Through-Going Muons at the Sudbury Neutrino Observatory

by

Tyron Tsui

B.Sc., The University of British Columbia, 2000
M.Sc., The University of British Columbia, 2002

A THESIS SUBMITTED IN PARTIAL FULFILMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

The Faculty of Graduate Studies

(Astronomy)

The University of British Columbia

(Vancouver)

April 2009

Abstract

This thesis presents results from the through-going muon analysis at the Sudbury Neutrino Observatory. Over 1229.26 days, 77376 direct cosmic ray muons above $\cos\theta_{zenith} > 0.4$ and 514 neutrino-induced muons in the range $-1 < \cos\theta_{zenith} < 0.4$ were detected. A log-likelihood analysis was performed on the neutrino-induced muons to produce a flux measurement and constrain atmospheric neutrino mixing parameters. The best fit oscillation values were $\Delta m^2 = (2.6 \pm 2.0) \times 10^{-3}$ eV^2, $\sin^2 2\theta = 1.0 \pm 0.1$. The observed muon flux was found to be $(2.48 \pm 0.25) \times 10^{-13}$ cm^{-2}s^{-1}sr^{-1} resulting in an overall flux normalization 1.22±0.09 times the theoretical estimates from the Bartol group. The quoted uncertainties are 1-D marginalized. The no-oscillation hypothesis was ruled out at 99.8% confidence level. The observed number of direct cosmic ray muons was converted into a vertical depth intensity. The resulting fits to the intensity as a function of depth were $I_0 = (0.93 \pm 0.05) \times 10^{-6}$ cm^{-2}s^{-1}sr^{-1}, 2.33 ± 0.31 km.w.e. and $\alpha = 5.62 \pm 0.40$.

Table of Contents

List of Tables

List of Figures

Glossary and List of Acronyms

AV Acrylic Vessel that holds the D_2O

Background event An event defined to be not of physical interest

Cherenkov cone The cone of light produced by a charged particle traveling faster than the local speed of light.

Cut A quantity describing some aspect of an event which was used to distinguish signal from background

Contamination Background events that passed the cuts into the signal sample

Direct Cosmic Ray Muon A muon produced in the atmosphere by a direct cosmic ray

Downward-going Muon A muon produced in the atmosphere by a direct cosmic ray

Fly's Eye Air-fluorescence detector experiment designed to detect and track cosmic ray events above 10^17 eV.

LEAP The Low Energy Anti-proton experiment. A balloon-borne experiment searching for low-energy cosmic ray anti-protons.

MC Monte Carlo

NHit Number of hit PMTs in an event

Neutrino-induced Muon A muon produced by an atmospheric neutrino interacting with the rock surrounding SNO.

PDF Probability Distribution Function

PMT Photomultiplier Tube

PSUP Phototube Support Structure

RMS Root Mean Square

Run A period of time in which the detector was running in a particular state

Sacrifice Signal events that were removed by the cuts

Signal event An event defined to be one of the main physics events that were analyzed

SNO Sudbury Neutrino Observatory

SNOMAN SNO Monte Carlo Analysis Software

Stopping muon A muon that stopped in the detector

Through-going muon A muon that was created outside of the detector and travels at least all the way through the detector.

Upward-going Muon A muon produced by an atmospheric neutrino interacting with the rock surrounding SNO.

Acknowledgments

There are many people that I have met and worked with that have helped with the material presented here. I would first like to thank Chris Waltham for taking the time back at the very beginning to meet with me and helping develop my interest in the muon analysis at SNO. Our many meetings over the years have kept me moving through uncertain steps and excited through the small achievements. I was always grateful for quick insights he brought to problems I dwelled on for what felt like ages and the concise questions he asked which could shed light on a solution or an alternate path to take. The small UBC SNO meetings were also very helpful in me keeping a steady pace and to ensure thoroughness in my work. I thank Scott Oser for organizing the meetings and the other "social visits" to remind me to keep moving. Scott and Rich Helmer have both provided invaluable help as extremely resourceful local members of the collaboration.

The SNO postdocs at UBC have been extremely helpful in keeping our local software in great condition and been readily available for quick questions and helpful answers. Thomas Kutter helped to get me on my feet and learning all the parts and problems with the SNO software. Juergen Wendland and Blair Jamison have been very important in helping me through various software and analysis problems.

I would also like to thank my PhD committee. In the meetings over the years, they have helped a tremendous amount with understanding the tasks that need to be done and in what areas I needed the most help in. Bill McCutcheon, Mark Van Raamsdonk and Stan

Yen have all kept me in check and moving.

I am grateful for everyone in the SNO muon working group. Tom Walker and Charles Currat provided motivation and alternate analyses to what I was doing. A special thanks goes to Joe Formaggio for taking lead and keeping the analysis on track and Chris Kyba for the extremely productive work sessions and discussions. I would also like to thank Nathaniel Tagg, whose initial work on through-going muons at SNO was a great help in the work completed here. At the SNO site, both Noel Gagnon and Fraser Duncan were both amazing wealths of knowledge with respect to the inner workings of the detector and site procedures.

In addition to all the academic guidance I have received over the years, I am extremely thankful for all of the support from my family and friends. I have been so lucky to have my parents, sister, grandparents, uncles, aunts, cousins, friends and my wife with me along the way. Their encouragement and advice was so very appreciated, when all I could say was that I had "about a year left", three years in a row. All of you have a part in making me who I am and helping me accomplish this. I thank my mother, who has always been an amazing source of love and support. I thank my wife, Bonnie, who was always there for me through whatever frustration, excitement and joy that came along with this thesis.

Chapter 1

Introduction

Cecil Powell called it "a thin rain of charged particles" when he accepted his Nobel Prize for Physics in 1950. The prize was for the improvements he made in the photographic method of studying high-speed particles and his discoveries regarding mesons made with this method. These particles are now called cosmic rays and since their discovery in 1913, have played a significant role in the development of our understanding of high-energy particle interactions. This thin rain was the source of particles observed at the Sudbury Neutrino Observatory (SNO) and are the subject of this thesis.

In this introductory chapter, I will provide a summary of cosmic ray research with an emphasis on muons and neutrinos produced in the cosmic ray showers as they are the focus of this work.

1.1 Cosmic Ray Showers

Cosmic rays are high-energy particles incident on the Earth's atmosphere from astrophysical sources. These primary particles are roughly 90% protons, 9% helium nuclei, 1% electrons and a small fraction of heavier nuclei. They can have energies up to 10^{21} electron volts (eV) and interact with our atmosphere to create showers of secondary mesons and leptons. Provided here are a short history of cosmic ray research and a description of cosmic ray properties.

1.1.1   Early Work

Cosmic rays were discovered between 1911 and 1913, when Victor Francis Hess carefully measured levels of ionizing radiation at varying altitudes in a series of manned balloon flights. This came at a time when investigations into radiation became numerous following the discoveries of X-rays in 1895 by Wilhelm Roentgen and radioactivity in 1896 by Henri Becquerel. Early measurements of radiation levels as a function of altitude were puzzling since it was thought at the time that the only source of ionizing radiation came from elements in the Earth, like uranium and thorium. The levels did decrease as Father Thomas Wulf showed in 1910 with measurements up the Eiffel Tower, but the reduction was lower than expected [1].

The advancement of cosmic ray research relied upon three main bases. First is an inventory of all the particles involved, next an understanding of the interactions between the particles and finally more precise and varied methods of detecting the particles. Hess took one of the first steps after he increased the precision of his measuring equipment and then systematically measured the radiation levels at the surface of the Earth up to 5.3 km. He showed that radiation levels decreased up to roughly 1 km, but above that they began to increase. The recorded levels taken at 5 km was twice that of those measured at sea level. For that work and concluding that there was radiation incident on our atmosphere from space, Hess won the Nobel Prize in Physics in 1936. Confirming Hess's results, Robert Millikan named the radiation "cosmic rays" in 1925 [2]. While Millikan was credited with initially naming cosmic rays, Compton was the first to correctly suggest that cosmic rays were charged particles [3].

The radiation detected at ground level consisted both a soft component of electrons and photons and a hard component. The composition of the hard component was unknown until 1936. The soft and hard labels refer to the ability of the particle to penetrate matter.

Observations at a wide range of altitudes showed that the showers developed as the cosmic rays interact with our atmosphere. Studying cosmic rays at lower altitudes proved easier since experimental equipment did not have to be moved high up onto mountains or lifted up into the atmosphere with a balloon. Of course, knowing the composition of the shower at the top was fundamental in understanding that the primaries were charged particles and were the source of the shower of secondaries.

A number of discoveries came together in the late 1930s and 1940s that helped solve the puzzle. The introduction of cloud chambers allowed for particles to be detected along with their tracks. The shower particles ionized the material in the chambers as they passed through the detectors and the amount of ionization could be measured by the thickness of the track. Geiger-Müller counters also became available and made particle detection easier as the devices simply produced a pulse after a charged particle passed through it.

Primary cosmic rays were identified as protons and heavier nuclei in high altitude balloon flights in nuclear emulsions. A balloon experiment in 1948 confirmed that the primary cosmic rays were mostly protons with some helium nuclei and a small fraction of heavier nuclei [4]. This information and ground level observations of the soft component provided a beginning and an end, but the interactions and particles in between were yet to be understood. The missing pieces at this point were identifying the muon and pion and developing a theory of the electromagnetic interactions involved in the shower.

Carl D. Anderson discovered the muon in 1936 (the hard component of the shower) when he observed that certain cosmic rays curved when passed through a magnetic field [5]. Particles with the same velocity curved more than protons, but less than electrons. Assuming they had the same charge as an electron, they had to have some intermediate mass.

In 1937, a paper by Bhabha and Heitler was published describing how cosmic ray showers

developed in a cascade production of γ-rays and electron pairs [6]. This paper provided the theoretical ground work for understanding the electromagnetic cascade and was supported by Bruno Rossi in 1934 and Pierre Auger in 1939 [7]. They independently observed nearly simultaneous events in detectors separated by a horizontal distance. Auger suggested that high-energy primary cosmic rays could interact with nuclei in the upper atmosphere to create an extended shower of particles that would be detected at sea level over areas of hundreds of square kilometres.

One of the last discoveries was that of the pion. Pions decaying into muons and then to electrons were directly observed in 1947 in an examination of particle tracks in a nuclear emulsion at a high-altitude mountain station [8]. We will see later that the development of a shower from a single primary to the extended array of particles involves the production and decay of pions and muons. The first observations of these particles connected those made at high altitude of the primary and ground observations of the secondaries.

1.1.2 Energy Spectrum

Over roughly 30 years of study and advancement, cosmic rays were the main source for high-energy particle physics research. Investigating particle interactions at high-energy and the discovery of new particles happened generally all in the cosmic ray field up until the 1950s. It was at this time that advancement in technology allowed for the construction of large scale high-energy particle accelerators which provided a controlled environment to study high-energy physics.

