15: Radial distribution in units of $(\frac{r_{\rm fit}}{r_{\rm AV}})^3$ for reconstructed events of a simulation of neutron captures on deuteron, for Phase I (D_2O only). The neutrons were generated uniformly in space in the D_2O volume. Due to the large capture cross section in the AV and in the H_2O, many neutrons wander outside the fiducial volume of the detector and are not detected.

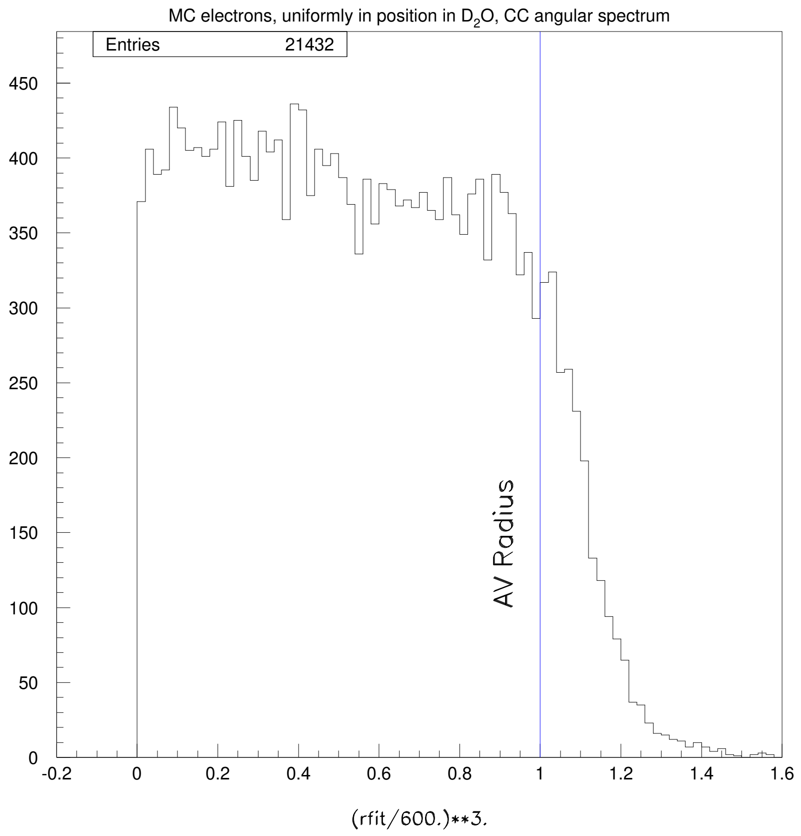

Figure 7.16: Simulation of radial distribution of events for electrons in units of $(r_{\rm fit}/r_{AV})^3$. We simulated events in the heavy water volume ($r_{\rm fit}/r_{AV} < 1.$) only. However, the finite resolution of the fitter causes events to leak into the light water volume, giving an apparent slope on the r^3 distribution.

The next feature is the dip near $r/r_{AV} = 1$, i.e., at the edge of the acrylic vessel. Besides the resolution and neutron effects described above, we have observed this effect in the detector simulation of electron events. This dip is due to the complicated optics near the AV, and is also reproduced in the simulation, as shown in Fig. 7.17. The next feature is the large rise outside the AV, followed by a sharp cutoff at $r^3 \simeq 3 \times 10^8\ cm^3$. Two effects are contributing to this shape. First, in the light water, backgrounds such as PMT $\beta - \gamma$s are rapidly rising in numbers from the AV until they peak at the PSUP. At roughly 7 m, the light concentrators start cutting off the response of the PMTs. This leads to the characteristic drop-off of the N_{Hit} spectrum seen in Fig 7.17. Fig. 7.18 illustrates the effect of concentrator cutoff pictorially.

This cutoff does not explain the sharp rise in the number of events at the AV. The first temptation is to attribute this rise solely to radioactivity in the light water. While this is certainly important, acceptance effects can also contribute to this type of peaking. A large bias is introduced by using N_{Hit} as an energy estimator for the threshold. From the ^{16}N source, we know that the light output of events in the light water is about 10% higher than in the heavy water, due to UV light absorption in the AV. Fig. 7.17 demonstrates that, using mono-energetic electrons, the SNO detector can produce similar effects. While the initial distribution is flat in r^3, placing an N_{Hit} cut on the data produces a shape near the AV like what is seen in Fig. 7.14. The peaking in the light water is therefore at least partially due to an acceptance effect.

Fig. 7.19 shows an r^3 distribution of the same data set with an additional $N_{\mathrm{Hit}} > 60$ cut. The large peak in the light water has disappeared, and the amount of data per unit volume is now higher in the heavy water than in the light water. A credible interpretation of this change is that by cutting lower N_{Hit} data, the ratio of signal to noise has been increased, as radioactive backgrounds have a spectrum that peaks at

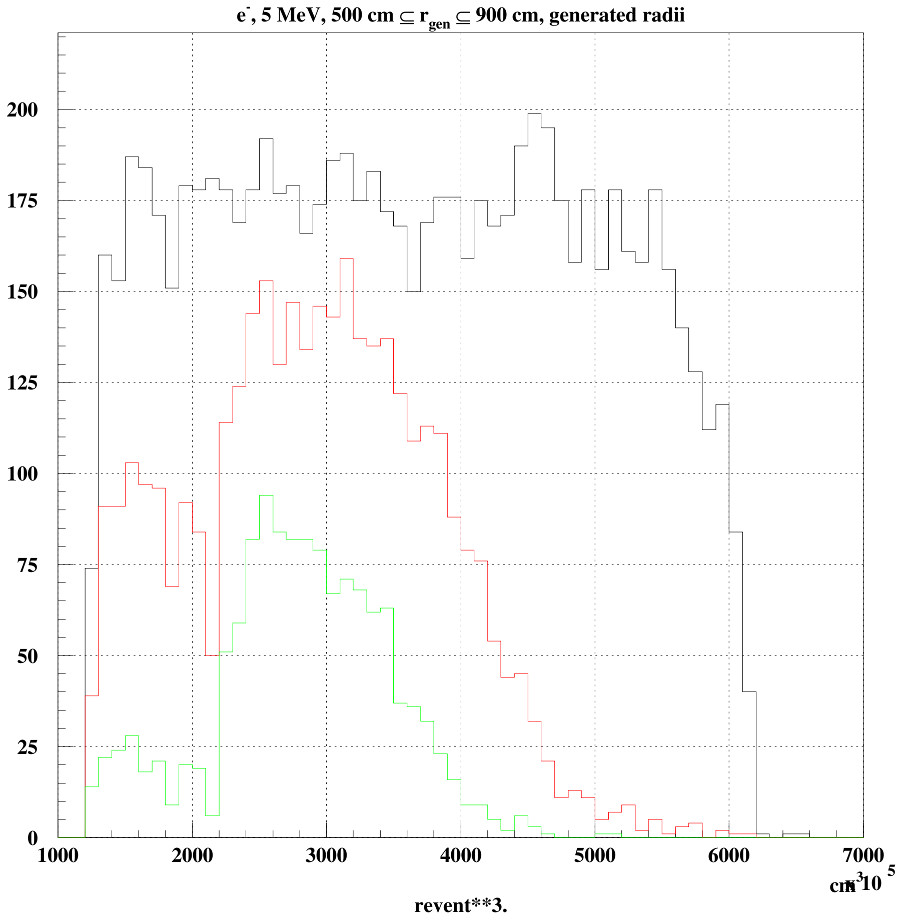

Figure 7.17: MC r^3 distribution for mono-energetic electrons generated uniformly in position and isotropically in direction inside and outside the AV, showing the effect of N_{Hit} cuts and of detector geometry. Events were generated with 500 cm$< r <$ 900 cm, and no reconstruction cuts were applied. The black (top) line is the generated distribution, the red (middle) and green (lowest) lines are $N_{\mathrm{Hit}} \geq 50, N_{\mathrm{Hit}} \geq 40$ cuts, respectively. The effect of N_{Hit} as an energy estimator is visible by the different reduction in the light and heavy water. Geometric effects near the AV are also visible. In units of r^3, the edge of the AV is at $r^3_{AV} = 216 \times 10^6\ cm^3$, where the dip in r^3 is visible.

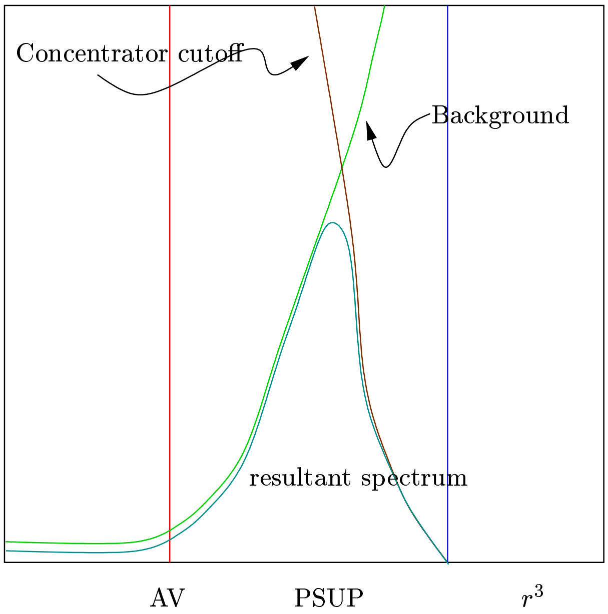

Figure 7.18: Effect of concentrator cutoff *vs* rapidly rising backgrounds

low $N_{\rm Hit}$.[6] Fig. 7.19 strongly suggests that we have an excess of events in the heavy water, which is consistent with neutrino interactions. It is interesting to note that the "signal" extends one bin (roughly 30 cm in r) outside the AV. A similar effect is visible in Fig. 7.14, where the bin closest to the AV radius inside the AV appears to be an extension of the signal in the light water. Both of these effects are due to the finite position resolution of our fitter.

We conclude from the r^3 distributions that the physics events look different in the light and heavy water.

7.1.3 $N_{\rm Hit}$ distribution

We will now discuss the $N_{\rm Hit}$ distribution. Fig. 7.20 shows an overlay of a Monte Carlo simulation on data. The Monte Carlo $N_{\rm Hit}$ plot is scaled by a factor 1.063 to

[6] See Fig. 6.5 for a raw $N_{\rm Hit}$ spectrum, which shows this peaking.

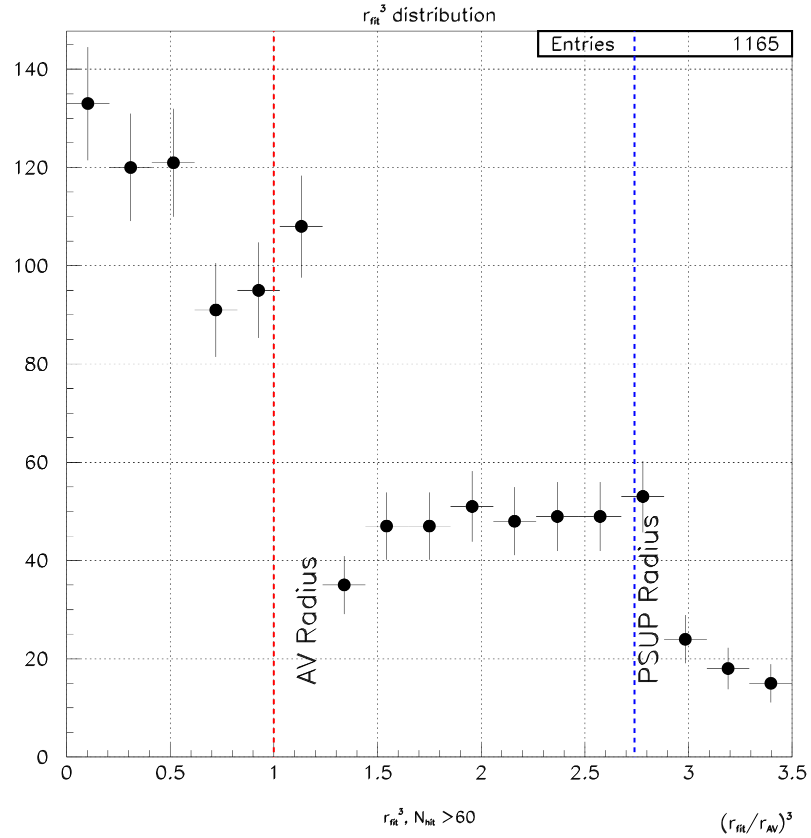

Figure 7.19: r^3 distribution with an additional cut $N_{\mathrm{Hit}} > 60$. Since radioactive backgrounds are at lower energies, this cut reduces the background contribution.

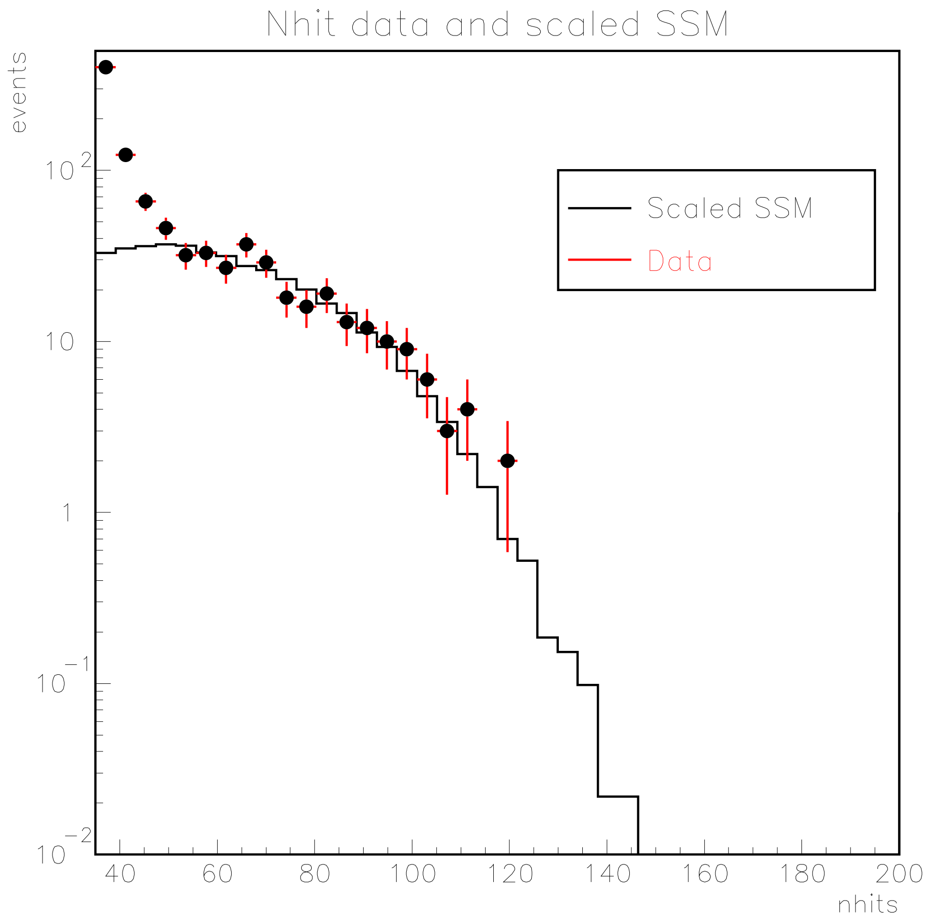

Figure 7.20: Overlay of data and MC N_{Hit} spectra. The MC spectrum was obtained by running 7 years of SSM data and scaling N_{Hit} to match the known discrepancy between MC and real ^{16}N data. The MC histogram was scaled to have the same area above N_{Hit} =55 as the data. The data is in the bars, the solid histogram is the MC. The fiducial volume cuts used for the N_{Hit} distribution is $r_{\text{fit}} < 540$ cm.

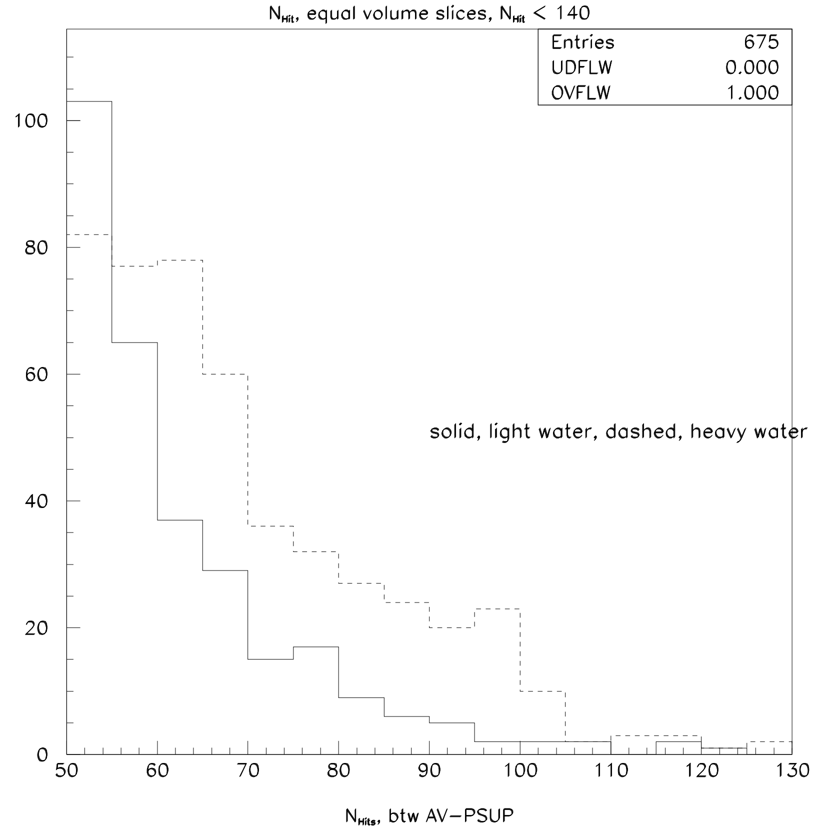

Figure 7.21: Comparison of equal volume slices in the central part of the detector *vs* a slice in the light water. The energy response in the light and heavy water is not the same, due to UV absorption in the panels of the AV.

reproduce the correct N_{Hit} behavior at the center for ^{16}N data. The overall vertical scale of the MC distribution was scaled to match the tail of the data, in order to facilitate a general shape comparison. We see that our data has good general shape compatibility with the MC prediction above 50 hits.

In Fig. 7.21, we can compare the N_{Hit} spectrum of two equal volume slices inside and outside the heavy water. Due to the energy response in the CC reaction and the absence of D_2O in the light water, we expect and see a shoulder at higher N_{Hit} in the

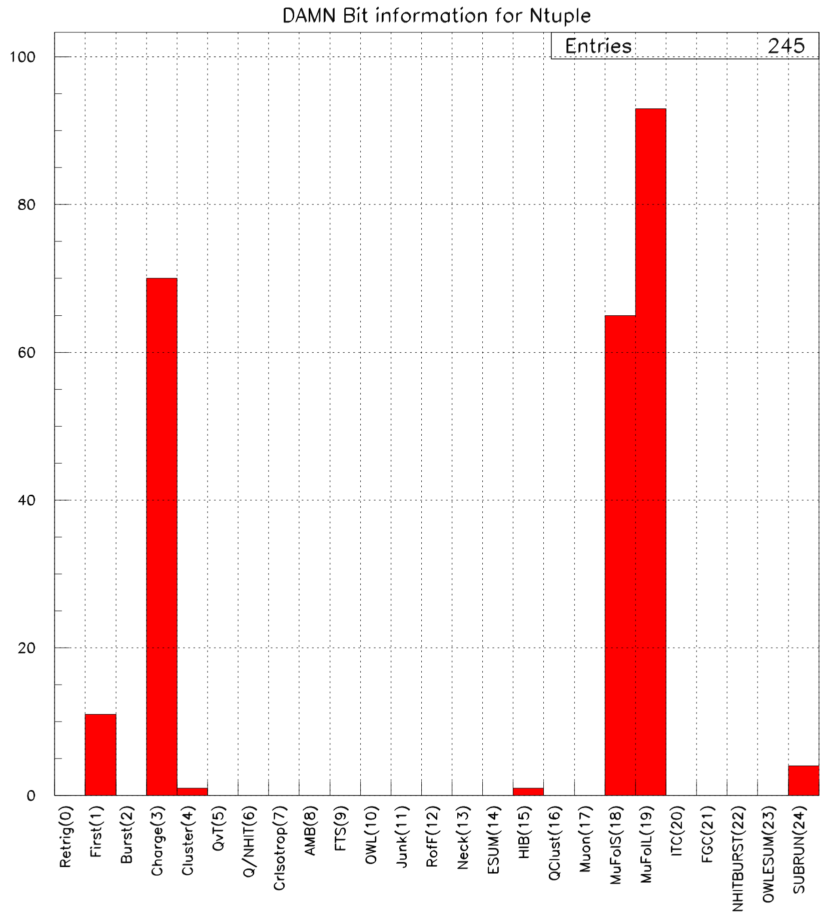

Figure 7.22: Remaining DAMN bits in the selected, reconstructed data.

heavy water plot. This is consistent with a neutrino signal in the D_2O.

7.1.4 Data sanity checks

Many other tests can be made to make sure the data is reasonable. A few are presented below. Fig. 7.22 shows a distribution of the remaining DAMN bits for the selected data sets. Since not all cuts were used in the selection mask, some events that were chosen were tagged by other cuts. The cuts that tagged the most events in our data sample are the charge and muon follower cuts. The presence of the charge cut might

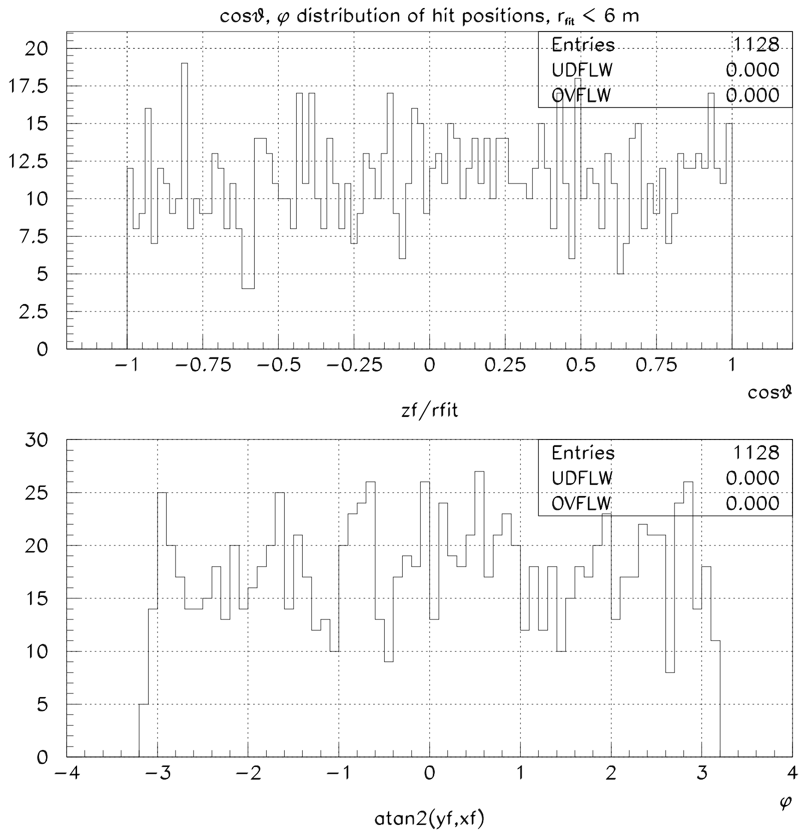

Figure 7.23: Distribution of reconstructed hits in $\cos\theta$ and ϕ.

indicate the existence of a bad board in the system. The muon follower cuts will tag a certain fraction of time, as they act to increase the dead time of the detector. The remaining tags do not present any indication of flaws in the data set.

Fig. 7.23 shows the distribution of hit positions in the D_2O in $\cos\theta = \frac{z_{\text{fit}}}{r_{\text{fit}}}$ and ϕ, for $\tan\phi = \frac{y_{\text{fit}}}{x_{\text{fit}}}$. Both distributions are flat, indicating that there aren't any obvious hot spots in the heavy water.

7.1.5 The data

We conclude from the preceding discussion of the $\cos\theta_{\odot}, r^3$, and N_{Hit} distributions that the cleaned data is consistent with neutrino data. The $\cos\theta_{\odot}$ distribution shows a clear ES peak, and the strength of the CC slope signal is consistent with the SNO's ability to measure it. In the high $N_{\mathrm{Hit}} \geq 60$ region where we will be doing our CC flux measurement, there is good indication that there are few backgrounds.

7.2 Charged Current flux

In this Section, we will present a normalized measurement of the CC flux ratio with the SNO. Due to the incomplete understanding of the SNO detector at this time, we will make extensive use of the Monte Carlo and also take ratio between data and Monte Carlo such that correlated errors will cancel. The ratio of the two numbers, the measured and predicted CC flux, gives us a model-dependent test of the solar model and of the neutrino oscillation hypothesis as a solution to the SNP. By taking the ratio, many of the associated systematic errors due to resolution and threshold cancel, making the measurement more robust despite the immature state of the detector calibrations and analysis. None the less, this is a preliminary measurement which is meant to point the way for more complete measurements to come, rather than an official SNO physics result.

7.2.1 Outline of Method

We will extract a charged current total number of counts from our data sample of 128.5 live days of running. We will then repeat the analysis identically with the

MC	$\mathcal{R}$	Run No.	Data	$r(cm)$
42.93	1.070	10490	45.93	0
42.87	1.071	10480	46.12	100
43.52	1.073	10482	46.68	200
42.72	1.080	10353	46.12	170
43.02	1.100	10357	47.30	300
43.84	1.087	10484	47.64	300
43.85	1.098	10486	48.15	400
43.29	1.104	10488	47.79	500

Table 7.2: ^{16}N Data/MC scalings at various positions. $\mathcal{R}$ is the ratio between data and MC, and r is the distance of the source from the center of the detector, in cm.

SNOMAN Monte Carlo for identical live time, with the null hypothesis of SSM and no new neutrino physics. The simulation consists of CC and ES interactions only, and we will estimate the number of neutrons (both from NC interactions and from deuteron photo-disintegration backgrounds) in the data by fitting.