The interest in cosmic ray research progressed to higher energies partly because accelerators provided a better platform to study certain aspects of high-energy physics and partly because it was a natural progression of the field. Higher energy primary cosmic rays provided more extensive cascades, both in the number of particles produced and surface
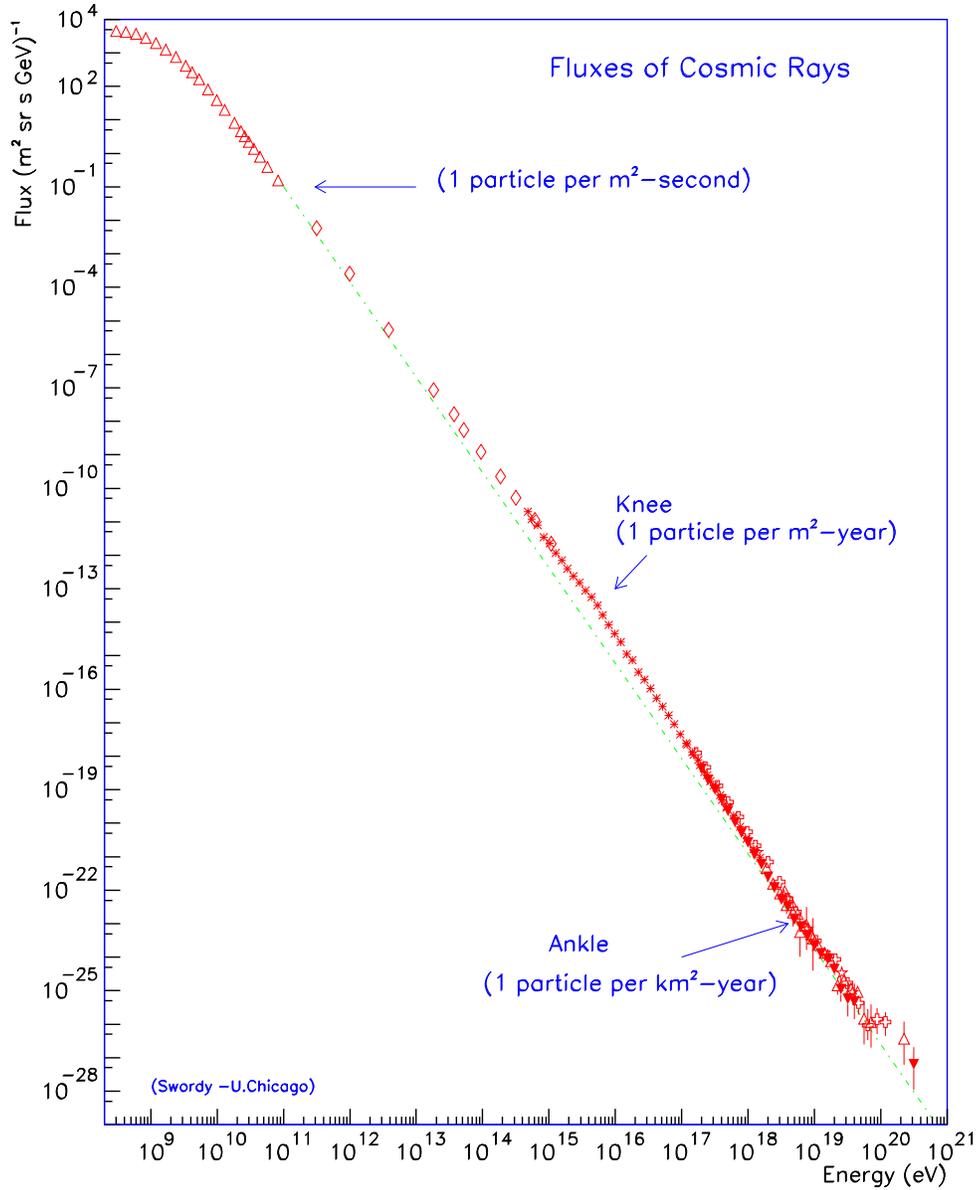
Figure 1.1: Energy spectrum of cosmic ray primaries. The green line shows E^{-3} for reference. The lower energy results are from single detector balloon or satellite experiments while the higher energy results are from ground arrays. This is a collection of data compiled by Swordy *et al* [9].

5

area covered on Earth. The current motivation behind cosmic ray studies is to try to understand the details of the large cascades and astrophysical sources that generate such high acceleration [10].

The energy spectrum of primary cosmic rays is shown in figure 1.1 which includes a green dashed line with E^{-3} for reference. The data are an accumulation of measurements taken over the last 100 years. The Large Hadron Collider (LHC) is the largest and most powerful accelerator, colliding proton beams at $7{\times}10^{12}$ eV per beam. Cosmic rays collide with essentially fixed targets in the atmosphere. To obtain energies for particle creation similar to LHC, the cosmic rays need an energy of roughly 10^{17} eV.

The regions at low energy and high energy are well populated with data from a number of experiments. The intermediate energy range $(10^{11} - 10^{14}$ eV) is not as well represented. Data in this region are difficult to obtain. The lower energy measurements were taken with single detector balloon or satellite experiments which use a magnetic field to bend and detect the primaries (e.g. LEAP [11]). The high energy events pass right through those types of detectors and are tracked by large ground arrays detecting the air showers they produce as they travel through the atmosphere (e.g. Fly's Eye [12]). The intermediate energy range particles are too energetic to be detected by the single detectors and too low energy to produce large showers. That energy range is best probed underground by atmospheric neutrino-induced muon detectors. The energy spectrum shown in figure 1.1 follows an inverse power law given by

$$\frac{dN}{dE} \propto E^{-(\gamma+1)}. \tag{1.1}$$

It is common to write the exponent as $(\gamma+1)$ which represents the differential spectral index, where γ is the integral spectral index. There are really only two features in the spectrum that span over 10 orders of magnitude in energy and over 30 orders in flux. The

"knee" at 10^{16} eV is a steepening in the slope to $\gamma \sim 2.3$ and the "ankle" at $10^{18.5}$ eV is where the slope flattens to $\gamma \sim 1.7$. These 2 transitions are attributed to having different sources accelerating the particles. Those with energy lower than the knee are believed to be accelerated in remnants of supernovae and those higher up to the ankle accelerated in higher energy galactic events like actual supernova explosions. Events with energy greater than the ankle are thought to be extragalactic in nature and originate from the centres of active galaxies.

The flux of events around the ankle and at higher energies is so low that their study requires large complex arrays of detectors spread over hundreds of kilometres. Currently, the Auger Project in Argentina is the largest detector array intended to observe the highest energy cosmic rays and large showers. The array covers over 3000 km^2 and combines two detection methods. On the ground are 1600 water Cherenkov detectors which will detect charged secondaries. In addition, 4 atmospheric fluorescence detectors will track the glow from the particle shower. While older experiments had to piece together data from various sources at a number of altitudes, Auger was designed to observe most or all of the development of a shower. In November of 2007, the collaboration presented results showing a strong correlation of the direction of their 27 highest energy events with the locations of active galactic nuclei [13].

1.2 Muons

Current experiments are capable of studying extensive air showers and tracking the cascade of particles produced. We now look specifically at muons from these showers. Muons are produced from pion and kaon decay where the main decay branches are

$$\pi^{\pm} \rightarrow \mu^{\pm} + \nu_{\mu}(\bar{\nu}_{\mu}) \qquad (\approx 100\%) \tag{1.2}$$

$$K^\pm \to \mu^\pm + \nu_\mu(\bar{\nu}_\mu) \qquad (63.5\%) \tag{1.3}$$

The decays shown in 1.2 have a positive pion decaying into a positive muon and muon-type anti-neutrino and a negative pion decaying into a negative muon and muon-type neutrino. The particle produced from the decay of the negative particle is shown in parentheses. The muons can decay through 1.4 and produce neutrinos. This process will be described in further detail in the following section.

$$\mu^\pm \to e^\pm + \nu_e(\bar{\nu}_e) + \bar{\nu}_\mu(\nu_\mu) \qquad (\approx 100\%) \tag{1.4}$$

Relative to the other particles in the cosmic ray shower, muons are long-lived and have a small interaction cross section. While muons have historically been referred to as the penetrating or hard component of cosmic ray showers, the fact that they are charged and hence easy to detect makes them the dominant signal in cosmic ray observations [14].

1.2.1 Properties

Over the past 50 years, the study of muons in cosmic ray showers has developed from using them to probe the Earth's magnetic field to understanding hadronic interactions and cascade development in the showers themselves. The study of muons from cosmic ray showers helps to uncover the shower composition, development and the elemental abundance of the primary cosmic ray. Considering that muon production from cosmic rays comes from pions and kaons, Gaisser [14] provides a semi-empirical formula for the observed number of muons for a given incoming angle, θ, at the surface of the Earth.

$$\frac{dN_\mu}{dE_\mu} \approx \frac{0.14E_\mu^{-2.7}}{\text{cm}^2 \text{ s sr GeV}} \left\{ \frac{1}{1 + \frac{1.1E_\mu \cos\theta}{115 \text{ GeV}}} + \frac{0.054}{1 + \frac{1.1E_\mu \cos\theta}{850 \text{ GeV}}} \right\} \tag{1.5}$$

The general form of this equation involves the original primary inverse power law multiplied by the two terms inside the braces. The first is the contribution from pion decay and the second is that from kaon decay. The 115 GeV and 850 GeV values are the energies at which the pion and kaon are more likely to interact than to decay (critical energy). While there is a general agreement, the predicted numbers at energies lower than roughly 10 GeV are higher than the data. This deficit is attributed to energy loss effects which are not accounted for in equation 1.5 and become significant at lower energies. Pion decay dominates as the source of muons for most energies, while the kaon contribution increases as a function of energy and is as high as 27% at energies over 1 TeV.

Up to 100 GeV, where the decay probability of pions and kaons is high, the muon spectrum closely matches that of the parent. At higher energies, the atmosphere is not thick enough for the pions to decay, and the muon spectrum steepens by one power of the energy. Pions have a decay length $(c\tau_\pi)$ of 7.8 m and at 1 TeV this becomes 55 km.

Muons can also provide clues to the nature of the primary cosmic ray. Consider two primaries of equal energy, one a proton and the other an iron nucleus. The two will produce muon energy distributions that differ. The proton produces a muon shower from one energetic particle while the iron nucleus breaks up resulting in a large number of lower energy interactions.

1.2.2 Propagation and Energy Loss

High energy muons from cosmic ray showers can be observed deep underground. This helps to reduce many of the backgrounds observed at the surface of the Earth, but understanding how the muon propagates through rock and loses energy along the way becomes very important. Muons lose energy as they pass through matter both continuously and discretely. The following formulae describe the energy loss of high-energy muons where ionization and

radiative processes dominate. Continuous energy loss is from ionizing the material it is passing through and is approximately given by

$$\frac{dE}{dX} \approx -2 \ \frac{\text{MeV}}{\text{g cm}^{-2}}.$$ (1.6)

Generally, the energy loss per unit length is divided by a density to provide a value independent of the material density and give the units shown above. The energy loss due to ionization for relativistic particles is roughly constant and equation 1.6 is good to better than 5% over GeV scale energies.