Detailed procedure

We will now discuss in detail the procedure for extracting a CC rate. The cuts are similar to those that we used to present evidence for solar neutrinos in the first part of this chapter, but some differences exist.

In both the data and MC analyses, we apply an N_{Hit} threshold to the data, rather than an energy threshold as described in the previous chapter. However, the Monte Carlo has not been fully calibrated and therefore the Monte Carlo must be scaled in order to eliminate the observed differences between it and the data in the ^{16}N calibration. This scaling is based on comparisons between data and Monte Carlo at several locations in the detector. In Fig. 7.24, we see an example of this scaling. We use a final value which is volume-weighted and derived from the data presented in

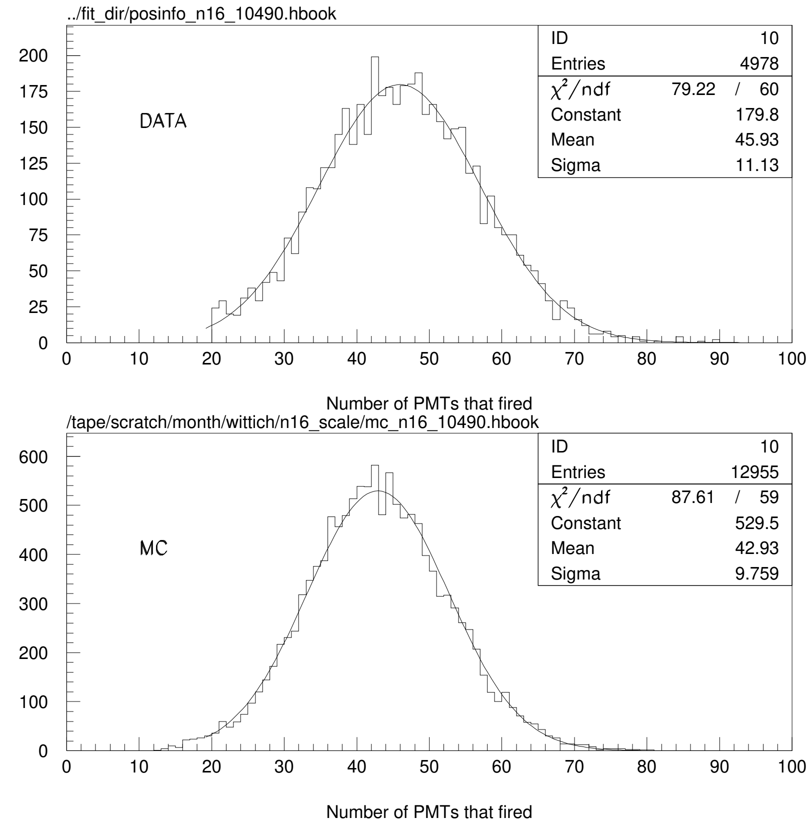

Figure 7.24: ^{16}N data-MC scaling.

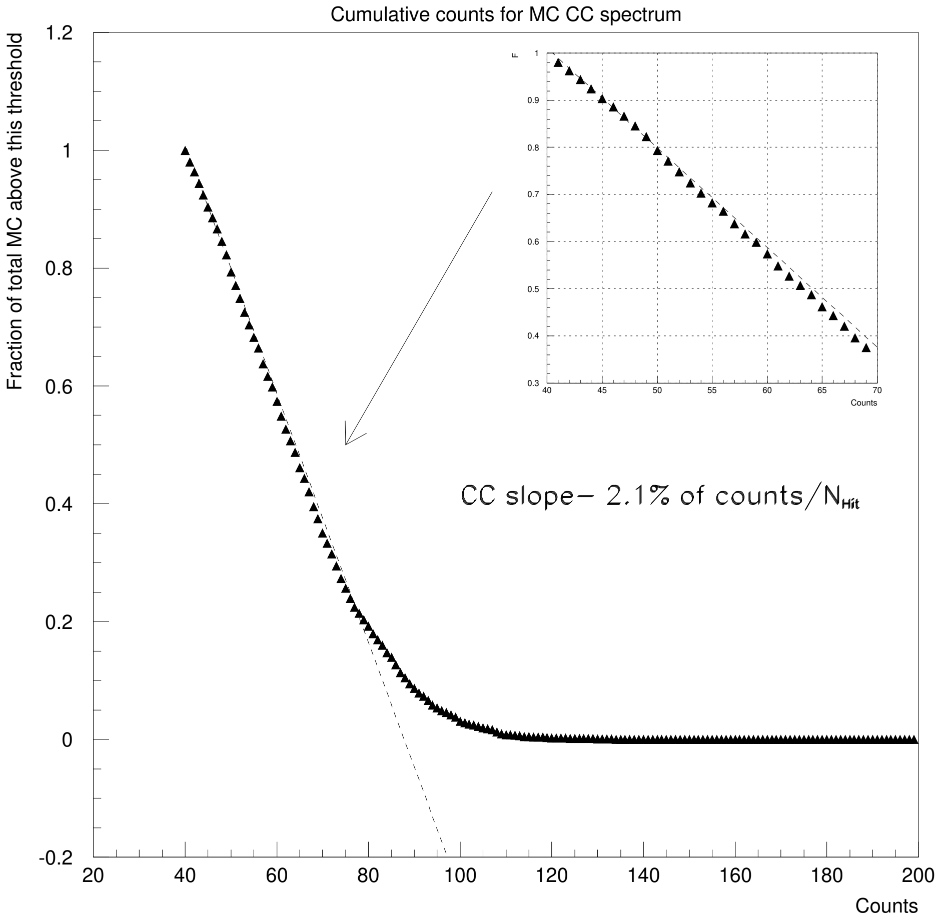

Figure 7.25: Fraction $\mathcal{F}$ of the spectrum above a threshold N_T. In the region around 60 hit PMTs, the total measured flux changes by 2.1% for each unit increase in N_{Hit} threshold.

Tab. 7.2.

$$\mathcal{R} = \frac{\mu_{data}}{\mu_{mc}} = 1.102 \pm 0.034 \tag{7.4}$$

where the error comes from the observed variation in the ratio of the N_{Hit} response to ^{16}N in data *vs.* simulation as a function of position. As discussed in Sec. 6.1.8, we apply an N_{Hit} threshold at $N_T > 60$ hits for the data and $N_T^{MC} = \mathcal{R}N_T = 54.5$ for the Monte Carlo. The error in this number will be discussed below.

The $N_{\rm Hit}$ threshold scaling is discrete, that is, it is non-sensical to apply a non-integer scaling to the data. We must understand the uncertainty associated with this discreteness. Fig. 7.25 shows the fractional rate of change of the integral flux as a function of threshold, that is

$$\mathcal{F}(N_T) = \int_{N_T}^{\infty} \frac{dN}{dN_{\rm Hit}} dN_{\rm Hit} \tag{7.5}$$

We see that near $N_{\rm Hit}$ =60, a one count change in the threshold changes the predicted MC total count by 2.1%. We therefore quote a 1.1% error due to the discreteness of the data-MC $N_{\rm Hit}$ scaling.

We reconstruct all events in our final data and MC samples with the FTE(G) fitter. We keep only events which the fitter successfully reconstructs and which reconstruct within 5.5 m of the center of the vessel. This fiducial volume cut reduces the number of deuteron targets by roughly 25%. However, this conservative fiducial volume cut insures that we are in a region where backgrounds are lower and the acceptance is well understood.

The $N_{\rm Hit}$ threshold we have chosen leaves us with an expected 6% neutron contribution from NC neutrons and general deuteron photo-disintegration background (Sec. 6.1.8). We choose to correct for this contribution by estimating its size based on the $N_{\rm Hit}$ distribution. We generate a MC shape of CC- and ES-induced electrons, in the proportions predicted by the SSM. We then scale this shape to match the data via a χ^2 fit at high threshold ($N_T > 70$), where we know the neutron contamination to be small. Fig. 7.26 shows an overlay of the data and scaled MC distributions. The difference between the integral count of these two distributions below $N_{\rm Hit}$ =70 provides an estimate of the neutron contamination in our signal data set. Based on

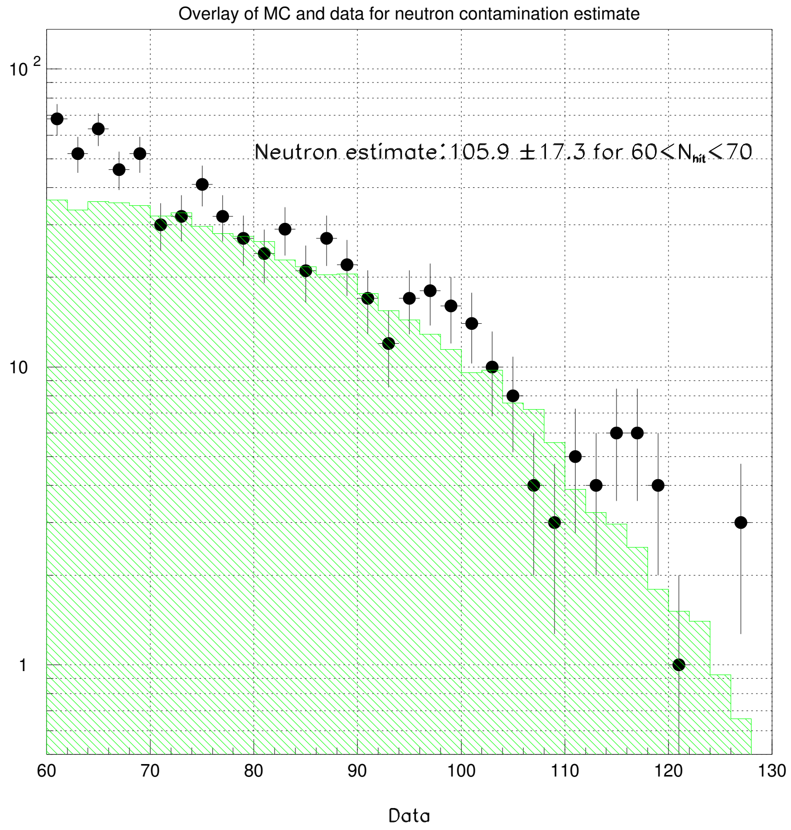

Figure 7.26: Fitted Neutron contribution to the $N_{\rm Hit}$ spectrum. The shaded spectrum is a simulation of $CC-ES$-induced electrons, and the points are the data. The usual cuts were applied to the data and simulation. We obtained the scale factor by a χ^2 fit of the data and simulation at $N_T > 70$, where the neutron contamination is known to be small. The estimated number of neutrons based on the spectral differences below 70 hits is 105.9 ± 17.3.

this fit, the neutron contamination is 105.9 ± 17.3. This number is clearly heavily dependent on the MC-data $N_{\rm Hit}$ scaling discussed above. We investigate the effect of varying this scaling and determine that the estimate of neutron background varies by roughly 10 counts. We take this to be a systematic error on the neutron contribution.

A MC simulation predicts roughly 70 counts above our current threshold for our live time, which we are in excess of. The quoted goal for radioactivity in SNO is 10% of the SSM of neutrons. While the detector is currently operating close to these limits, the measured excess in neutrons appears consistent with radioactive backgrounds.

Fig. 7.26 also shows an excess in the regime near $N_{\rm Hit} = 90$. This excess could be statistical, or it could be due to a high energy γ-ray backgrounds. Further study is required to determine its origin.

Next, we plot the $\cos\theta_\odot$ distribution and use it to separate the ES background from the CC signal. The functional form of the signals is based on a MC study. *A priori*, we expect the CC signal to have a $1 - \frac{\cos\theta}{3}$ angular distribution with respect to the sun, and the ES signal to be strongly forward peaked.

Figs. 7.27–7.28 show the angular function based on two MC studies. For the CC spectrum, we fit a first order polynomial and measure the slope to be $m_{cc} = -0.276 \pm 0.014$, as compare to the predicted $-\beta/3 \simeq -1/3$. The smearing is caused by the finite resolution of the detector and the fitter. The functional form for the ES is as described in Eq. 6.18.