Energy loss from discrete radiation processes is proportional to the muon energy and therefore become more important at higher energies. These discrete processes include bremsstrahlung, pair production and electromagnetic interactions producing hadrons. Combining 1.6 with a term to account for these gives

$$\frac{dE}{dX} = -\alpha - \beta E_\mu,$$ (1.7)

where $\alpha = 2$ MeV/g cm^{-2} and $\beta = \beta_{br} + \beta_{pair} + \beta_{em}$. The summed fractional energy loss from the three discrete processes is β and roughly 4×10^{-6} cm^2/g for rock. The value $\epsilon = \alpha/\beta \sim 500$ GeV provides the critical energy at which the energy loss from ionization equals the loss from radiation. Significantly above 500 GeV, energy loss is dominated by discrete processes. With a given starting energy (E_μ^0), one can calculate the average energy of a muon (E_μ) after traveling through X g/cm^2 of rock

$$E_\mu = (E_\mu^0 + \epsilon)e^{-\beta X} - \epsilon.$$ (1.8)

We can find the minimum initial energy for a muon to penetrate to depth X by setting

$E_\mu = 0$ and solving equation 1.8 for E_μ^0. This gives

$$E_\mu^{min} = \epsilon(e^{\beta X} - 1). \tag{1.9}$$

Predicting the number of muons observed at a given depth with only continuous energy loss is straightforward. At higher energies the discrete processes become more and more significant. While these processes are understood, predicting how they affect the muon spectrum becomes more involved, as each process has a certain probability to have the muon lose a significant amount of its energy in only one interaction.

Current methods for predicting muon fluxes at a given depth use Monte Carlo simulations which account for all of these processes. Generally 1D models are used where every particle is traveling in the same direction and decays and interactions produce particles which propagate in one line. This is sufficient for the simulation of high-energy events as the assumption holds true, but as one goes to lower energies more and more secondaries have a relatively sizable transverse momentum and therefore, only 3-D simulations account for the lower energy spread in the shower. We used 3-D simulations to model the muon flux at SNO; the details are provided in chapter 3.

The standard presentation of the data and theory is to show a depth-intensity plot as in figure 1.2. This shows data and the fitted function to these data for the Large Volume Detector in Gran Sasso [15]. The incoming angles of the muon events were converted to vertical depths through rock. The LVD experiment has been live since the middle of 1992 and is an array of liquid scintillator tanks. The facility is 900 m above sea-level in a highway tunnel under Gran Sasso. The detector array, under a complex overburden, consists of 840 tanks with a combined target mass of roughly 1 kton. In between the tanks a secondary array of gas-filled wire tubes allow for muon tracking.

With a power law spectrum of muons at the surface (equation 1.5) and equation 1.9, we
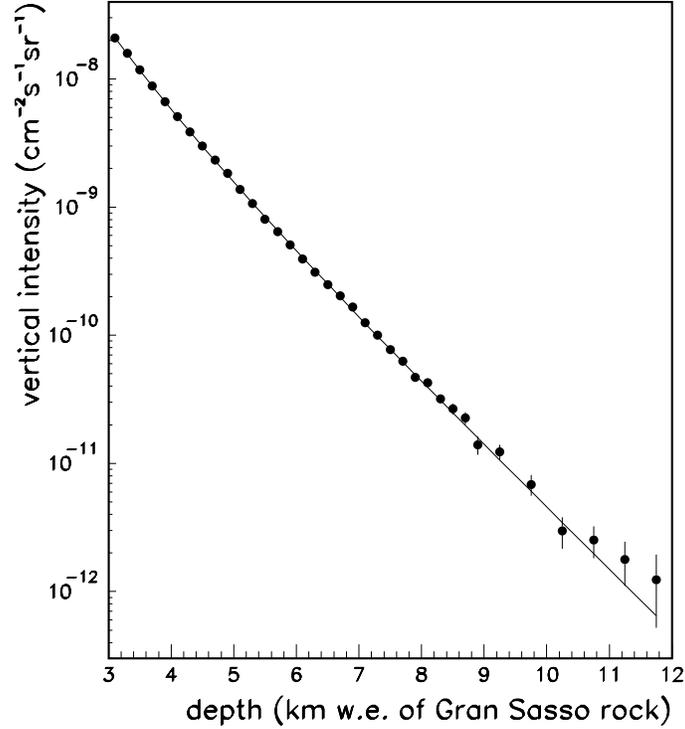
Figure 1.2: Depth-vertical muon intensity plot of the Large Volume Detector (LVD) muon data along with their best fit to the data [15].

can understand the parts of

$$F_{\mu}^{vertical} \propto \frac{\epsilon^{-\gamma+1}}{\gamma - 1} \times e^{-(\gamma-1)\beta X} \times (1 - e^{-\beta X})^{-\gamma+1}. \tag{1.10}$$

This gives the flux of muons at a certain vertical depth, X. As with equation 1.5, a constant of proportionality is left out of the form. Measurements at various depths and incoming angles can be converted into this vertical flux. The first factor originates from the primary cosmic ray spectrum at the surface of the Earth and shares that spectral index $(\gamma + 1)$.

12

The second factor dominates at large depths and comes from equation 1.9. The last factor is large and approaches unity at increasing depths. Data from other experiments will be presented below along with our analysis. Deeper and more precise measurements are needed to constrain uncertainties in the modeled energy loss processes and in understanding muon propagation.

1.3 Atmospheric Neutrinos

As can be seen from equations 1.2 and 1.3, the most abundant particles produced in cosmic ray showers are neutrinos. Only interacting weakly, neutrinos are even more penetrating than muons and difficult to detect. Here we will describe neutrinos from the shower.

1.3.1 Production, Ratios and Fluxes

Detailed studies of atmospheric neutrinos began to increase in number in the 1980s, when experiments like Kamiokande [16] and IMB (Irvine, Michigan, Brookhaven) [17] were being developed to search for proton decay. Atmospheric neutrino events are the strongest background in proton decay studies since the signal from both events are most dominant around 1 GeV. At the time, a large effort was put into understanding the theoretical flux of atmospheric neutrinos as well as measuring it.

Two neutrino production methods from cosmic rays were seen in decays 1.2 1.4. A clear prediction from these decays is the ratio of muon-type neutrinos to electron-type neutrinos. Counting the neutrinos produced in the $\pi \to \mu \to \nu$ chain shows

$$\frac{\nu_\mu + \bar{\nu}_\mu}{\nu_e + \bar{\nu}_e} = 2 \tag{1.11}$$

This ratio is fairly robust over most energies and is only slightly affected by kaon decays. At

energies above a few hundred GeV, the $\frac{\nu_\mu}{\nu_e}$ ratio is strongly dependent on the pion to kaon ratio which is only known with an accuracy of roughly 30%. At GeV scale energies, it is a convenient measurement for experiments to make since taking a ratio of events generally removes many systematic errors in an analysis. A summary of the data available for this ratio is given in table 1.1 with accompanying text in the following section.

The overall neutrino flux from pion and muon decay is calculated in the same manner as for muons. Gaisser accounts for the meson decay probabilities and the decay kinematics to give

$$\frac{dN_\nu}{dE_\nu} \approx \frac{0.0096 E_\nu^{-2.7}}{\text{cm}^2 \text{ s sr GeV}} \left\{ \frac{1}{1 + \frac{3.7 E_\nu \cos\theta}{115 \text{ GeV}}} + \frac{0.38}{1 + \frac{1.7 E_\nu \cos\theta}{850 \text{ GeV}}} \right\}. \tag{1.12}$$

As for the equivalent muon expression, the multiplying term has the $E^{-2.7}$ form from the primary cosmic spectrum and the two terms in braces are the pion and kaon contributions. The charged pion decay chain dominates the GeV range neutrino production, while the kaon contribution increases with energy. Pions are increasingly likely to interact before decay as energy increases. Kaon decays become the larger source of neutrinos above 115 GeV, where pions generally interact before decaying, and increase to contribute 3 times more neutrinos than the pions at the highest energies [10].

1.3.2 Detection and Data

Calculating the observed flux of muons underground involves taking the surface spectrum and propagating the muons through rock. The case for neutrinos is different in that the Earth is transparent to neutrinos up to roughly 10^5 GeV. At this energy and above the atmospheric neutrino flux is many orders of magnitude lower than the peak values. Atmospheric neutrinos up to this energy can therefore be detected deep underground where backgrounds like surface cosmic rays are significantly attenuated by the overburden of rock.

Those produced in the atmosphere on the opposite side of the Earth travel all the way through the Earth and are detected as coming up from the bottom of the detector. The underground neutrino flux is the same from above as from below. While the detector is closer to the surface and atmosphere on one side of the earth, the same angle directly opposite, covers more of the atmosphere. There is roughly a 2:1 ratio in the flux of neutrinos coming in horizontally compared to that of the flux coming vertically. This is caused by the increased path length through the atmosphere for the pion decay chain which produces neutrinos.

We can estimate the characteristics of a detector needed to observe a statistically significant sample of atmospheric neutrinos. The charged current interaction $(\nu_i + N \rightarrow \ell + X)$ is a common neutrino detection method. The neutrino of type i interacts with a nucleon in the detector to produce a corresponding lepton (ℓ) that is observed. The atmospheric neutrino flux is roughly 1 cm^{-2} s^{-1} and the neutrino interaction cross-section is of order 10^{-38} cm^2. The event rate can be written as the flux multiplied by the cross-section and the amount of interaction material available $(F_\nu \times \sigma_\nu \times N_A)$. Combining this gives

$$\frac{1 \ \nu}{\text{cm}^2 \ \text{s}} \times \frac{10^{-38} \ \text{cm}^2}{\text{nucleon}} \times \frac{6 \times 10^{32} \ \text{nucleons}}{\text{kT}} \times \frac{3.15 \times 10^7 \ \text{s}}{\text{year}} \sim \frac{180 \ \nu \ \text{events}}{\text{kT year}}. \tag{1.13}$$

A 1 kton detector operating for 1 year would detect approximately 180 events. Two detectors of this size, both built in the early 1980s to search for proton decay, were Kamiokande and IMB. Kamiokande was a 4.5 kT water Cherenkov detector, constructed 1 km underground in a Kamioka mine. IMB, an 8 kT detector, was located 600 m underground in a salt mine in Painesville, Ohio. Both experiments placed lower limits on the proton lifetime, but also made a measurement of the muon to electron neutrino ratio (equation 1.11). Kamiokande [16] and IMB [17] reported a ratio significantly smaller than the expected value

Experiment	Double Ratio
IMB	$0.54\pm0.05\pm0.11$
Kamiokande	$0.60\pm0.06\pm0.05$
Soudan 2	$0.68\pm0.11\pm0.05$
Super-K (sub-GeV)	$0.66\pm0.02\pm0.05$

Table 1.1: Double ratio results from 4 different experiments. The ratios were a comparison of the observed muon to electron ratio to that of the expected numbers. While consistent with each other, the significant deviation from a value of 1 required further investigation. Statistical and systematic errors for each are shown respectively.

of 2. The results were presented as a double ratio

$$R2 = \frac{R(\mu/e)_{observed}}{R(\mu/e)_{predicted}}. \tag{1.14}$$

Table 1.1 shows the results from Kamiokande, IMB, Kamiokande's successor, Super-Kamiokande (Super-K) and Soudan 2. Soudan 2 was also searching for proton decay in the 1990s, but was a 1 kT iron tracking calorimeter. The errors shown are statistical and systematic respectively. The reported double ratios were all significantly far from a value of 1, which is expected if the observed and predicted numbers agree. The disagreement motivated atmospheric neutrino study and larger next-generation detectors were built to resolve this puzzle along with more extensive proton decay searches. One such detector, Super-K, began observations in 1996 with a 50 kT volume and greater photocathode coverage. One possible early interpretation of the results was that the neutrino cross-section in water was not well known. Results from Soudan 2 that agreed with the water experiments helped to rule that out, since different instruments and experimental techniques were used. The observed double ratios revitalized interest in an idea proposed much earlier: neutrino oscillations.