Finally, we perform a fit to the $\cos\theta_\odot$ distribution based on the ES and CC shapes described above, i.e.,

$$\mathcal{F}(\cos\theta_\odot) = \alpha f_{es}(\cos\theta_\odot) + \beta g_{cc}(\cos\theta_\odot) \tag{7.6}$$

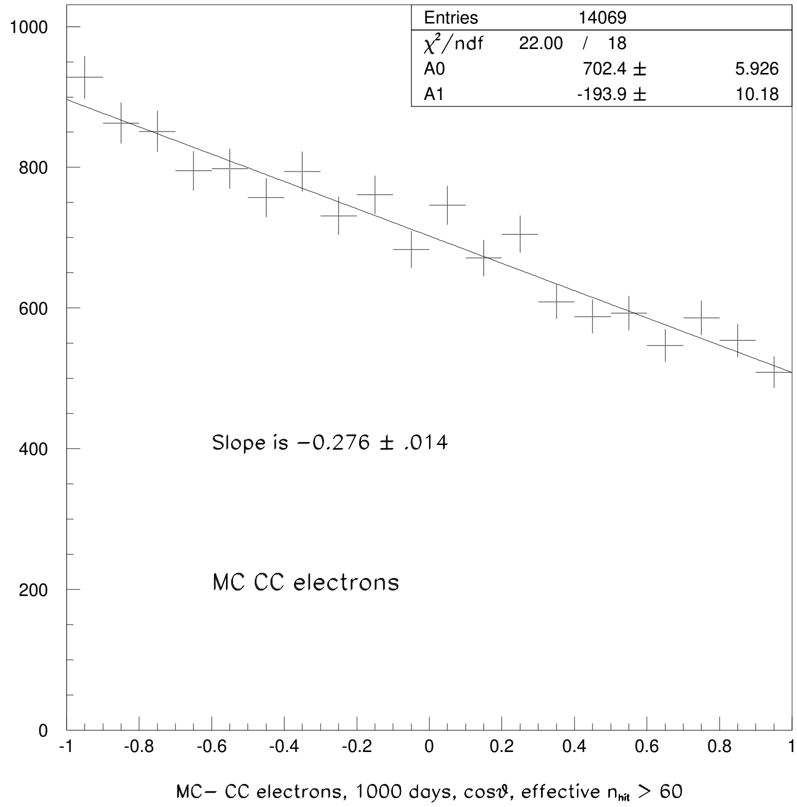

Figure 7.27: Angular resolution of the CC spectrum, based on a 1000 day equivalent CC Monte Carlo. We applied the same N_{Hit} and r_f cuts here as in the data.

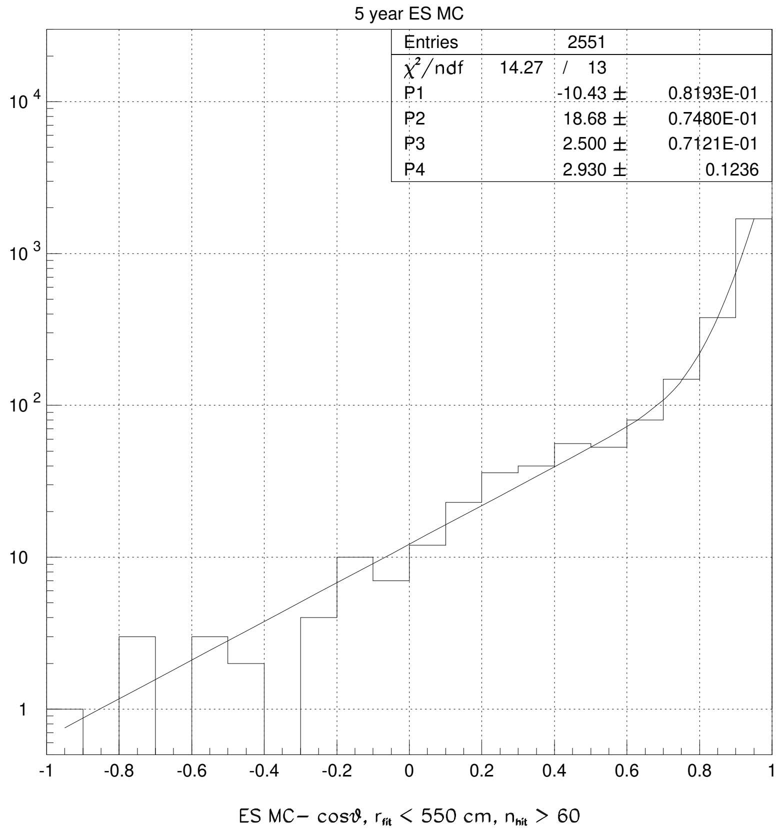

Figure 7.28: Angular resolution of the ES spectrum, based on a 5 year ES MC. We applied the same N_{Hit} and r_f cuts here as in the data.

where

$$
\begin{aligned}
f_{es}(\cos\theta_\odot) &= \exp\left(p_1 + p_2\cos\theta_\odot\right) + \exp\left(p_3 + p_4\cos\theta_\odot\right) && (7.7)\\
g_{cc} &= 1 + m_{cc}\cos\theta_\odot && (7.8)
\end{aligned}
$$

and fit for α and β. Integrating $\beta\ g_{cc}(\cos\theta_\odot)$ gives the CC flux.

Assumptions

This measurement is based on a set of assumption which we will list below.

- **No instrumental backgrounds.** We assume that above our current analysis threshold, the residual contamination of instrumental backgrounds is small. This claim is based on the signal data presented in the previous section. Analysis techniques to estimate these backgrounds are under development.

- **High energy γ-ray background is small.** Based on the estimate of this background in [46], we claim that high energy γ-rays contribute an upper limit of 14.1 ± 5.1 events in the signal region, and that this estimate is a good estimate of all physics backgrounds in our data set.

- **Sacrifice at high energies is small.** The sacrifice, that is, the amount of data that is cut by data cleaning, has been studied in great detail by the data cleaning groups for low energy events, i.e., events in the region at and below our analysis threshold. Currently, none of the calibration sources that would allow a study of this at the full CC energy have been deployed, and hence no measurement of the sacrifice at higher energies has been made. None the less, the measurements at the tails of the ^{16}N calibration data indicate that the sacrifice is not large at higher energies.

- **The excess at low N_{Hit} is due to neutrons.** This assumption is reasonable, based on the current understanding of what neutrons look like at SNO from the ^{16}N calibration sources. However, the shape of the other backgrounds is not yet well understood and could invalidate this assumption.

- **The effects of N_{Hit} scaling between data and Monte Carlo is understood.** It is unlikely that all differences between the data and the Monte Carlo are eliminated by the scale factor, nor is it obvious that this scaling doesn't affect the final result. For instance, fitters tend to be more effective at higher N_{Hit}. By applying a scale factor, we are effectively applying a lower N_{Hit} threshold to the MC than to the data. Therefore, more MC data occurs in a region where the fitter is less effective. Nonetheless, we assume that we understand the effects of this scaling.

7.2.2 Fits to Data and MC distributions

Fig. 7.29 shows the $\cos\theta_{\odot}$ distribution for the final data set, with all cuts applied. The parameters P1 and P2 show the fitted number of CC and ES events with statistical errors, respectively. Fig. 7.30 shows the same for a 1285 day (ten times our live time) simulation. We see that the ratio of data to simulation for our data set is

$$\frac{R^{\text{Data}}}{R^{MC}} = 0.34 \pm 0.02(\text{stat}) \tag{7.9}$$

where the quoted error is statistical only.

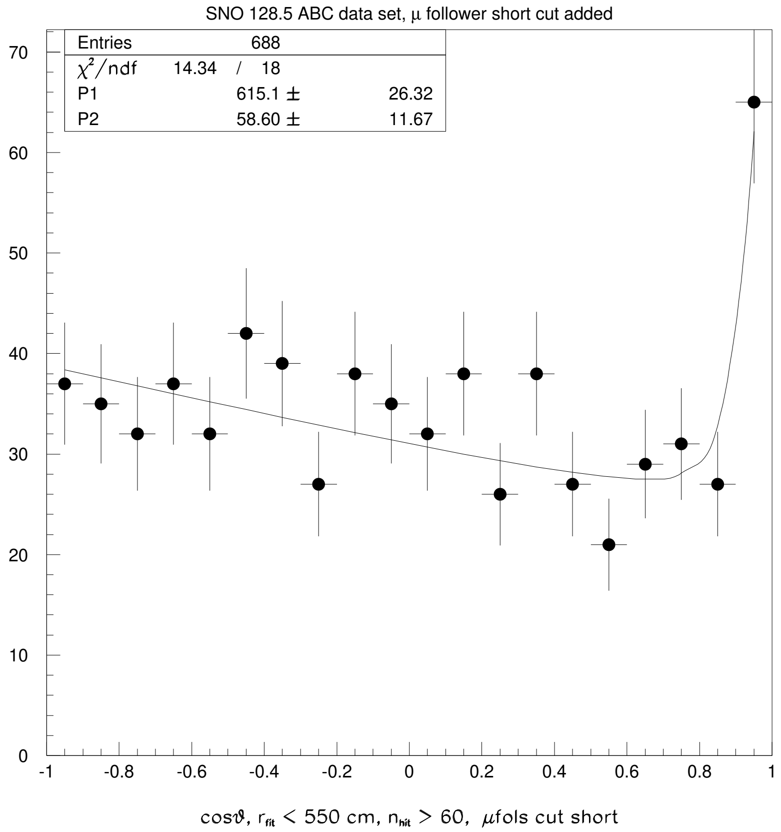

Figure 7.29: Final $\cos\theta_{\odot}$ for 128.5 live days of data. P1 is the fit number of CC events, with statistical error only.

Background	Counts
neutrons	105.9 ± 17.3
Final Count	
	509.2 ± 31.5

Table 7.3: Backgrounds subtracted from the CC data total counts.

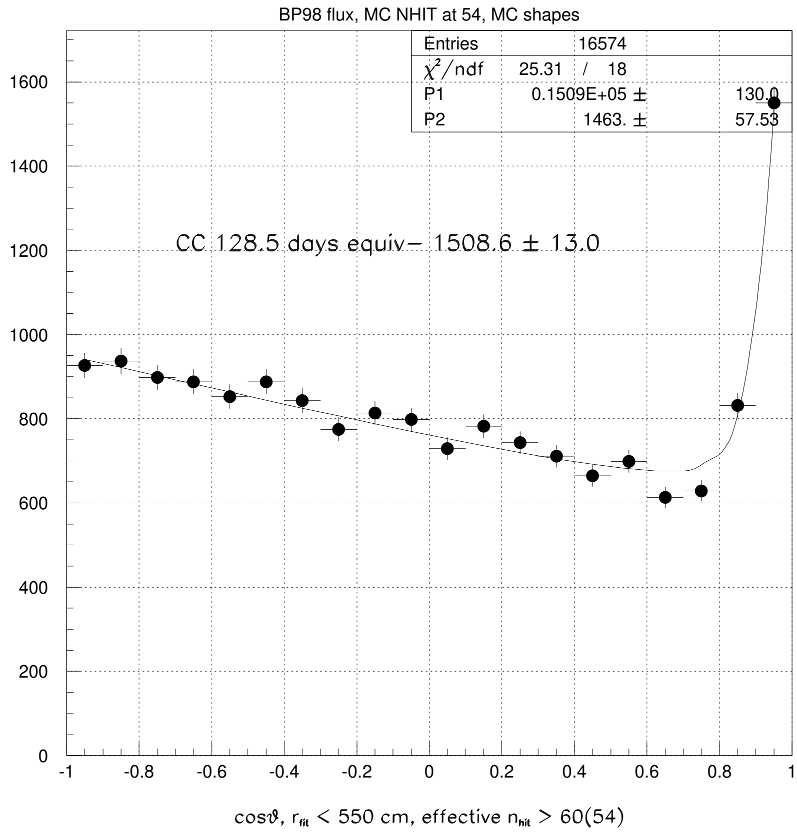

Figure 7.30: Monte Carlo $\cos\theta_\odot$ distribution for 1285 live days, for the null hypothesis of SSM flux, no new neutrino physics.

7.2.3 Systematic Errors

When estimating the systematic errors on this measurement, we have to consider the errors on the basic building blocks for the measurement. As we are taking the ratio of a measured *vs.* predicted rate, many of the errors approximately cancel. Therefore, we must consider in particular those errors that affect either the numerator or the denominator of our ratio, but not the other.

In order to calculate an integrated rate for the denominator of our ratio, we need to know the flux, the interaction cross sections, the number of targets, and the live time. We must consider each of these in turn, in addition to the errors that are due to this particular measurement.

Errors that only affect the simulation

The errors that affect only the simulation consist of the assumptions that are inputs to the model of the SNO detector. These include basic physics such as cross sections and more mundane topics such as the live-time.

SNOMAN uses neutrino – deuteron cross sections calculated by Kubodera and Nozawa [62]. Due to approximations made in the calculations, the authors cite an uncertainty of 5% for the $\nu - d$ cross section in the solar neutrino energy range. Note that for SNO's signature measurement, the ratio CC/NC, some of the uncertainties cancel and the ratio of cross sections is only uncertain to about 1% [62].

For the prediction of solar neutrino fluxes, we use the standard solar model of Bahcall *et al* [63]. The overall 1 σ uncertainties in the flux due to the SSM are $^{+19\%}_{-14\%}$ for the ^{8}B neutrinos.

The total number of targets N is a function of the isotopic purity of the heavy water, as well as the uncertainty in the fiducial volume. The isotopic purity is measured by the SNO water group to be 99.917 ± 0.005 [21]. This uncertainty is negligible in light of the ones already considered. The effects of fiducial volume errors due to fitter resolution approximately cancels in the ratio of the measurement and the prediction, as we apply the same analysis in both cases.

As mentioned above, the N_{Hit} scaling between data and Monte Carlo has associated with it a 3% uncertainty. As demonstrated in Fig. 7.25, for each one N_{Hit} change in the threshold, the measured rate changes by 2.1%. Given a threshold of 54.5 hits, this corresponds to a 3.5% systematic uncertainty in the total number of counts. In addition, this uncertainty is correlated to the uncertainty caused by the discrete nature of the energy threshold, and hence the total uncertainty in this scaling is increased by 1.1% to 4.6%.

Errors that only affect the data

These errors could, for instance, cause a reduction in the measured number of counts *ab initio*. Possible sources of error in this category include the SNO hardware trigger and the data cleaning.