1.4 Oscillations

In the original Standard Model of particle physics neutrinos were massless. This is because only left-handed neutrinos have been observed, and fermions are thought to have mass because of interactions with the Higgs field involving both right and left-handed versions of the fermion. The handedness or helicity of a particle gives the sign of the spin of the particle in the direction of motion. Non-zero neutrino masses bring up questions regarding the origins of neutrino mass and mass mechanisms in general.

The atmospheric neutrino anomaly was not the only evidence for neutrino oscillations. A deficit was also detected in the number of neutrinos coming from the Sun by Ray Davis. His experiment used chlorine as a target to detect solar electron-type neutrinos in the Homestake gold mine in the late 1960s [18]. As with the study of atmospheric neutrinos, evidence for neutrino oscillations being the cause of the solar neutrino deficit grew over the years as larger and newer experiments confirmed the deficit. The most compelling evidence for a non-zero neutrino mass came in early 2001/2 from SNO [19]. SNO was designed to detect three types of neutrinos and distinguish electron neutrinos from the others. A detailed history of the solar neutrino problem is given by Bahcall [20].

1.4.1 Oscillation Theory

In the early 1960s, Maki, Nakagawa and Sakata [21] proposed the possibility of oscillations between muon and electron-type neutrinos (tau's were unknown). Reference is also often made to Pontecorvo's work around roughly the same time. While also developing neutrino oscillations, Pontecorvo was investigating neutrino-antineutrino oscillations [22]. At the time, neutrinos were thought to be massless particles which move at the speed of light. This gives neutrinos a fixed flavour state; electron-type neutrinos exist and propagate purely as electron-type neutrinos. The idea of neutrino oscillations originated from mixing in

quarks and suggested that if neutrinos had small and different masses, neutrino flavour mixing would be possible. Neutrinos would be created or annihilated as flavour eigenstates $(\nu_e, \nu_\mu, \nu_\tau)$, but propagate through space as mass eigenstates (ν_1, ν_2, ν_3). Different mass eigenstates allow for a phase difference to develop over the distance traveled. This would result in an oscillation in the detected neutrino flavour.

To develop this concept, we can approach the problem as in reference [23] and consider two neutrino flavours (ν_e, ν_μ). While born and detected as ν_e and ν_μ, each are a linear combination of two mass eigenstates so that the wave functions

$$
\begin{aligned}
\nu_\mu &= \quad \cos\theta \times \nu_1 + \sin\theta \times \nu_2 \\
\nu_e &= -\sin\theta \times \nu_1 + \cos\theta \times \nu_2.
\end{aligned}
\tag{1.15}
$$

are orthonormal states. This unitary transformation ensures the mixing conserves the total number of neutrinos. The propagation of the mass states through space is given in the following equation (1.16) where $\hbar = c = 1$.

$$
\begin{aligned}
\nu_1(t) &= \nu_1(0)e^{-iE_1 t} \\
\nu_2(t) &= \nu_2(0)e^{-iE_2 t}
\end{aligned}
\tag{1.16}
$$

To determine the oscillation probability of one type of neutrino being detected with the same flavour with which it was born, the ratio of the intensity over time to the initial intensity needs to be expressed. Muon neutrinos initially have $\nu_\mu(0) = 1$ and $\nu_e(0) = 0$. This provides an expression for the initial mass states by inverting the equations in 1.15,

$$
\begin{aligned}
\nu_2(0) &= \nu_\mu(0)\sin\theta \\
\nu_1(0) &= \nu_\mu(0)\cos\theta
\end{aligned}
\tag{1.17}
$$

and provides

$$\nu_\mu(t) = \cos\theta\nu_1(t) + \sin\theta\nu_2(t). \tag{1.18}$$

Now the amplitude, A_μ is given by

$$A_\mu = \frac{\nu_\mu(t)}{\nu_\mu(0)} = \cos^2\theta e^{-iE_1 t} + \sin^2\theta e^{-iE_2 t}. \tag{1.19}$$

Finally, the intensity is

$$\frac{I_\mu(t)}{I_\mu(0)} = A_\mu A_\mu^* = 1 - \sin^2 2\theta \sin^2 \frac{(E_2 - E_1)t}{2}. \tag{1.20}$$

In the relativistic limit, we can write

$$E_i \simeq p + \frac{m_i^2}{2p}. \tag{1.21}$$

Using equation 1.21 along with $\Delta m^2 = m_2^2 - m_1^2$ and $t = L/c$, the probability of the detecting ν_μ after a distance L is

$$P(\nu_\mu \rightarrow \nu_\mu) = 1 - \sin^2 2\theta \sin^2\left(\frac{1.27\Delta m^2 L}{E}\right) \tag{1.22}$$

and

$$P(\nu_\mu \rightarrow \nu_e) = 1 - P(\nu_\mu \rightarrow \nu_\mu). \tag{1.23}$$

The energy E has units of GeV, the distance traveled L is in km and Δm^2 in eV^2. Experimental sensitivity for Δm^2 comes from the L/E ratio and the $\sin^2(2\theta)$ amplitude term depends on the statistics.

1.4.2 Observational Results

With 3 neutrino mass eigenstates, there are 3 possible mixing pairs and the simplified mixing matrix (1.15) becomes a 3×3 matrix. The only reason that the simplified mixing matrix is referenced so often is that one of the mixing angles (θ_{13}) is very small. Having one very small mixing angle allows the two-flavour scheme to be an accurate description. The full mixing matrix (named the Maki-Nakagawa-Sakata or Pontecorvo-Maki-Nakagawa-Sakata matrix) relating the three flavour states to the three mass states can be parameterized as

$$
U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & e^{i\delta}s_{13} \\ 0 & 1 & 0 \\ -e^{-i\delta}s_{13} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{1.24}
$$

In the above matrix, $c_{ij} \equiv \cos\theta_{ij}$ and $s_{ij} \equiv \sin\theta_{ij}$. The neutrino parameters associated with atmospheric neutrinos are generally given by Δm_A^2 and θ_A or Δm_{23}^2 and θ_{23}. Those associated with solar neutrinos are Δm_S^2 and θ_S or Δm_{12}^2 and θ_{12}. The subscript numbers refer to the two mass eigenstates in each coupling. Table 1.2 summarizes the current observed oscillation values. The solar (θ_{12}) results have been a combination of the reactor antineutrino experiment, KamLAND and water Cherenkov experiments (SNO, Super-K). The values quoted in the table are the best current estimates from SNO 2008 [24]. The atmospheric values are from MINOS 2008 [25]. Super-K is 1 km underground in a mountain mine site in Kamioka, Japan. The detector is a stainless steel cylinder holding 50 ktons of ultra-pure water. Cherenkov light that is produced in the water is detected by the array of over 11 000 PMTs. Designed to search for proton decay, the experiment also detects solar and atmospheric neutrinos. With the large target volume and PMT coverage and the long livetime, Super-K has been the leader in terms of particle exposure and collected statistics.

The MINOS experiment was designed to study the oscillation of neutrinos produced

by the Fermilab main injector. There are two detectors, one near the neutrino production beam and the other 735 km away in an underground mine in northern Dakota. The two detectors have similar characteristics, but the far detector has a 5.4 kton target while the near detector is just under 1 kton. Each detector consists of alternating planes of magnetized steel and plastic scintillator. The magnetic field allows the separation of positively and negatively charged events. The main measurement MINOS was designed to make was the disappearance of muon-type neutrinos that were produced in the beam and detected at the far detector.

The mass-squared difference results from the solar and atmospheric experiments suggest that two of the three mass states are significantly closer to each other than the third. The last mass-squared difference (Δm_{13}^2) is usually quoted as approximately equal to Δm_{23}^2 within current experimental errors. The CHOOZ experiment constrains $\sin^2(2\theta_{13})$ to be less than 0.17 [26]. A direct measurement of this small angle would help to complete the neutrino oscillation scheme as well as guide future experiments that hope to further probe neutrino properties. Experiments like T2K [27] and Double Chooz [28] aim to measure θ_{13}. A non-zero value for θ_{13} allows for full three-flavour neutrino oscillations and a CP violating phase term which is proportional to the size of θ_{13}. The phase term allows for a difference in the oscillation probability of neutrinos and antineutrinos. CP symmetry asserts that swapping a particle with its antiparticle (charge conjugation) and switching its parity has no effect on physical laws or interactions. If this were an absolute law, there would be no way in which the current forms of matter could exist from the Big Bang as all matter and antimatter would annihilate. CP violation was already detected in kaon [29] and B mesons [30], but the effect is not large enough to account for the matter-antimatter asymmetry observed in the Universe. Experiments needed to measure the CP violating phase for neutrinos are next-generation type and their design will depend on the constrained measurement of θ_{13}.

Pairing	Mixing Angle	Mass Squared Difference (ev^2)
1-2	$\theta = 34.4^{\circ}{}^{+1.3}_{-1.2}$	$(7.59 \pm 0.20) \times 10^{-5}$
2-3	$\sin^2(2\theta) = 1^{+0.0}_{-0.05}$	$(2.43 \pm 0.13) \times 10^{-3}$
1-3	$\sin^2(2\theta) < 0.17$	$\Delta m_{13}^2 \approx \Delta m_{23}^2$

Table 1.2: Summary of observed neutrino oscillation parameters.

If θ_{13} is found to be roughly equal to or greater than 0.01, intensified conventional neutrino beams should be able to probe CP violation. If θ_{13} is found to be significantly less than 0.01, experiments may need to rely on "neutrino factories" to observe it [31]. These factories will be designed to produce a very pure and intense source of neutrinos, from muon decays, that will be directed at another detector on the other side of the Earth.