Any inefficiencies in the trigger rate would cause a reduction in the amount of data that is recorded in SNO and hence would cause an underestimate of the total rate. For the data taken in this thesis, the trigger threshold for the main N_{Hit} trigger was in the range $N_T = 18 - 22$. The trigger efficiency at high thresholds where we are doing our analysis is 100% [64].

Events removed by the data cleaning will also cause a reduction in the rate. As discussed in 6.1.1, the data cleaning group uses triggered sources that produce Čerenkov light to estimate the amount of sacrifice, that is, the amount of data that is

Quantity	Size (%)
$\nu - d$ cross sections	5.0
Solar Model	$^{+19.0}_{-14.0}$
N_{Hit} Scaling	4.6
Number of targets	0.0
High E γs	2.3
n subtraction	1.8
Data Cleaning Sacrifice	0.5
Trigger Efficiency	0.0
Total Exp.	5.5

Table 7.4: Systematic errors for the CC rate ratio. The final line lists the total number excluding the cross section and solar model uncertainties.

cut by the data cleaning. The current best estimate on the sacrifice, as based on the ^{16}N source, is 0.5% [45]. However, the ^{16}N source only covers the lower energy range of the data used in this thesis and therefore the true sacrifice is currently unknown. The conservative estimate is to assume extrapolation of the current constant amount to higher energies.

Total Systematic error

Adding all the systematic errors listed in Tab. 7.4, the total error for our measurement is

$$\frac{R^{\text{Data}}}{R^{\text{MC}}} = 0.34 \pm 0.02(\text{stat}) \pm 0.02(\text{sys}) \pm 0.02(\text{cross section}) \begin{matrix} +0.06 \\ -0.05 \end{matrix} (\text{solar model}) \quad (7.10)$$

where we have split off the theoretical errors due to the solar model and the $\nu - d$ cross sections. The largest single error unique to this measurement is due to the N_{Hit} scaling.

Implications

What implications does a particular result for a CC rate from SNO have for solar neutrinos and neutrino oscillations in general? Recall that the current state of affairs is that solar neutrino experiments have measured a suppression of the solar neutrino flux ($\nu_e \rightarrow \nu_x$, where x can be active or sterile neutrinos), and that Super Kamiokande has evidence for neutrino mass based on the atmospheric neutrino results ($\nu_\mu \rightarrow \nu_\tau$, where transitions to ν_s are ruled out at the 90% CL [65]). More controversially, LSND has claimed to observe neutrino oscillations in the channel $\overline{\nu}_\mu \rightarrow \overline{\nu}_e$.

The experimental results must satisfy the relation

$$\sum \Delta m_{ij}^2 = 0. \tag{7.11}$$

The current ranges of Δm^2 as favored by the three experiments are, as shown in Fig. 2.5, $\Delta m^2 \sim 10^{-1}$ eV^2 for LSND, $\Delta m^2 \sim 10^{-3}$ eV^2 for atmospheric neutrinos, and $\Delta m^2 \sim 10^{-6}$ eV^2 for solar neutrinos in the case of the currently favored MSW solution. Since these numbers do not satisfy Eq. 7.11, we must introduce a fourth light neutrino. LEP limits rule out more than three light active neutrinos based on the invisible width of the Z boson [66], so the fourth neutrino must be non-interacting, that is, sterile.

Super Kamiokande has presented a ratio of measured *vs* predicted ^{8}B solar neutrino flux which indicates a strong deviation from the SSM prediction [67].

$$R = 0.465 \pm 0.005 \text{ (stat.)} \begin{array}{c} +0.015 \\ -0.013 \end{array} \text{(sys.)} \tag{7.12}$$

As Super K is a water Čerenkov detector which is only sensitive to the ES reaction for solar neutrinos, the ratio in Eq. 7.12 is a combined measurement of the ν_e and ν_μ, ν_τ fluxes at the earth. The ratio of ν_e to $\nu_{\mu,\tau}$ in this flux is unknown. Since SNO directly measures the electron neutrino flux from the sun, we can use the information from the SNO CC flux ratio to estimate the fraction of non-electron-neutrinos in the Super K ES flux. We shall discuss three possible scenarios.

If SNO measures a ratio that is identical to Super Kamiokande, we conclude that all neutrinos that Super Kamiokande sees are, in fact, electron neutrinos and the solar neutrino problem is due to $\nu_e \to \nu_s$. This conclusion would be confirmed by SNO measuring a CC/NC ratio that is one. While by itself this result would also appear to be consistent with the claim that the SSM flux prediction is wrong, this assertion has been theoretically ruled out when one includes the other solar neutrino experiments such as the Chlorine and gallium experiments [7, 8]. This result would also be consistent with the LSND result.

If SNO measures a ratio that is greater than the Super Kamiokande measurement, two possible causes spring to mind. The first explanation is that the Super K result is wrong. This appears unlikely, though not impossible. Alternatively, a more likely explanation is that the theoretical $\nu_e + d$ interaction cross section calculations are incorrect, leading, for instance, to an underestimate of the expected rate in the SNO detector. However, the experimental cross section for $\overline{\nu}_e + d \to e^+n + n$ has been measured experimentally to about 20%, in a measurement whose central value agrees well with the Kubodera-Nozawa predicted value [68]. This lends confidence in the calculated cross sections.

If SNO measures a number that is less than the Super Kamiokande ratio, we can conclude that the difference between the SNO and Super K numbers come from the solar neutrinos that have oscillated to ν_μ and ν_τ. In this case, the CC/NC ratio would be less than one, indicating oscillation into active neutrinos explains the solar neutrino problem. This result would be incompatible with the LSND result, as the values of δm^2 in SNO that would produce this result are much smaller than the LSND measurements ($\delta m^2 \sim 10^{-6}$ eV^2).

The results quoted here seem to prefer this interpretation.

7.3 Conclusion

In this thesis, we have described the work done on the development of the SNO detector, from construction to a first analysis of SNO data. With the CC/NC ratio, SNO has an exciting opportunity to test the neutrino oscillation hypothesis as a solution to the solar neutrino problem in a model-independent fashion. In addition, with the CC reaction, SNO has the ability to measure any distortions of the ^{8}B spectrum due to possible MSW effects.

The work done in this thesis has spanned the initial prototyping of the electronics for the SNO PMT readout, to their production testing and final system installation. In addition, work was done to help design and implement a complete data acquisition software system for electronics testing and commissioning, and a data flow system to insure the validity of the data in its movements to long term storage.

Since the turn on of SNO, the focus changed to understanding the detector, concentrating on instrumental backgrounds and first signal extraction. Now, in this thesis, we present some analysis illustrating that SNO has data that is consistent with a ^{8}B-ν_e induced signal from the Sun. This evidence was presented, and a preliminary version of a CC flux ratio measurement was presented as a road map for future, more complete analyses that will be performed by the collaboration when the detector calibrations have been completed.

Appendix A

Sequencer State machine

In this appendix and the following ones, we describe details of some of the electronics I have have worked on during my tenure at Penn. The first appendix describes the details of one of the Xilinx programs on the FE card. We introduced the sequencer and described its function in the data flow in Ch. 5. The details are recorded here.

A.1 Purpose

The sequencer is a state machine used on the front end cards of the SNO electronics. Its function is to move the data from the analog store arrays on the SNO QUSN7 CMOS chip[1] through digitization and into the 4 MB SIMMs of memory on each of the FECs, where it can then be read out by the DAQ computers. The sequencer acts independently of the DAQ computers. In combination with the analog storage arrays on the QUSN7 CMOS chip and the FIFO controller (see Appendix B), the sequencer is the motor of the SNO front end card.

[1]For a complete description of the CMOS chip, see [69]

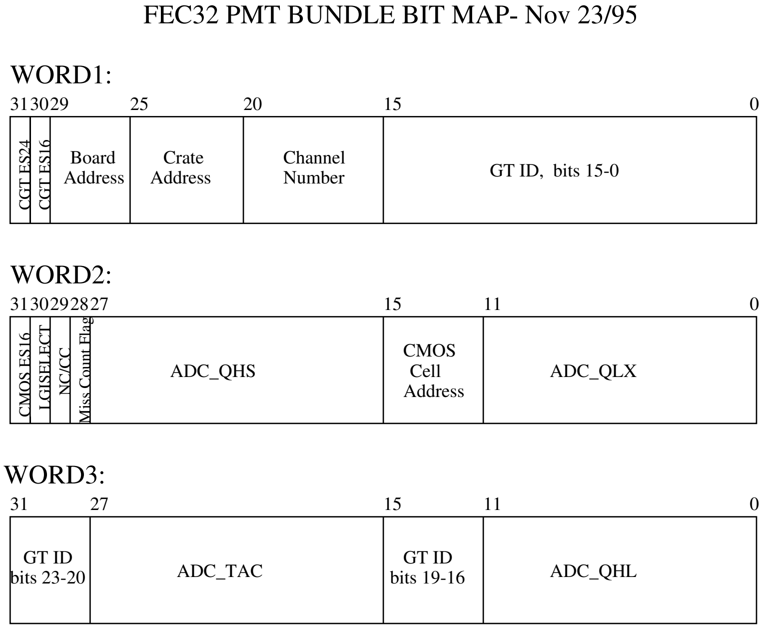

Figure A.1: SNO data bitmap for PMT data

A.2 Physical Implementation

The sequencer is implemented in a Field Programmable Gate Array (FPGA) model XC3060 made by the Xilinx Corporation. The FPGA is configured by the DAQ computers after power up of the FECs, which allows dynamic reconfiguration of the chips during running for testing purposes. This feature was immensely useful during FEC verification and commissioning. The sequencer is referenced to a 16 MHz clock, which determines the unit of time on which it operates.

A.3 Description of functionality

The sequencer's job is to move PMT data to the FEC memory so that it can be read out by the DAQ computers. The bulk of the bits are the digitized PMT information

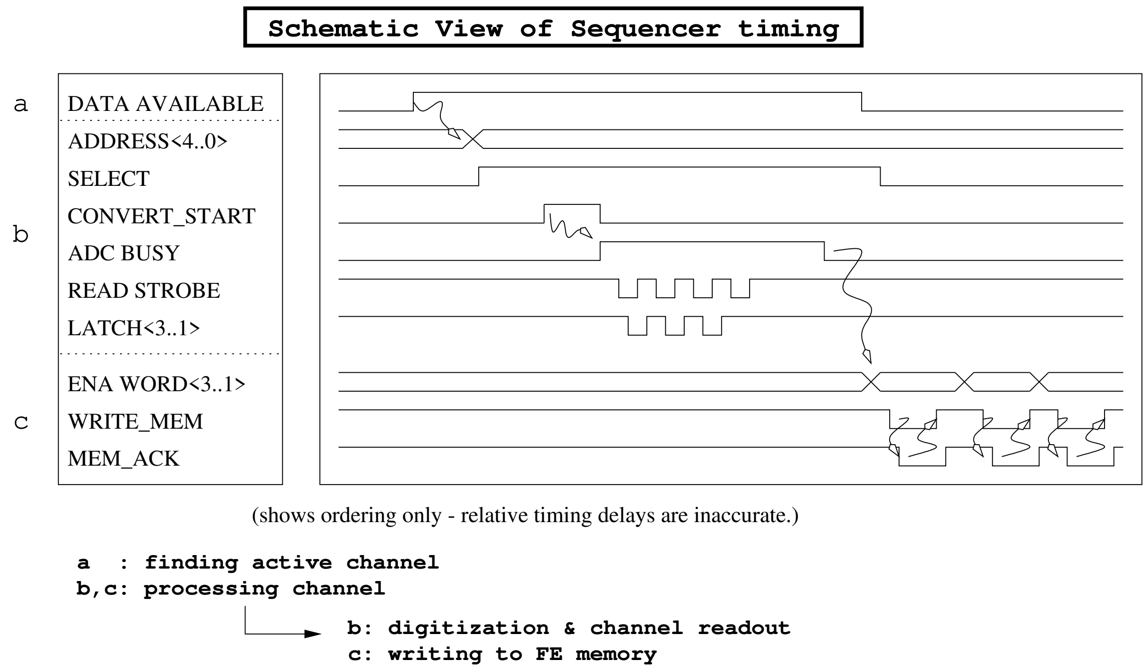

Figure A.2: Schematic description of sequencer timing

(ADC_TAC for the time-to-amplitude conversion, and ADC_QHL, ADC_QHS and ADC_QLX for the three charge conversions). In addition, a number of status words are added to the data. This includes the status of several test registers, and the global event identifier, which is needed during the event building phase. As all PMT bundles arrive at the event builder asynchronously, a tag is required to keep track of which event each PMT bundle belongs to. This tag is a counter on each channel which records the number of global trigger (GT) signals seen by that channel, and is called the global trigger identifier (GTID).

The above information is collated into three 32 bit data words, which together make up the PMT bundle shown in Fig. A.1. These three words are written into the local FEC memory.