With only the mass squared differences measured, the mass hierarchy is presently unknown. The two possibilities are a "normal" hierarchy where the third mass state is significantly greater than the the first two states and an "inverted" hierarchy where the third mass state is significantly smaller than the first two states. The absolute mass of neutrinos cannot be given by oscillations. The upcoming KATRIN experiment [32] is designed to measure the mass of the electron neutrino directly. The experiment is a next generation tritium beta-decay experiment which will measure the electron energy spectrum. A neutrino with non-zero mass would result in a lower maximum energy and a distortion in the high energy tail of the distribution since some of the energy available goes into the mass of the neutrino.

1.5 SNO Contributions and Thesis Organization

We have outlined the current state of cosmic ray research and briefly alluded to areas in which data from SNO will contribute. The work presented here will be developed into three results. The first is a measurement of the direct cosmic ray muon flux at SNO depth. The

second and third are related to the atmospheric neutrino-induced muons. Extracting the atmospheric oscillation parameters from our data will provide a nearly model-independent consistency check of existing data. We will also calculate an overall flux normalization for atmospheric neutrinos and provide a direct measurement of the flux of unoscillated atmospheric neutrinos.

As a heavy water Cherenkov light detector 2 km underground, SNO is in a unique position to study cosmic ray physics. We explained earlier that muons and neutrinos are the most penetrating components of cosmic ray showers. Neutrinos have an extremely small chance of interacting as they travel through the Earth and the highest energy cosmic ray muons can travel up to 5 km through rock. These direct cosmic ray (or "downward") muons can be observed with an incoming angle of up to $66°$ from zenith ($\cos 66° = 0.4$). A muon coming exactly downward towards SNO has the shortest path from the surface of the Earth, while the path-length increases as the zenith angle increases. This relationship between incoming angle and distance traveled through rock is simple because of the flat overburden at SNO. At mountain sites like Super-K the analysis is more complex. We saw in figure 1.2 that current experimental limits on the depth at which these muons are observed is roughly 12 km.w.e.. The significant depth of SNO will allow for measurements of the flux through more than 15 k.m.w.e.. Our measurements will help to constrain muon-rock interaction uncertainties at the highest energies.

While the direct cosmic ray muons can only been seen up to zenith angles of $66°$, this allows for muons produced by atmospheric neutrinos (or "upward muons") to be detected $24°$ above horizontal and at all angles below. When integrated over zenith angles up to $66°$ the direct cosmic ray muon flux overwhelms the atmospheric neutrino-induced muons by roughly $150\times$. In an upward facing cone with an opening angle of $66°$ the two components are indistinguishable.

The upward muons we observed reveal information regarding the atmospheric neutrinos produced in cosmic ray showers. While smaller and more statistically limited than Super-K, the depth of SNO provides a unique look at atmospheric neutrino oscillations. In the region above horizontal that is free of direct cosmic ray muons, SNO detects muons produced by atmospheric neutrinos that have not traveled far enough through the Earth to oscillate. This allows SNO to make a unique measurement of the absolute underground flux of atmospheric neutrinos and fit for oscillation parameters in one analysis. Previous experiments have not been deep enough to shield them from the overwhelming direct cosmic ray muon flux and allow the observation of the unoscillated atmospheric neutrinos above the horizon. Without these neutrinos, experiments had to rely on theoretical estimates of the flux. Doing this while estimating oscillation parameters in a disappearance experiment introduces uncertainties we do not have. Current estimates of the neutrino flux uncertainties are roughly $\pm 15\%$ and depend strongly on neutrino energy. The direct measurement of the unoscillated neutrino flux provided here, gives an experimental means to constrain this flux for future neutrino oscillation analyses. The largest contributor to the upward-muon flux comes from deep inelastic scattering of neutrinos. The uncertainty on the cross-section for this process is strongly constrained by high energy neutrino-beam experiments to about $\pm 3\%$.

The D_2O dataset was previously analyzed and presented in Tagg's thesis [33] in 2001. Aspects of the work described in that thesis provided a number of bases for the present analysis performed by the muon working group. The thesis helped in the understanding of muon event properties and associated biases, described the first through-going muon fitter at SNO and developed cuts for through-going muons.

The following chapters will be organized as follows. First, I will describe the detector and software programs used to simulate and analyze events. In chapter 3, the muon events seen at SNO will be characterized and the reconstruction algorithms detailed. Next, chapter

4 will present the simulated datasets and event cuts that were developed from the simulated events. Chapter 5 contains information on the real datasets and background analysis. In chapter 6 the analysis method is presented along with results and conclusions. Finally in chapter 7, a summary and discussion on future work is given.

Chapter 2

The Sudbury Neutrino Observatory

SNO has been extensively described in [34]. The intent here is not to reproduce all of the details involved with how the detector works, but to provide sufficient information to help the reader better understand the analysis to follow. We start with an overview of the experimental phases of SNO then continue with a description of the physical location of the detector and how the site and surroundings affect and facilitate the analysis. The properties of the detector will also be presented and finally the detector software and simulation code are detailed.

2.1 The Experimental Phases of SNO

The detector was designed to investigate solar neutrino oscillations and the livetime was divided into three different phases. Table 2.1 summarizes the phases. The first phase called the "D_2O phase" ran from November 2, 1999 to May 31, 2001. The D_2O target volume provided the means to measure the flux of electron-type neutrinos and the flux of all types of neutrinos coming from the Sun. All previous experiments only measured the ν_e flux and saw a deficit. To support the idea that the ν_e's from the sun were oscillating into other types of neutrinos, both the ν_e flux and total neutrino flux had to be measured. Within

the D_2O, three reactions with neutrinos were possible. The first was the charged current reaction

$$\nu_e + d \rightarrow p + p + e^-. \tag{2.1}$$

Here the incoming electron-type neutrino (ν_e) interacts weakly with the deuterium nucleus (d) in an exchange of a charged W boson to change the neutron in the deuterium into a proton (p) and the neutrino to an electron. The electron is ejected at a speed greater than the speed of light in water, producing Cherenkov light detected by the photomultiplier tubes (PMTs).

In the second reaction,

$$\nu_x + d \rightarrow p + n + \nu_x. \tag{2.2}$$

Here any type of neutrino (ν_x) interacts with the deuterium in an exchange of the neutral Z boson. This breaks apart the deuterium and releases a neutron which eventually captures onto another nucleus. The capture produces γ-rays which scatter electrons and again produces the detected Cherenkov light. The reaction does not favour any particular type of neutrino.

The last reaction had

$$e^- + \nu_x \rightarrow e^- + \nu_x. \tag{2.3}$$

This electron scattering process was the primary detection method in other light water experiments as it is not unique to heavy water. Unlike the neutral current reaction described in 2.2, this reaction favours ν_e interactions by a factor of six.

In the second phase 2 tonnes of NaCl salt was added to the D_2O. This increased the sensitivity for neutron detection as the cross-section for neutron capture described in reaction 2.2 is three orders of magnitude higher on Cl than D_2O. Neutron capture on Cl also produces multiple γ-rays that allows the use of isotropy parameters to further distinguish

Phase	Duration	Note
D_2O	19 months	Initial phase
Salt	24 months	2 tonnes NaCl added to the D_2O
NCD	23 months	40 proportional counters installed

Table 2.1: Summary of the three main phases of the SNO experiment.

neutron capture events from charged current electron events. There were no expected effects on the consistency of the muon analysis between the first two phases. The consistency is investigated in the following chapter. This "salt phase" ran from July 26, 2001 to August 29, 2003.

In the final phase, 40 proportional counters (36 filled with ^{3}He and 4 filled with ^{4}He) were suspended in the D_2O. The counters detected neutrons capturing on ^{3}He. The ^{4}He counters were installed to help calibrate backgrounds. This provided an independent method of measuring the neutrons released from the neutral current reaction. This "NCD phase" (Neutral Current Detectors) ran from January 3, 2005 to November 20, 2006.

Muon events produce significantly more light than solar neutrino events and fire thousands of PMTs. The muon track fitting routine that was used relies on both timing, charge and hit pattern information from all of the PMTs that fired. This makes the muon analysis relatively insensitive to the changes in the detector through the phases of SNO. The physical effect of having 40 NCDs in the D_2O was the shadowing and reflecting of light which would otherwise reach PMTs. The changes in muon event properties through the phases are examined in the following chapter.

2.2 Location

The detector is located on the 6800 foot level of Vale Inco's Creighton nickel mine in the city of Greater Sudbury, Ontario. The mine is Canada's deepest at roughly 2.2 km and the

Region	Material	Chemical Composition	Density (g/cm^3)	$< Z/A >$
Hanging Wall	Norite	$O_{45\%}Na_{2\%}Mg_{3\%}$ $Al_{9\%}Si_{26\%}Ca_{6\%}Fe_{7\%}K_{1\%}$	2.84 ± 0.10	0.491
Foot Wall	Granite/Gabbro	$O_{47\%}Na_{2\%}Mg_{1\%}$ $Al_{9\%}Si_{30\%}Ca_{5\%}Fe_{4\%}K_{2\%}$	2.83 ± 0.10	0.495

Table 2.2: Summary of hanging and foot wall properties. The cross-sectional view of SNO and the surrounding rock types are shown in figure 2.1.

detector is part of the the deepest underground science facility in the world. The detector was shut down at the end of 2006. The underground site was then expanded into adjacent excavated areas to provide space for a number of next generation experiments well suited to the depth of the site. SNOLAB is now the name for the entire facility and currently includes PICASSO and DEAP, both dark matter search experiments as well as future experiments like SNO+ [35].

The SNO cavity was excavated from the surrounding norite rock. Norite is igneous (solidified molten) rock and differentiated from granite by being rich in magnesium and iron [36]. A fault line 70 m from the detector divides the surrounding rock into two regions referred to as the hanging wall and the foot wall. As figure 2.1 shows, the detector sits in the hanging wall. The foot wall is mostly composed of granite and is sloped $25°$ from vertical. To better understand the local surroundings, 74 bore samples were taken at various depths. The data are shown in figure 2.2. This provided a measurement of the local rock density as a function of depth. Averaging over the samples gives a hanging wall density of 2.84 ± 0.10 g/cm^3 and a foot wall density of 2.83 ± 0.10 g/cm^3.