The PMT bundle also includes the geographic information of the relevant crate,

ADC data	
TAC	digitized time-to-amplitude conversion
Q_{HS}	high gain, short integrate charge conversion
Q_{HS}	high gain, long integrate charge conversion
Q_{LX}	low gain, adjustable time length charge conversion
Identifiers	
GTID	unique event identifier
Board address	originating channel's board address
Crate address	originating channel's crate address
Channel address	originating channel address
Cell address	CMOS capacitor number used to store channel information, 0-15
Error bits	
CMOS ES16	synchronization error in the sixteen bit rollover of the GTID counter
CGT ES16/24	synchronization error in the 16/24 bit rollover of the CGT counter
Missed Count	error bit indicating the channel has missed data
Status bits	
LGISELECT	bit indicating whether the Q_{LX} low gain integration integrates over long or short times.
NC/CC	unused.

Table A.1: PMT data bundle definitions

card and channel where the bundle originated. This allows the identification of the PMT which fired. Additionally, the CMOS cell number used in storing the event is recorded. The analog storage array in the QUSN7 chip has sixteen separate storage capacitors. Each of these capacitors has a unique offset and must be individually calibrated. An identifier is included which shows which of the capacitors was used when the tube fired in order to allow the proper calibration to be applied. It is labeled 'CMOS cell address' in Tab. A.1.

The 'missed count' bit indicates a channel–level loss of data. In the case that the internal buffers of the CMOS chip are full when a discriminator firing coincides with a GT signal, the CMOS chip cannot store this new data – it is lost. When this occurs, an internal 'missed count' counter is incremented and the 'missed count' bit is set in the data stream to alert the DAQ computers immediately and the physicist doing analysis at a later date to this loss of data. The rest of the information includes error bits (CGTES16, CGTES24, CMOSES16) and status bits (LGISELECT, NC/CC) and is described in the PMT bundle description in Tab. A.1.

The sequencer is logically split into three different blocks at the top level. One block looks for new data to process, while the second processes the channel with new data. The final block is a control block to provide auxiliary functions. These parts will be described in detail below. A cartoon timing diagram to accompany the text is included in Fig. A.2.

A.3.1 Finding active channels

When a channel has valid data to be stored, a signal called DATA_AVAILABLE is sent from the QUSN7 to the sequencer indicating that that chip has data to be digitized.

The sequencer monitors each of the 32 channel DATA_AVAILABLE flags, using a priority–encoding scheme to decide which channel to process first. The priority–encoding scheme currently used is very simple and is also *unfair*. Priority is based on channel number - the highest channel number is serviced first, with others coming in numerical order. The scheme is unfair because it does not take into account whether a channel has been waiting longer than other channels when applying the priority scheme; rather, it applies the priority scheme strictly. For instance, if channel 15 fires, followed by channel 3, followed again by channel 10, the order of processing[2] will be in descending numerical order: first channel 15, followed by channel 10, followed by channel three. This can lead to pathological situations where high-rate channels could prevent others from being read out, leading to a loss of channel data when the analog storage arrays of the neglected channels fill up. By monitoring the data stream for high rate channels (i.e., channels that take part in every event), the DAQ computers can detect these pathologies and act to minimize their negative effects. One possible plan of action is to disable the channel's readout in the sequencer.

If a misbehaving channel is found, its DATA_AVAILABLE signal can be masked out via a setting in one of the sequencer control registers. The mask hides the DATA_AVAILABLE channel from the sequencer's active channel finding routine, effectively disabling it by preventing its data from getting read out. However, its front end chips, the integrator and discriminator, will still be active, and in the interest of minimizing possible channel-to-channel crosstalk, further steps can be taken when disabling a channel to silence its front end.

[2] I am assuming that the time progression is sufficiently rapid such that channels 3 and 10 fire before the sequencer finishes processing channel 15.

A.3.2 Processing an active channel

Once a chip has data available, control is passed to the next top-level block which processes the chip. Here, the analog data is read out, and the PMT bundle is assembled and written into the local DRAM memory. This block is made up of three sub blocks.

1. Digitizing data
2. Writing data to memory
3. reading out CMOS status registers

While the first two are used during normal data taking, the third is used for diagnostic purposes.

Digitizing the data

At this point, the channel in question is activated by asserting its SELECT line. On selection, the CMOS chip gates its four stored analog values onto signal lines to the analog to digital converters (ADCs). At this point, a 1 μs wait is imposed to allow the output to settle on the ADC inputs, to account for the slew rate of the intermediate op-amps. After this time has passed, the sequencer asserts the ADC's CONVERT_START signal, the ADCs latch the signals at their inputs and start the digitization, a process which takes 14 clock cycles of the ADC's 8 MHz clock.[3]

Once the digitization is underway, the sequencer extracts the digital information pertaining to this PMT bundle from the CMOS chip. The chip stores its digital

[3]The MAX120 ADCs are sample-and-hold devices. Once digitization starts, it is impervious to the digital noise caused by sequencer processing. This was verified during the FEC design and debugging.

data in internal registers, which are accessible via an 8 bit data bus and can be addressed via five address bits. When the CMOS chip receives a READ_STROBE signal, data from the register selected by the address bits is presented on this external data bus. If the internal register is more than eight bits wide, successive blocks are presented on following strobes. During data taking, a register at address 0x0 called 'Auto-increment' is selected. In 'Auto-increment', three eight bit words are extracted from the chip via three READ_STROBEs and latched into three external eight bit registers. A fourth READ_STROBE signal indicates to the CMOS chip that the data at the current location has been read out and that it can mark the present memory locations and capacitors as unused. In a similar fashion, the channel's geographic crate, slot, and channel information is stored in intermediate registers. The crate- and board-level geographic data is determined by jumpers on the backplane, while the channel-level data is just the encoded information used when finding an active channel, as described in A.3.1.

At this point, the sequencer waits for the ADC_BUSY signal from the AD converters to be de-asserted, indicating that the data has been converted. At this stage, the data can be written into the FE memory and control is passed to the next section.

Writing the data to memory

The PMT bundle data is now ready to be written into the FE memory. This is in itself a simple operation, but a few complications exist. The memory is set up to look like a dual-ported memory, i.e., two places can request access to it at the same time. In our case, the two players are the sequencer writing data into memory, and the DAQ computers reading data out of memory. The arbitration between these two requests is handled by a commercial dual-ported memory controller.

The next complication is the book-keeping of which memory location to write to. In order to simplify the work for the DAQ and the sequencer, the pointers into memory of the most recently read memory location, the most recently written memory location, and the amount of free store is managed by another FPGA called the FIFO controller. It mimics a 'first in, first out' stack arrangement where the sequencer puts data onto the top of the stack, while the DAQ removes data from the bottom of the stack. It is described in detail in Appendix B. The main advantage of this method is that neither the DAQ nor the sequencer need to know which memory locations they are writing to - this is taken care of transparently by the FIFO controller.

From the sequencer's point of view, the data is written into memory in three writes of 32 bits each. A request is made to the memory controller, and the memory controller arbitrates access attempts between the front end and the DAQ. The memory requests themselves are gated by the FIFO controller. In the unlikely case that the memory is filled, the sequencer will hang until memory becomes available. The FE can continue to acquire data during this time until its CMOS analog storage arrays fill up, however, so no data is lost. At this point, the cycle can return to the beginning, Sec. A.3.1.

Reading out CMOS status registers

In addition to the data described above, the CMOS chip holds some status information that can be used for diagnostic purposes. For instance, a count of the total number of discriminator firings is stored on the CMOS chip. By reading this counter out at fixed intervals, the DAQ can measure the channel's noise rate.

The data in these registers is accessed in the same way as described in Sec. A.3.2. Access to the normal data taking register, the auto-increment register, is forbidden in hardware in this mode, as the data thus read out would not be stored in FEC memory and hence lost.

There is a possible conflict between the action of reading out a CMOS status register and normal data taking processes described above. Preference is given to normal data taking, and in case of collision, the sequencer sets an error bit in the top bit of the output word of the register read, indicating that an error has occurred. It is up to the DAQ to test for this error bit.

A.3.3 Control functionality

The control functionality described above is managed through a set of 32 bit registers. These registers allow the disabling of undesirable channels by masking out their DATA_AVAILABLE signal from normal processing. Read-only access to the current snapshot of the DATA_AVAILABLE flag for all channels is also granted.

A.3.4 Test functionality

Several test modes exist to allow software emulation of the sequencer, or access to a test DAC which feeds into the ADCs to test their performance. These programs were most often used during development and production testing of the boards. Currently, the only alternate program in regular use is the TESTMODE2 program.

Test Mode 2, a.k.a. CALDAC

When the sequencer is programmed in test mode 2, it is a passive device. All outputs, such as the lines that interact with the ADC, can be toggled under computer control.

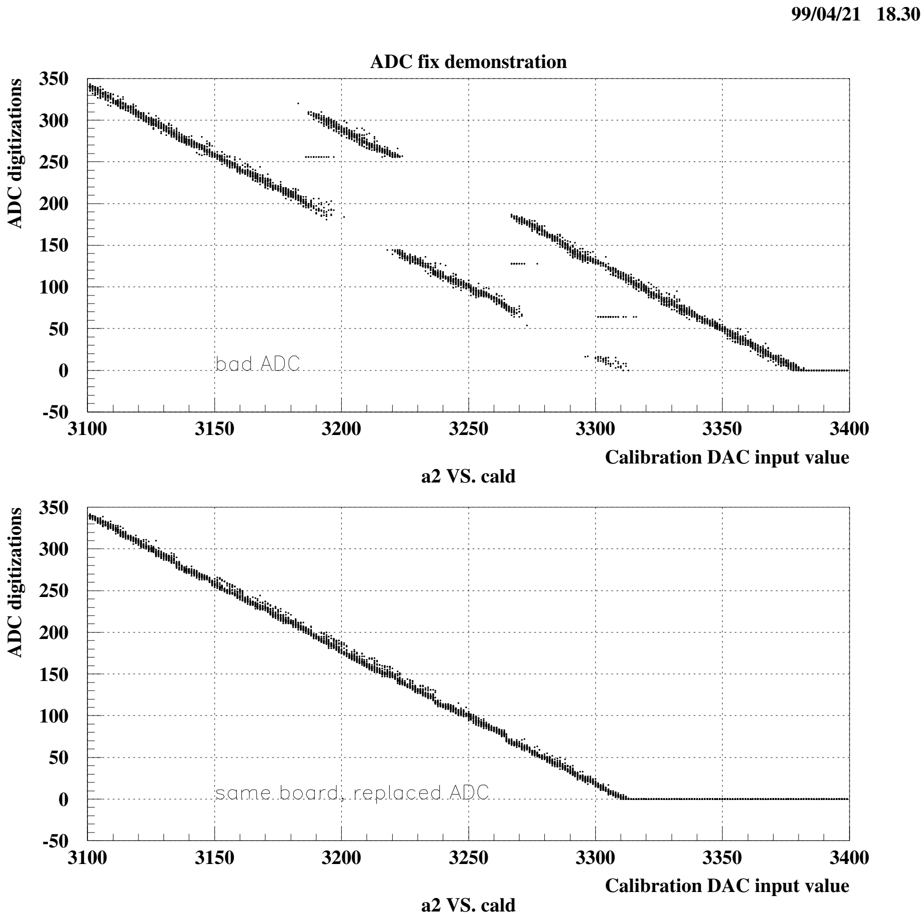

Figure A.3: A typical example of the non-linear response of the MAXIM 120 12 bit ADCs.

In addition, if the mode is turned on in the GENCSR,[4] a 12 bit calibration DAC on the FEC is connected to the ADC inputs. This test mode allows testing of the linearity of the commercial ADCs. A failure mode has appeared in these commercial ADCs where the chip produces discontinuous jumps in the output of the ADCs. Initially, the problems appeared in ill-behaved TAC digitizations. We unexpectedly traced it back to the commercial ADCs using the CALDAC program of the sequencer. This

[4]General Control and Status register.

behavior appears to be caused by faults in the commercial manufacturing process of the ADCs. A typical example of the failure mode is shown in Fig. A.3. The CALDAC program remains the single best diagnostic for this problem.

Appendix B

The FIFO controller

B.1 Functionality and Purpose

The FIFO controller is a device which is used in conjunction with the commercial dual-ported memory controller and commercial DRAM memory to take care of all memory addressing during normal data taking. The function of the FIFO controller is to minimize the DAQ programming task and reduce the number of bus accesses in reading and writing to memory. In addition, it prevents the DAQ computers from reading out undefined memory areas and prevents the sequencer (see Appendix A) from overwriting memory that has not been read out.

B.2 SNO FEC memory model

The memory model of the SNO front end card is best thought of as a simple first-in, first-out stack. While data is continuously added to the top of the stack from the PMTs by the sequencer, it is simultaneously being removed from the bottom of the stack by the DAQ computers. The job separation is strict: the sequencer is a *writing client*, which never removes data from the stack. The DAQ is a *reading client*, which never puts data on the stack. The act of deleting and inserting items from the stack is independent as long as there is data to be removed, and as long as there is enough room in the stack to accommodate new entries.

While commercial FIFOs which fill these needs are available, they are expensive. The affordable ones are small in size, and not intended for large-scale systems where the stack size can be in the hundreds of thousands of entries long. Affordable memory has a simple random-access memory model, where the location and type of action (e.g., read or write) must be specified for each access.

We split the implementation of the stack-like memory model into two parts: dual memory access arbitration and pointer management. In order to fill the first of these requirements and allow two sets of accesses, a commercial dual-ported memory controller (Phillips F1765) was purchased. It arbitrates access for up to 4 MB of memory from two clients, in order to prevent collisions. To fulfill the other requirements, some logic, either in hardware or software, is required to manage several pieces of information.