Detectors around the world are located in a variety of rock types. A standard measure of depth that allows comparisons to be made between experimental results in different rock is metres water equivalent (m.w.e) and is given in equation 2.4, where ρ is the rock or water density and h is the vertical depth. The conversion is from a depth through rock to
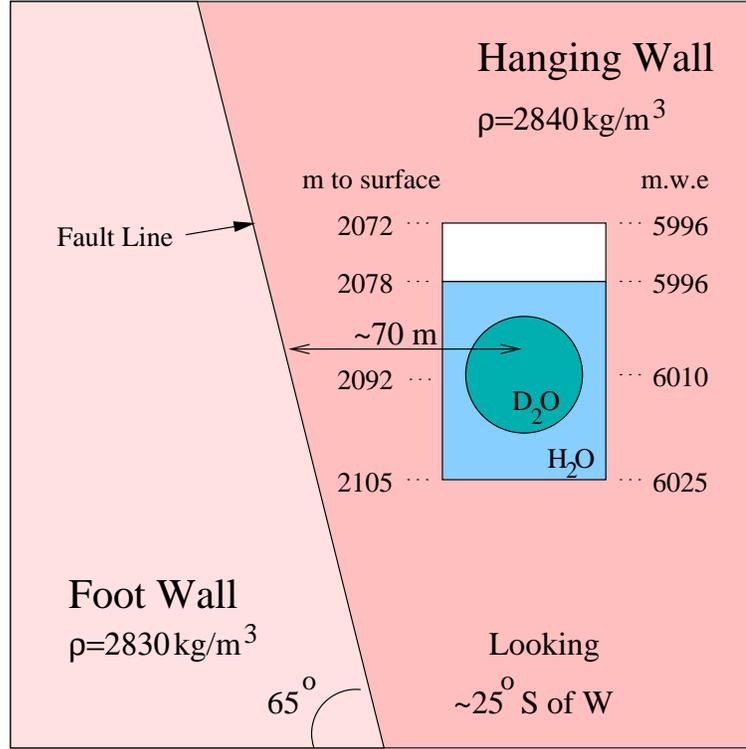
Figure 2.1: The rock surrounding SNO. The depth of the detector and position in the hanging wall are shown along with the fault line separating the hanging wall from the foot wall.

a standardized depth through water.

$$D_{m.w.e.} = \frac{\rho_{rock}}{\rho_{H_2O}} \times h \tag{2.4}$$

The overburden at SNO is generally flat. Within a 5 km radius on the surface above the detector, the topology is confined between 300 and 320 m above sea level. The vertical depth from the surface to the centre of the detector is 2092 ± 6 m or 5890 ± 200 m.w.e.. While the uncertainty in the physical measurement of depth is small, the larger uncertainty in the rock density propagates through to the depth value as quoted in metres water equivalent.

30
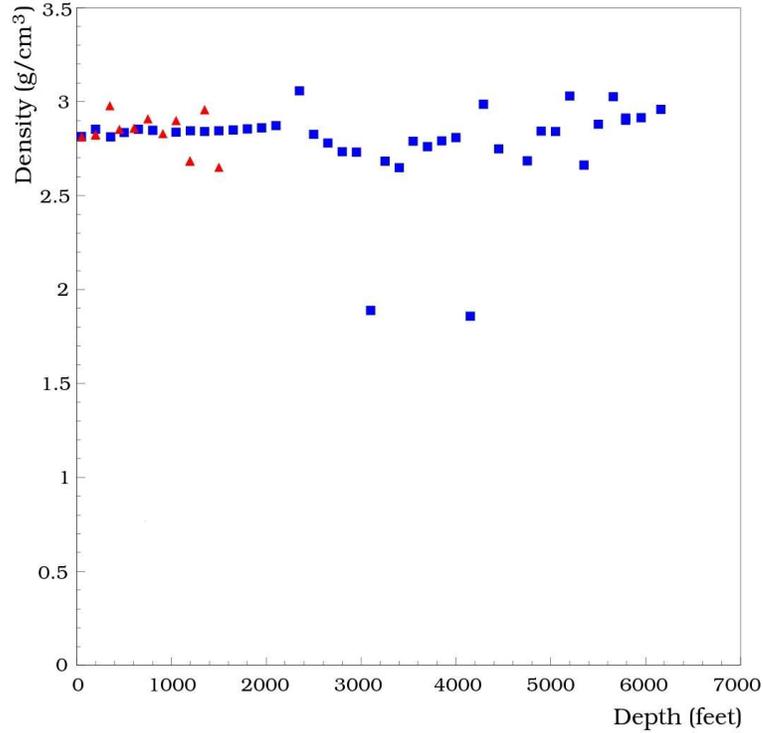
Figure 2.2: Bore sample data of the rock density near SNO. The blue squares indicate bore samples taken along the hanging wall, while triangles indicate density measurements in the footwall. The two outliers are thought to be samples of fractured rock as there are no types of rock with densities so low.

Measuring and understanding the muon flux at SNO requires precise depth and density values. Two prominent effects of increasing rock density include increasing muon flux attenuation and a higher number of targets for interactions. The rock type and density are important in simulating muon propagation and observed flux as the largest uncertainties are from the interaction cross-sections.

The local geomagnetic environment is dependent on the latitude of SNO and determines the minimum momentum of cosmic ray primaries that can penetrate the field and reach the top of the atmosphere above the detector. These geomagnetic cutoffs depend on the

incoming direction and in no direction are the cutoffs above 40 GeV/c as calculated in [37]. The downward muons observed at SNO are generally produced with more than 1 TeV and therefore no correction needs to be applied to the cosmic ray flux here. Neutrino-induced muon events, however have lower energies and latitudinal dependencies. This is accounted for in our simulations as the input neutrino flux uses data that is specific to the latitude at SNO.

2.3 Detector

The detector cavity is roughly barrel-shaped with a height of 34 m and a radius at the equator of 11 m. A platform separates the top deck from the detector underneath. The deck is 6 m high and contains the electronic racks, computers and supports for the detector below. The lower portion of the cavity holds the physical detector and contains the H_2O and D_2O volumes, PMTs and support structures for both. Figure 2.3 shows a schematic of the detector.

A charged particle propagating through the detector at speeds higher than that of light in water will produce Cherenkov light. For the through-going muon analysis, both the D_2O and H_2O are the medium in which the Cherenkov light was produced. The difference in the refractive indices of D_2O and H_2O is at the half-percent level, so having the two distinct volumes of water is not very significant for the muon analysis. Cherenkov light produced in the water was detected by the array of PMTs. The signals from the PMTs travel up into the data acquisition (DAQ) electronics and are processed to form "an event". The following subsections will summarize the parts of the detector.

OUTER
H$_2$O

INNER
H$_2$O

D$_2$O

ACRYLIC
VESSEL

PSUP

17.80 M DIA.

PSfrag replacements

$\hat{q}$
$\hat{t}$
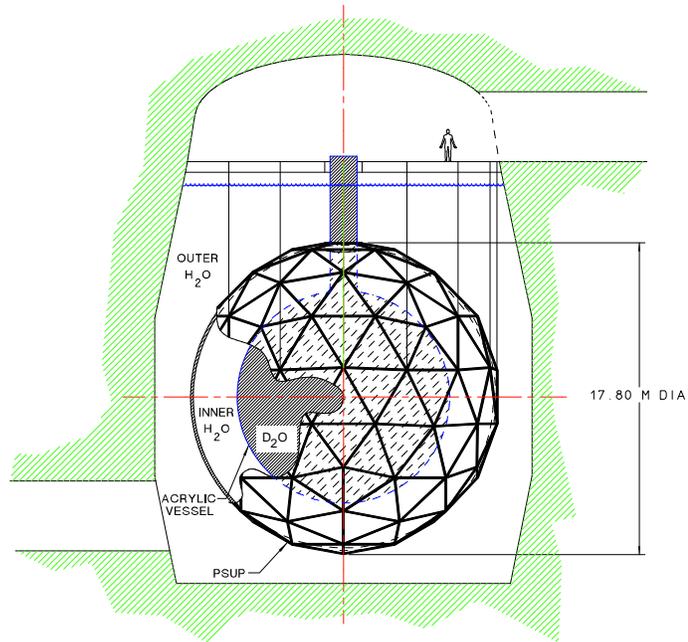
Figure 2.3: A schematic of SNO showing the cavity, support structure, acrylic vessel and water volumes from [34].

2.3.1   Water

Everything below the deck was submerged in water and can be separated into 3 regions as shown in figure 2.3. The innermost region had a spherical acrylic vessel (AV) 12 m in diameter, containing 1000 tonnes of D$_2$O. This was the target volume for the solar neutrino analysis. An 18 m diameter PMT support structure (PSUP) surrounded the AV and acted as a mounting surface for the inward looking PMTs. The region between the AV and PSUP was filled with H$_2$O and labeled "Inner H$_2$O" in figure 2.3. The PSUP was also a light and water barrier and separates the inner H$_2$O volume from the outer H$_2$O region.

The purpose of the 30 tonne AV was to isolate the D_2O from the surrounding H_2O. The 122 ultraviolet transmitting acrylic panels were bonded together to form the AV sphere. All but the 10 equatorial panels were roughly 5.5 cm thick. The equatorial panels were thicker (11.4 cm) since they contained the grooves for the ropes by which the entire AV was suspended. A 6.8 m long neck connected the deck and the top of the AV. The neck allowed for piping and calibration devices to be inserted into the D_2O volume.

There were two distinct water systems, one for the H_2O and the other for the D_2O. While the incoming potable H_2O was filtered and purified, ultrapure water leaches out soluble components when in contact with solid surfaces and can also support biological activity. The H_2O was continuously circulated to remove ions, organics and suspended solids. The D_2O was also recirculated through a reverse osmosis system to maintain its purity. On loan from Atomic Energy of Canada Limited, the D_2O was returned to the Ontario Hydro Bruce heavy water plant at the end of 2006.

Depending on the geometry of the event, Cherenkov light produced in the detector must pass through a number of combinations of D_2O, H_2O, and acrylic. Attenuation and scattering lengths in these media affect the total amount of light detected and the arrival times of those photons. The attenuation lengths are dependent on wavelength and are on the order of 50 m for water and 0.5 m for acrylic. Extensive optical calibrations were completed to estimate these values and are described below.

2.3.2   Photomultiplier Tubes

The photons from the Cherenkov light are detected by the PMTs. Figure 2.4 shows a schematic of a SNO PMT. The photocathode is a thin metallic coating on the inside of the front of the tube and is held at ground voltage. Electrons from the photocathode are freed by means of the photoelectric effect and directed towards the array of dynodes.
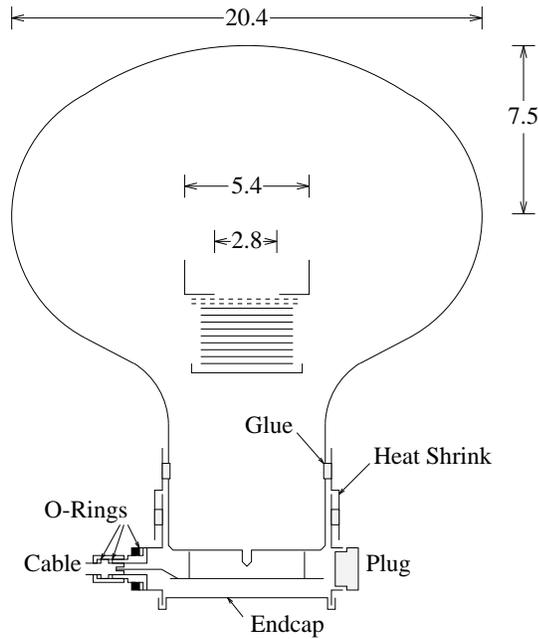
Figure 2.4: A SNO PMT and its waterproof enclosure. The solid horizontal lines are the dynode stack and the dashed lines make up the focusing grid. The dimensions are shown in centimetres.

Each successive dynode is at a higher voltage which increases the number of electrons by secondary emission. The final shower of electrons hit the anode and a signal is passed up into the DAQ system. The final dynode voltage is held roughly at 2 kV.

There are 9456 Hamamatsu R1408 PMTs of 20 cm diameter mounted on the PSUP looking in towards the D_2O that cover roughly 54% of 4π, an additional 91 PMTs that face outward from the sphere and 23 PMTs hanging from the deck. These face the outer H_2O and serve to detect background light originating in the outer region from radioactivity in the surrounding rock.

The PSUP is the geodesic sphere that functions as the main support system and the array of panels in which the PMTs are anchored. The housings the PMTs sit in hold both the PMTs and the surrounding light concentrators. The dielectric-coated aluminum

concentrators increase the effective area of the photocathode to maximize the number of detected photons and limit the incoming angle of photons so each PMT "looks" towards the centre of the detector. Flat panels hold 7 to 21 PMTs each and are placed into a repeating array within the PSUP structure. The entire sphere was suspended from bearings on the deck with 15 stainless steel wire rope cables.

The signals from each PMT feed through 32 m of custom coaxial cable into one of 19 SNO crates. A crate refers to the electronics and hardware responsible for processing the signals from 512 PMTs and supplying high voltage to the PMTs. Crates are housed in racks either alone or in pairs. Each rack has one high voltage supply, but the feed to each PMT goes through a resistor so that the high voltage cut from each PMT by removing the resistor.

The DAQ system processes, temporarily stores and digitizes the analog signals from the PMTs. The chain was required to provide timing to better than a nanosecond, charge measurements that cover a wide dynamic range and rate capability in excess of 1 kHz. The 4 MB of on-board memory allows for the buffering of large event bursts and the temporary storing of data.

To reduce the amount of data and processing, the SNO detector was designed to trigger on signals with an energy large enough for the event to be considered as a possible physics event. So while the PMTs fire at several hundred Hertz, events are only recorded above either a threshold in the number of hit PMTs or in overall charge. In the standard running mode, events with greater than roughly 20 PMTs hit trigger the detector. This threshold allowed the detector to be 100% efficient at detecting events in the D_2O with energies greater than 3.5 MeV [38]. For our analysis, with muon energies in the GeV range, losing events because of the trigger threshold is not a concern.

Overall event timing accurate to 100 ns was provided by a commercial GPS system with

a 10 MHz signal. This accuracy allows for a calculation of the effective livetime as well as correlations of SNO events with other detectors around the world. The surface GPS receiver communicated to the electronics underground through a fiber optic cable. A second, highly stable 50 MHz clock offering more precise inter-event timing was also used.

2.4 Calibrations

With the variety of media, hardware and electronics involved with SNO, many calibration procedures were necessary to describe the response of the detector. The few that are described here had the largest effect on the muon analysis. Through most of the lifetime of SNO, there was no direct calibration source for the muon analysis. The calibration of individual components of the detector were replied upon to help understand the properties of muon events.

The DAQ electronics were calibrated in time and charge. The charge offsets are the digitized values when there is no charge input from the PMT. Each channel was triggered to fire and the resulting charge output was recorded as the zero charge. The PMT timing is calibrated with timed pulses to each of the PMTs and the trigger. This measured the relative PMT differences in signal travel time and PMT response times.

Each PMT, though nominally identical to all the others, was distinct in properties like efficiency, turn-on voltage and gain. The PMT array was calibrated to account for PMT variations with an isotropic light source known as the "laserball". The laserball is a 10 cm diameter quartz sphere filled with 50 μm glass beads suspended in a silicon gel. A nitrogen laser up on the deck pulses 337.1 nm light into the ball through a set of fiber optic cables. Two neutral density filters are also in place to allow the intensity of light to be varied. The beads scatter the light and provides a light source with small and known anisotropy.

Since the properties of the source are known, individual PMT responses were recorded

PSfrag replacements
$\hat{q}$
i
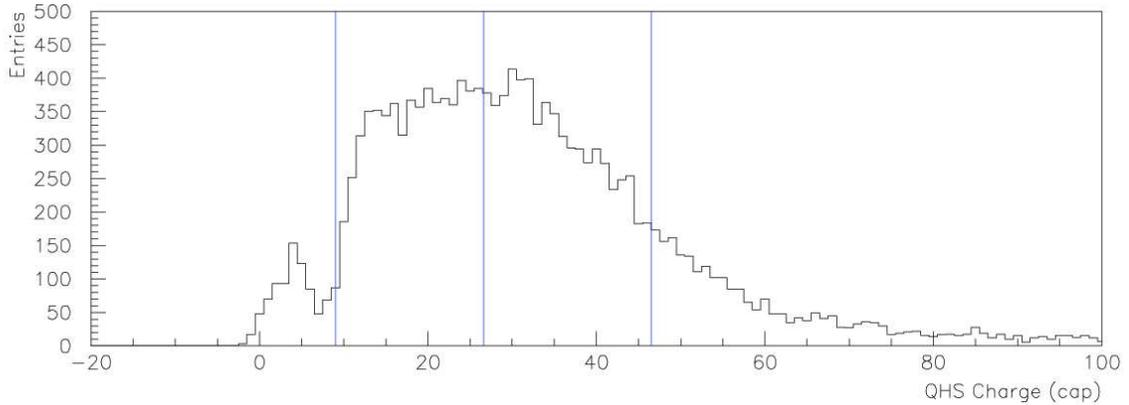
Figure 2.5: This shows a charge distribution for a single PMT. The vertical lines from left to right represent the discriminator threshold, peak and high half point. This figure was taken from [39].

and compensated for while analyzing the events. The standard configuration of the laserball has it positioned roughly in the centre of the detector, pulsing at 45 Hz for 70 minutes and at an intensity that attempts to ensure that any PMT only detects one photoelectron in each event.

The gain of a PMT was obtained by measuring the "high half point" of the non-Gaussian PMT charge distribution of laserball events. The high half point is defined as the point at which the charge distribution falls to half of the maximum value. Figure 2.5 shows the charge distribution and related terms for a single PMT. The charge measure used in this thesis is gain scaled units (GSU). This was obtained by dividing the offset corrected charge counts by the high half point of the channel. There is no strict conversion from GSU to photoelectrons since one photoelectron can produce a distribution of detected charges, but one GSU is roughly equal to 1.2 photoelectrons. This is also the reason for using GSU later on in the analysis as it is a well defined value.

A muon detection system was installed above the SNO detector in the NCD phase of

the experiment and had a livetime of 95 days. The system was designed to provide an external calibration source for the muon events. It consisted of a series of 128 single-wire chambers. The wires were arranged into four planes and the entire system was triggered by three scintillator panels. A muon passing through the system deposits energy in the scintillator and ionizes the gas in the wire chambers. The specific wires that were hit and the timing associated with the hits provided the information to reconstruct the muon path through the system. The system was able to track muon hits in the wires with roughly a 1 cm uncertainty. A sample of 30 muons detected both in the external muon system and SNO were used to check current muon track fitting algorithm. The results are described below in the systematic errors section.

2.5 SNOMAN

SNOMAN is the "SNO Monte carlo and ANalysis" program [34]. This software was developed to analyze and simulate data as well as significant backgrounds. A design priority for the SNOMAN code was that all routines should be as modular as possible so as to simplify distributed development. The key feature of this design is that of a central data structure and a set of independent processors that communicate with it. This keeps the processing flow decoupled from the data flow and makes changing existing processors and adding new ones simpler. SNOMAN controls these processors and the manner in which they interact with the data.

The processors that read or modify elements of the data structure are called software units. Each software unit has a single initialization and termination entry point as well as one or more execution entry points. These execution points can depend on other software units in a non-circular manner. These software units rely on a database of information about calibrations, detector status, and configuration geometry to analyze and interpret the data.

Examples of software units used in this thesis include ones to calculate the livetime and to fit a track to a muon traveling through the detector. The livetime processor accepts as input a list of data runs then references the database to calculate the length of each run and include any information stored that may reduce the livetime. The muon fitter reads in event information such as PMT charge and time and references the database for details like which PMTs are on or off and calibration constants specific to the data run. Both of these processes will be described in more detail in the following chapter.

A major function of SNOMAN is to control and process MC simulations of signals and backgrounds. The entire process of simulating a detected particle can be broken down into the generation of the particle, propagation, generation of the Cherenkov light and its detection in various PMTs. The basic units of the data structure for the MC are the event vertex and track. The vertex represents an event, namely the creation of a particle or any interaction. The track represents the movement of the particle from vertex to vertex. For the analysis the muon propagation MC in SNOMAN is most relevant. In addition to modeling the status of the detector and referencing calibration and geometry constants, the MC simulates the propagation of muons through SNO by accounting for discrete and continuous energy loss processes, capture and decay.

The way in which SNOMAN is designed allows for efficient and well-established and tested software packages to be used where possible. Three such packages were used for the muon analysis: `nuance` for the atmospheric neutrino production and propagation, MUSIC for the muons simulations and MCNP [40] for neutron propagation.

2.5.1 `nuance`

The `nuance` software package developed by Dave Casper [41], simulates neutrino interactions and propagation. It was originally designed to model atmospheric neutrino interactions for

large underground detectors searching for proton decay like Super-K. SNOMAN used this to simulate the atmospheric neutrino flux at SNO. While previous calculations took all particles and interactions to be linear 1-D processes, Version 3 of `nuance` used here implemented full 3-D interactions and the most current hadronic cross-section measurements.

The source of the simulated neutrinos was the Bartol group three-dimensional calculation of the atmospheric neutrino spectrum [42]. The calculation involves tracking the neutrino production and interactions in the atmosphere and Earth from the cosmic ray primary spectrum. The main goal in the most recent effort was to account for the effects on the flux detected at a particular site from the geomagnetic field. Previous 1-D calculations only account for particles originally traveling towards a detector since all particles are assumed to travel in the same direction. The 3-D simulation allows for secondary particles to deviate from the original primary direction. The effects are really only seen at lower energy with roughly a 3% excess in the 3-D flux below 1 GeV compared to the 1-D calculations. Above that energy the 1-D and 3-D calculations give the same results within statistics. The group provided atmospheric neutrino flux results from the 3-D simulation for SNO that were used here.

The simulated neutrinos were propagated through the Earth by tracking the neutrino-nuclear scattering processes from 50 MeV to 1 TeV. These include quasi-elastic channels, resonance, deep inelastic, coherent, diffractive and elastic scattering. The models, properties and constants for each process were taken from different sources. The model and cross-sections for quasi-elastic scattering are from Smith and Moniz [43]. Bodek and Yang [44] provide the information for deep inelastic cross-sections and the Rein-Sehgal model was used for resonant processes [45]. By summing the cross-sections and rates of all exclusive channels, and then adding (inclusive) deep-inelastic scattering inside appropriate kinematic limits, the total cross-sections and event rates are obtained.