- memory location last written to by the writing client
- memory location last read from by the reading client
- separation between these two pointers
  - for the writing client, to prevent overwriting valid data.
  - for the reading client, to prevent reading out invalid data.

Additionally, care must be taken such that the reading client does not write into memory, incorrectly marking memory as containing valid data, and vice versa.

The FIFO controller attempts to take care of these issues and hide the details from the reading and writing client. It manages the relevant pointers and forbids the incorrect accesses.

B.3 Implementation

The implementation of the memory model is constrained by the tools that were available when this work was done. Modern FPGA programming allows different methods, from the very high level procedural descriptions, to truth tables, to more traditional gate-level designs. When the FPGAs for SNO were implemented, we did not have access to any of the higher–level design tools. Therefore, all designs were implemented using gate–level description. While this makes the implementation of the design harder, it is not a fundamental limitation.

B.3.1 Main Functionality

The internal representation of the memory used in the FIFO controller is a linear array, as shown in Fig. B.1. Four MB of memory allow one megaword[1] of memory locations, equivalent to $N_{\mathrm{Total}} = 2^{20}$ memory locations to be addressed. The controller has three separate counters to keep track of the read pointer, write pointer, and difference counter respectively. As the names indicate, the read, write and difference pointers indicate the last location read by the reading client, the last location written

[1] One word corresponds to four bytes and to 32 bits.

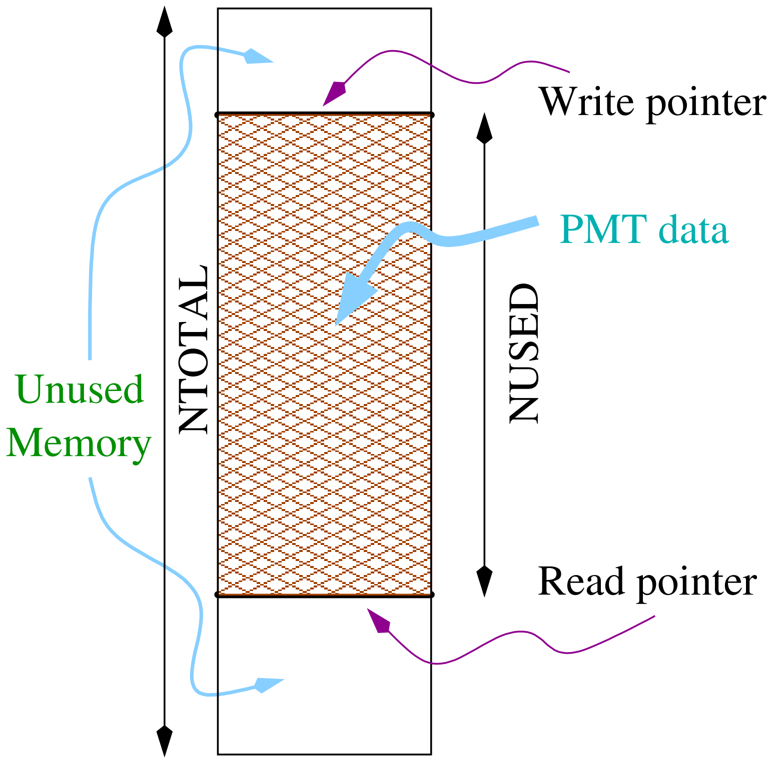

Figure B.1: FIFO memory layout

to by the writing client, and the total number of free locations available for writing.[2], [3] A more efficient implementation might have only kept track of the read and write pointers, and calculated the amount of free space based on the distance between these two pointers, but the number of distinct cases that would need to be handled was sufficiently large that we decided that efficiency would be sacrificed for simplicity of implementation.

The read pointer is a unidirectional twenty bit counter that is incremented on each read access. In case of a read lockout, a read–inhibit signal prevents the counter from incrementing. The inhibit condition is determined on the previous access cycle, so there are no race conditions between the access and the inhibit signal. The write pointer is implemented identically, with the appropriate substitutions of 'write' for 'read'.

The inhibit signals are determined by the difference counter. The counter is implemented as a twenty bit up–down counter, with the value incrementing on an access from the write client, and decrementing on an access from the read client. While the read and write pointers keep track of locations in memory where data is stored, the difference counter only keeps track of the amount of free store and determines when to set the read and write inhibit signals. When this counter reads zero, there is no valid data to be read out and the read–inhibit signal is set. When the counter reaches `0xfffff`,[4] the free store has been exhausted and no accesses by the writing client will be allowed.

If an access attempt is made during a lockout period, the access will hang up

[2]The FIFO controller counts memory locations in long words, not in the natural unit of the sequencer, the PMT bundle. Each PMT bundle corresponds to three long words.

[3]The number of free locations N_{Free} corresponds trivially to the number of available locations for reading, $N_{\mathrm{free}} = N_{\mathrm{Total}} - N_{\mathrm{used}}$.

[4]i.e., $2^{20} - 1$ in hexadecimal notation

indefinitely. The reading client will receive a time–out after a period of time, while the writing client waits until it can continue.

In the normal case of an uninhibited read or write access, the counters for read and write access are incremented by one and gated onto the external address bus which leads to the memory controller to serve as the address for the memory access. An access request is now allowed to be sent to the memory controller. Upon successful completion of the memory access, the client de-asserts its request, which indicates to the FIFO controller to release the address lines, having completed its task.

The inverse of the read–inhibit signal, i.e., read–enable, is made available via the backplane to a register on the XL2 interface board. By querying one register per crate, the DAQ computers can determine which, if any, FECs have data to be read out.

B.3.2 Diagnostic functionality

The current value of the read and write pointers and the value of the difference counter can be accessed in a read-only fashion from the DAQ. This allows the DAQ to decide how much memory to read out, rather than reading until it receives a time–out. The rate of increment of the read pointers gives a board-level diagnostic of the rate of coincidence of the PMT singles rates with SNO triggers. Loading of the counters is not implemented, and any attempts to write to the read-only pointer values will seem to succeed but have no effect, that is, they will not time out but the values will remain unchanged. One can reset the chip completely, setting the read, write, and difference pointers to zero, by toggling the Xilinx RESET line. All data in memory is effectively discarded by this action.

In addition to the normal functionality, all Xilinx chips have a power-down mode in which the chip suspends its operation, but retains its internal program and state. This feature is used on the SNO FE card to allow direct accesses to the memory, bypassing the FIFO controller. In this mode, full read/write access is given to both the reading and the writing client.[5] We use this mode for testing the FEC memory for faults. There is an unfortunate side effect to this mode that when the chip comes out of the power-down mode, it asserts a programming signal to indicate that it is coming back on-line. This signal causes all other FPGAs to expect reprogramming, and renders them unusable. A full reprogram of all FEC Xilinx chips is required at this point.

[5]What the sequencer does in this situation is undefined.

Appendix C

Timing of the SNO front end

In Ch. 5, we described the SNO front end and its three custom ICs. These high-performance chips include a discriminator and measurements of the time and charge of PMT signals.[1] In order to get maximum performance out of this chipset, each channel's chips must be individually set up. Of central importance in this setup procedure is the timing of the discriminator, which controls the integrator, and the co–ordination between the discriminator and the TAC. The timing in the SNO chips

[1]In addition to Ch 5, see [25] for a description of the discriminator-integrator-timer chip set and [69] for the documentation for the TAC and analog storage array.

time	ns
T_{TOT}	420
$T_{\mathrm{Q_{HS}}}$	70
$T_{\mathrm{Q_{HL}}}$	$330 + T_{\mathrm{Q_{HS}}}$
T_{deadtime}	20
T_{GTVALID}	400
Width of TR_{100}	93
TAC nonlinear regime	60

Table C.1: Representative values for SNO FE timing

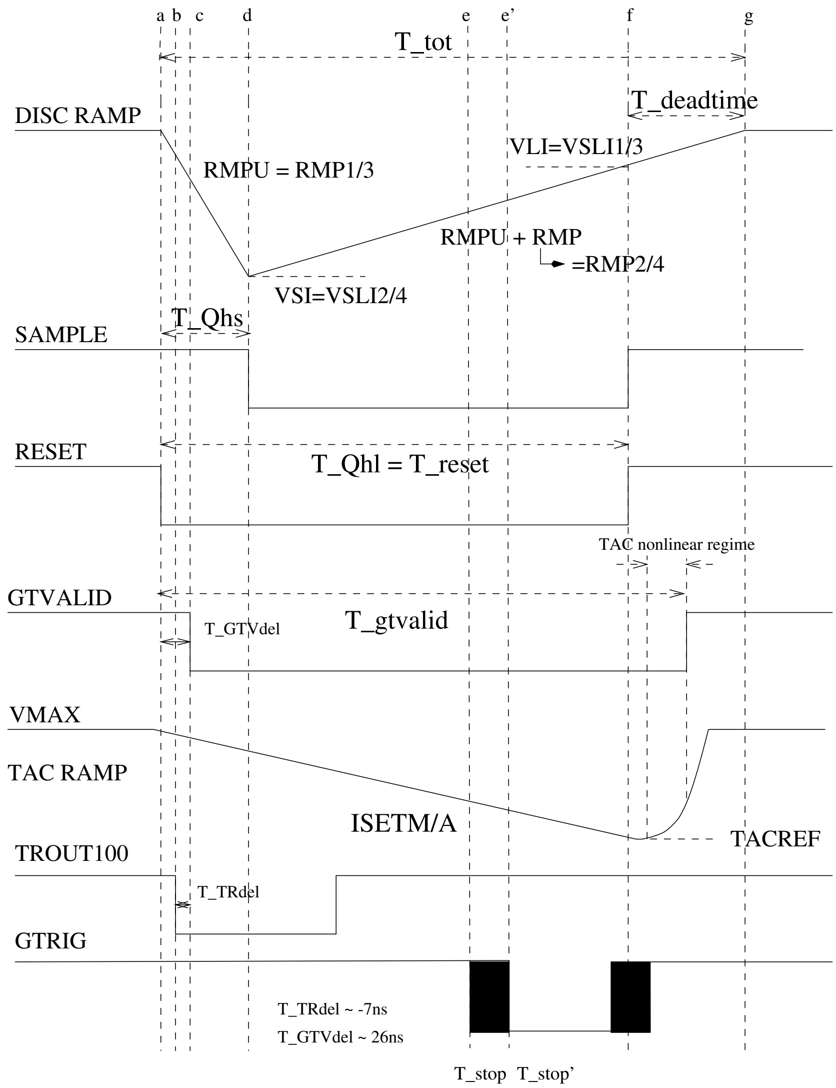

The tube displayed causes TROUT100 sum to cross the GTRIG threshold, causing a GTRIG.
Drawing *not* to scale.

Figure C.1: Timing diagram for the SNO front end. Representative values of the timing are given in Tab. C.1.

controls three different settings:

1. the total live time for a chip (T_{TOT}), i.e., the minimum time between successive firings

2. the short and long charge sample times ($T_{\mathrm{Q_{HS}}}$ and $T_{\mathrm{Q_{HL}}}$), which determine how long the integrator is active for short and long samples

3. the time in which a global trigger is valid (T_{GTVALID})

In addition, this chips-et produces the channel-level trigger primitives. A complete timing diagram is shown in Fig C.1. Physically, the functionality is split across three chips. There exist two Penn-designed and built bipolar chips that function as an integrator and discriminator-timer. The trigger and TAC functionality is in a CMOS chip which was jointly built by Queen's University and Penn, which also includes the analog storage array.

C.1 Integrator and discriminator timing

The integrator provides three separate charge measurements of the PMT signal, of two different gains and with varying integration times, as described in Sec. 4.4.2 and in Ch. 5. This part of the timing is controlled by a voltage ramp in the discriminator-timer called DISC_RAMP. The discriminator fires when a signal from the PMT reaches a value above the threshold. At this time, a voltage starts ramping down from a preset value. The firing of the discriminator marks the beginning of time for this event, labeled a in figure C.1. The rate of descent is controlled by a signal called RMPU. The short integrate time $T_{\mathrm{Q_{HS}}}$ is defined by when this signal crosses a certain value, called V_{SI}. The crossing time is labeled d in the figure. Upon crossing V_{SI}, the ramp signal starts ramping up. The rate of ascent is controlled by a combination of the RMPU signal previously mentioned and another voltage called RMP.

The long integrate time $T_{\mathrm{Q_{HL}}}$ is defined by when the ramp signal crossed a value defined by a signal called V_{LI}, labeled f in the figure. At this point, the chip resets itself, and is ready to accept pulses some time T_{deadtime} later (point g), which induces a channel-based dead time. Representative values for the timing are listed in table C.1. The times are determined by physical limits such as the light crossing time in the tank,[2] the time it takes a trigger signal to go round-trip from the FEC, to the central timing rack, and back to the FEC, and by the TAC non-linearity, discussed below.

A somewhat unfortunate choice was made in that the timing of the system is split across two of these chips: the discriminator-timer determines the timing of the charge integration and channel dead time, while the CMOS chip provides the timing for the trigger signal and the channel live time. Synchronization between these two chips must be provided manually. A better design choice for the final system would have been to allow the discriminator-timer to provide the timing for both the TAC and the integrator, thereby automatically synchronizing the system.

C.2 TAC and trigger timing

The TAC and trigger timing, as mentioned in the introduction to this chapter, is independent of the integrator-discriminator timing, except that the TAC time is initiated by the discriminator. The QUSN7 CMOS chip's timing is also determined by

[2] At $r_{\mathrm{PSUP}} \sim 8.4m$, the SNO tank crossing time is $T_x \sim 80ns$.

a voltage ramp. The rate of descent here is controlled by two current sources, I_{setm} and I_{seta}. Each of these signals is adjustable on a board level, i.e., 32 channels at a time. In addition, secondary digital adjustment sources can adjusted on a per-channel level.[3]

Once the ramp has started, initiated by the discriminator firing (labeled as a in figure C.1), a trigger signal TR_{100} is emitted (see label b) and sent to the central timing rack.[4] At point c in the figure, a signal called GTVALID is asserted. GTVALID controls the effect of the GT signal on the given channel.

If the arrival of a GT signal coincides with the assertion of GTVALID, the data presently being sampled by the integrator and discriminator will be stored and processed for inclusion into the data stream. In this case, the ramp is stopped at the arrival time of the global trigger, and the value of the ramp voltage at that time encodes the time between the discriminator firing and global trigger, as indicated at position e.[5] Once the integrator is finished with its integration, the voltages at the integration capacitors are stored in the next available set of analog storage arrays and the channel sets its DATA_AVAILABLE signal, indicating to the sequencer that it has data (see section A).

In the case where no GT arrives before the TAC ramp has reached the level designated by TACREF, the chip resets itself and discards the data from this firing, labeled f in the diagram. However, due to a design flaw in the QUSN7, there is a small window of opportunity after the TAC ramp has started to reset itself but before

[3]The TAC TRIM bits, see [69]

[4]For a complete description of the SNO trigger system, see [23].

[5]The reader who is still awake will notice that there are two times labeled e and e'. In the case of two otherwise identical events, the arrival time of the GT signal can vary by 20 ns. The global trigger signal is synchronized to a 50 MHz counter, which is used for inter-event timing. This synchronization leads to a 20 ns jitter in the arrival time of GT. However, it allows accurate event-by-event system level timing and does not imply any loss of accuracy.

the chip stops accepting triggers. The design flaw leads to a dual-valued TAC values for certain GT times. If a GT arrives in the time labeled on the diagram as *TAC non-linear regime*, the TAC will measure a value that could also have occurred in the time between labels e and f. However, the effect of this seemingly fatal flaw can be minimized. By tuning the values T_{tot} and T_{GTVALID} appropriately, the times in question here can be forced to be *early* light, i.e., light that happened well before the arrival of the GT signal. The double-valued times can be placed *before* the trigger and hence in the less interesting part of the trigger window, before the event that triggered the detector. For normal triggers, the light in this time slice will not be light from the physics event. However, this does constrain the values one can set the timing ramps to in table C.1.

A second design flaw in the CMOS chip leads to the "double zero" orphans described in Sec. 6.1.1. When a GT reaches the CMOS chip during the end of GT-VALID, the tag that identifies the event (the GTID) sometimes is stored incorrectly. The internal counter that tracks the GTIDs is a ripple counter, and during certain circumstances, the tag that is placed on an event will occur while the counter is 'rippling' and the upper and lower nybbles will be inconsistent. This flaw means that the SNO data stream will always contain a small contamination of orphans and a similar contamination of mis-paired children, that is, PMT bundles that are grouped with the wrong event.

Bibliography

[1] J. Bahcall, *Neutrino Astrophysics* (Cambridge University Press, Cambridge, 1989).

[2] J. Bahcall, available at `http://www.sns.ias.edu/~jnb/`.

[3] J. Abdurashitov *et al.*, in Suzuki and Totsuka [70], pp. 20–25.

[4] T. A. Kirsten, in Suzuki and Totsuka [70], pp. 26–34.

[5] Y. Fukuda *et al.*, Phys. Rev. Let. **77**(9), 1683 (1996).

[6] Y. Fukuda *et al.*, Phys. Rev. Let. **82**(12), 2430 (1999).

[7] N. Hata, S. Bludman, and P. Langacker, Phys. Rev. D **49**, 3622 (1994).

[8] N. Hata and P. Langacker, Phys. Rev. D **56**(10), 6107 (1997).

[9] B. Pontecorvo, Sov Phys JETP (6), 429 (1958).

[10] G. Ewan *et al.*, *Sudbury Neutrino Observatory Proposal*, SNO technical report SNO-STR-87-12, Queen's University (1987).

[11] H. H. Chen, Phys. Rev. Let. **55**(14), 1534 (1985).

[12] N. Schmitz, *Neutrinophysik* (Teubner, Stuttgart, 1997).

[13] C. W. Kim and A. Pevsner, *Neutrinos in Physics and Astrophysics*, vol. 8 of *Contemporary Concepts in Physics* (Harwood academic publishers, Amsterdam, 1993).

[14] F. Boehm and P. Vogel, *Physics of Massive Neutrinos* (Cambridge University Press, Cambridge, 1992).

[15] L. Wolfenstein, Phys. Rev. D **17**, 2369 (1978).

[16] S. P. Mikheev and A. Y. Smirnov, Soviet Journal Nuclear Physics **42**(6), 913 (1985).

[17] J. M. Conrad, in *International Conference on High Energy Physics (ICHEP98)* (1998), arXiv:hep-ex/9811009.

[18] J. V. Jelley, *Čerenkov Radiation and Its Applications* (Pergamon Press, London, 1958).

[19] C. M. Lederer and V. S. Shirley, eds., *Table of Isotopes* (John Wiley & Sons, New York, 1978), 7th ed.

[20] A. Poon, T. Burritt, C. Duba, S. Elliott, K. Heeger, and R. Robertson, *The Design of the SNO NCD Construction Hardware In-situ Monitoring Experiment (CHIME)*, SNO technical report SNO-STR-98-016, University of Washington (1998).

[21] J. Boger *et al.*, Nuclear Instruments and Methods **A449**, 172 (2000).

[22] T. Ekenberg, F. M. Newcomer, R. Van Berg, A. Biman, and R. L. Stevenson, in *IEEE Transactions on Nuclear Science* (1995), vol. V0042, pp. 925–932.

[23] J. R. Klein, M. S. Neubauer, R. Van Berg, and F. M. Newcomer, *The SNO trigger system*, SNO technical report SNO-STR-97-036, University of Pennsylvania (1997).

[24] R. Stevenson *et al.*, in *Conference Issue of the 1994 Transactions on Nuclear Science* (1994).

[25] F. M. Newcomer and R. Van Berg, in *IEEE Transactions on Nuclear Science* (1994), vol. V0042, pp. 745–749.

[26] R. J. Ford, *Calibration of SNO for the detection of* ^{8}B *neutrinos*, Ph.D. thesis, Queen's University, Kingston, Ontario (1998).

[27] D. S. McDonald, *Studies of the Sudbury Neutrino Observatory Detector and Sonoluminescence using a Sonoluminescent Source*, Ph.D. thesis, University of Pennsylvania (1999).

[28] R. J. Ford, *SNO* N_2*/Dye Laser and Fibre Optic Diffuser Ball System Manual*, Tech. Rep. SNO-STR-96-056, Queen's University (1997).

[29] R. J. Ford, *Fitting Optical Constants in SNO using a Pulsed* N_2*/Dye Laser*, SNO technical report SNO-STR-96-055, Queen's University (1997).

[30] W. R. Nelson, H. Hirayama, and D. W. Rogers, *The EGS4 code system*, Stanford Linear Accelerator Center (1985).

[31] The SNO Collaboration, *The SNOMAN User's Manual*, version 3.02 ed.

[32] The SNO Collaboration, *The SNOMAN Programmer's Manual*, version 3.02 ed.

[33] J. Beacom and P. Vogel, Phys. Rev. D **58**, 093012 (1998).

[34] P. Wittich, *Do flasher light emissions have considerable time width*, on the web at `http://www.hep.upenn.edu/~wittich/analysis/fl_twidth/flasher_twidth.html`.

[35] P. Wittich, *Data Cleaning: The Cluster Cut*, SNO technical report SNO-STR-2000-001, University of Pennsylvania (2000).

[36] J. R. Klein, *Data Cleaning: The Q Cluster Cut*, SNO data cleaning report, University of Pennsylvania (199x).

[37] J. R. Klein, *Data Cleaning: The QvT Cut*, SNO data cleaning report, University of Pennsylvania (2000).

[38] M. S. Neubauer, *Data Cleaning: The Fitterless Time Spread Cut*, SNO technical report SNO-STR-2000-002, University of Pennsylvania (2000).

[39] M. S. Neubauer, *Addendum to the FTS Report*, SNO Data Cleaning report, University of Pennsylvania (2000).

[40] N. McCauley, *Data Cleaning: The Burst Cut*, SNO Data cleaning report, University of Oxford (2000).

[41] N. McCauley, *Data Cleaning: The Neck Cut*, SNO Data cleaning report, University of Oxford (2000).

[42] R. G. Van de Water, *Data Cleaning: The Q/NHIT cut*, SNO data cleaning report, University of Pennsylvania (2000).

[43] J. R. Klein, *Data Cleaning: An Analog Measurement Board Cut*, SNO data cleaning report, University of Pennsylvania (1999).

[44] N. Tagg, *OWL Tube Cut Document*, SNO Data Cleaning report, University of Guelph (2000).

[45] J. R. Klein, personal communication.

[46] M. Yuk, *Evidence for Solar Neutrinos at SNO*, SNO data cleaning report, TOXENN (TRIUMF, Oxford, University of Pennsylvania) (2000).

[47] A. D. Marino, *Muon-induced spallation products in the SNO detector*, SNO technical report, Lawrence Berkeley National Laboratory (2000).

[48] N. Tagg, *Data Cleaning: The Muon Tagg*, SNO Data cleaning report, University of Guelph (1999).

[49] N. Tagg, personal communication.

[50] P. Skensved and B. C. Robertson, *Summary of Backgrounds in SNO*, SNO technical report SNO-STR-94-13, Queens University (1994).

[51] M. S. Neubauer, personal communication.

[52] J. D. Jackson, *Classical Electrodynamics* (John Wiley & Sons, New York, 1975), 2nd ed.

[53] C. Grupen, *Particle Detectors*, Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology (Cambridge University Press, New York, 1996).

[54] A. W. P. Poon, *Energy Calibration of the Sudbury Neutrino Observatory using Monoenergetic Gamma-Ray Sources*, Ph.D. thesis, University of British Columbia, Vancouver, British Columbia (1998).

[55] A. S. Hamer, *Energy Calibration of SNO for measurement of the Charged-Current Neutrino Reaction*, Ph.D. thesis, Queen's University, Kingston, Ontario (1999).

[56] S. J. Brice, *The Elastic Fitter*, SNO technical report SNO-STR-95-040, Oxford University (1995).

[57] M. E. Moorhead, *Grid Fitter for Reducing Tails in Spatial Distributions*, SNO technical report SNO-STR-95-042, Oxford University (1995).

[58] C. J. Jillings, *The Angular Resolution of SNO*, SNO technical report SNO-STR-98-010, Queen's University (1998).

[59] C. J. Jillings, *The Electron Scattering Reaction in the Sudbury Neutrino Observatory*, Ph.D. thesis, Queen's University, Kingston, Ontario (1999).

[60] M. Boulay, J. Farine, A. Hamian, and P. Wittich, *Data Selection for Neutrino Analysis*, SNO technical report SNO-STR, Queen's University, SNO, University of Washington, University of Pennsylvania (2000).

[61] R. G. H. Robertson, *Extraction of the Charged-Current Signal from the Angular Distribution*, SNO technical report SNO-STR-98-014, University of Washington (1998).

[62] K. Kubodera and S. Nozawa, International J. Mod. Phys. E **3**(101), 579 (1994).

[63] J. N. Bahcall, S. Basu, and M. H. Pnsonneault, Phys. Lett. B **443**(1) (1998).

[64] M. S. Neubauer and J. R. Klein, *Measurement of the SNO Trigger Efficiency*, SNO data cleaning report, University of Pennsylvania (2000).

[65] H. Sobel, *The study of atmospheric neutrinos with Super-Kamiokande*, as presented at the *XIX International Conference on Neutrino Physics and Astrophysics (Neutrino 2000)* (2000).

[66] K. Grotz and H. V. Klapdor, *Die Schwache Wechselwirkung in Kern-, Teilchen- und Astrophysik* (Teubner, Stuttgart, 1989).

[67] Y. Suzuki, *Solar neutrino results from Super-Kamiokande*, as presented at the *XIX International Conference on Neutrino Physics and Astrophysics (Neutrino 2000)* (2000).

[68] S. P. Riley *et al.*, Phys. Rev. C **59**, 1780 (1999).

[69] R. L. Stevenson, *SNO CMOS Analog storage Array*, Queen's University, Kingston, Ontario (1995).

[70] Y. Suzuki and Y. Totsuka, eds., *Neutrino 98: Neutrino Physics and Astrophysics*, vol. 77 of *Nuclear Physics B (Proc. Supplement)* (Elsevier Science, Amsterdam, 1999).