If a neutrino produced a detectable muon at SNO, `nuance` passed the muon to another algorithm PROPMU, designed by Lipari [46]. In PROPMU, muons are propagated in rock based on survival tables for muons as a function of penetration depth. While PROPMU is not as detailed as the direct cosmic ray muon propagation code described below, its accuracy in describing the muon flux is at the 5% level and consistency checks between the two programs were performed by the authors of MUSIC [47]. Figure 2.6 shows the survival probabilities of muons in each of the simulations. There were 100 000 muons produced for each energy bin for each of the two depths. The shape of the distributions agree but MUSIC values were roughly 1-2% lower than PROPMU. The difference was attributed to updated bremsstrahlung cross-sections. The `nuance` was originally developed and tested with PROPMU as the muon propagation code. Since the results from PROPMU were similar to those from the more detailed MUSIC code, `nuance` was kept in its original form, with PROPMU as the muon propagation processor. For our calculations, only muons with an energy greater than 600 MeV were simulated since muons with energy less than that do not produce a through-going muon at SNO. The rock used in the simulation had a density of 2.85 g/cm^3 and $Z/A = 0.48$. Both agreed within errors of the measured values shown in table 2.2.

2.5.2 MUSIC

We have discussed the basic approach to simulating muon propagation and reasons for understanding the processes involved with muons traveling through rock. For the simulation of direct cosmic ray muons, we used the 3-D MC code MUSIC (MUon SImulation Code) [47]. These simulations are necessary since the direct cosmic ray muons are a measured signal as well as a background in the atmospheric neutrino analysis. Using the observed surface flux of direct cosmic ray muons and this propagation software we can produce a realistic dataset
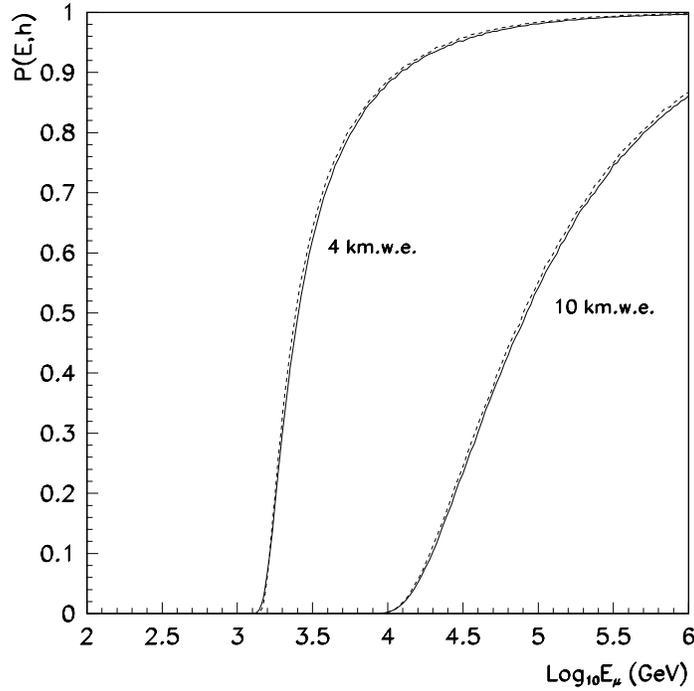
Figure 2.6: This shows a comparison of survival probabilities from MUSIC (solid line) and PROPMU (dashed line) through two different depths of rock. This figure was taken from [47].

of muon events. This allowed for the checking of the event rates and angular distribution of these events with those that were observed.

The accuracy of any particle propagation simulation is restricted by uncertainties in the interaction cross-sections and computational simplifications necessary to speed up the simulations. MUSIC accounts for all high energy muon interactions and uses the most recent and accurate cross-sections available. The general method is to propagate a muon through a layer of rock of known chemistry and density while accounting for the energy loss processes as in equation 1.7. The energy loss due to ionization is treated as a continuous process while

the other interactions are treated discretely. The magnitude of the other processes scale with energy and are bremsstrahlung, pair production and muon-nucleus inelastic scattering.

Starting with the Gaisser parameterization of the surface muon flux (equation 1.5), the code takes a muon of known initial energy, starting point, and direction then simulates muon interaction points spaced according to the mean free path. The energy loss processes are applied and the lateral and angular displacements are evaluated. If the muon made it through the layer of rock, the final energy, direction, and arrival point are stored. The results from the simulations are provided in the next chapter.

Chapter 3

Event Reconstruction

The details of a muon event seen at SNO were used to reconstruct the event through the detector and estimate the position the muon entered the detector and the direction it was traveling.. Here I present a description of the properties of observed muon events and the event fitting algorithm.

3.1   Muon Events

Muons traveling faster than the speed of light in water will produce Cherenkov light continuously along the track in a cone opening up in the direction of motion. This light is detected by the PMTs with a distinct pattern. A cone intersecting a sphere produces a ring and since the majority of muons seen at SNO have enough energy to travel all the way through the detector, the ring of hit PMTs fills in. If the muon stops in the detector the hit pattern will only be partially filled in. Figure 3.1 schematically shows a muon traveling through the detector.

Figures 3.2, 3.3 and 3.4 show the same simulated muon on the standard event display used at SNO (XSNOED). The flat map versions of the muon in figures 3.3 and 3.4 usually make it simpler to view and understand the muon event. This particular muon travels approximately downward and through the centre of the detector. Both figures are of the same muon, with the first showing the charge distribution and the second showing the time
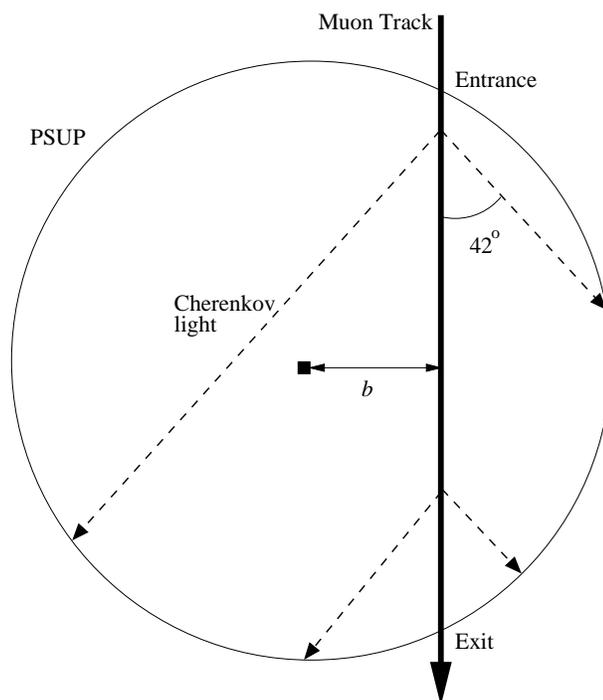
Figure 3.1: A schematic of a through-going muon event at SNO. The muon track is the thick vertical line and the circle represents the PSUP. Cherenkov light is produced continuously along the track with an opening angle of roughly $42°$. The impact parameter is labeled b.

distribution of the PMTs. Each coloured point represents a PMT. Highest charges detected by the PMTs are those near the exit point and are shown in figure 3.3 as the over-scaled pink points. Muons deposit energy in the detector at a constant rate. When the muon is near the exit point the energy is deposited into a smaller number of PMTs and therefore each PMT near the exit point records a relatively high charge. Most of the PMTs in the top portion of the detector were hit by reflected light and so they register smaller charges and later time values. The timing and colouring scheme of figure 3.4 has a low resolution, but the division of direct (prompt) light and reflected late light is apparent. The first PMTs hit are generally those near the entrance point of the muon. While for this muon, the entrance
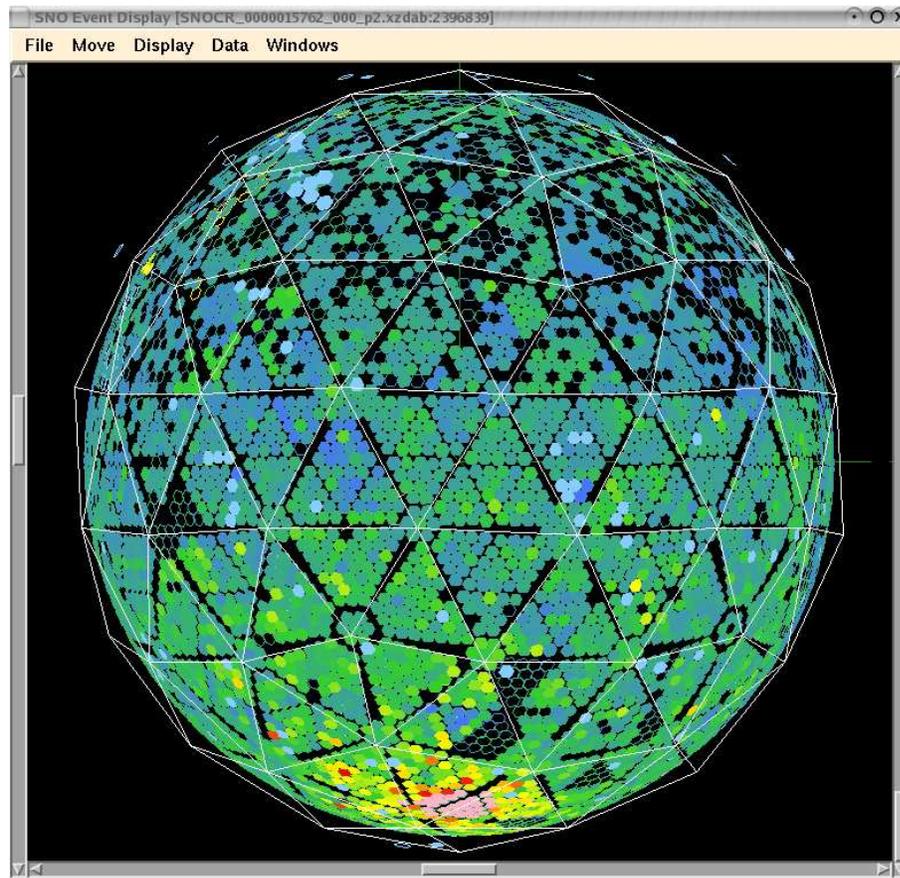
Figure 3.2: A through-going MC Muon XSNOED Sphere Charge Map. Each coloured point represents a PMT. Blue coloured points represent low charge while red and over-scaled pink represent high charge. The high charge group of PMTs near the bottom is the exit point of the muon.

point is in the upper portion of the detector filled with reflected late light, we will show below that early PMTs can still trace the entrance point well.

The PMT timing and charge values provide information used to reconstruct the muon track. While the entrance point and exit point are relatively easy to define, the energy of a through going muon is not a property we can extract. Muons of different energy traveling the same track through SNO, will deposit roughly the same amount of energy in the detector.

47
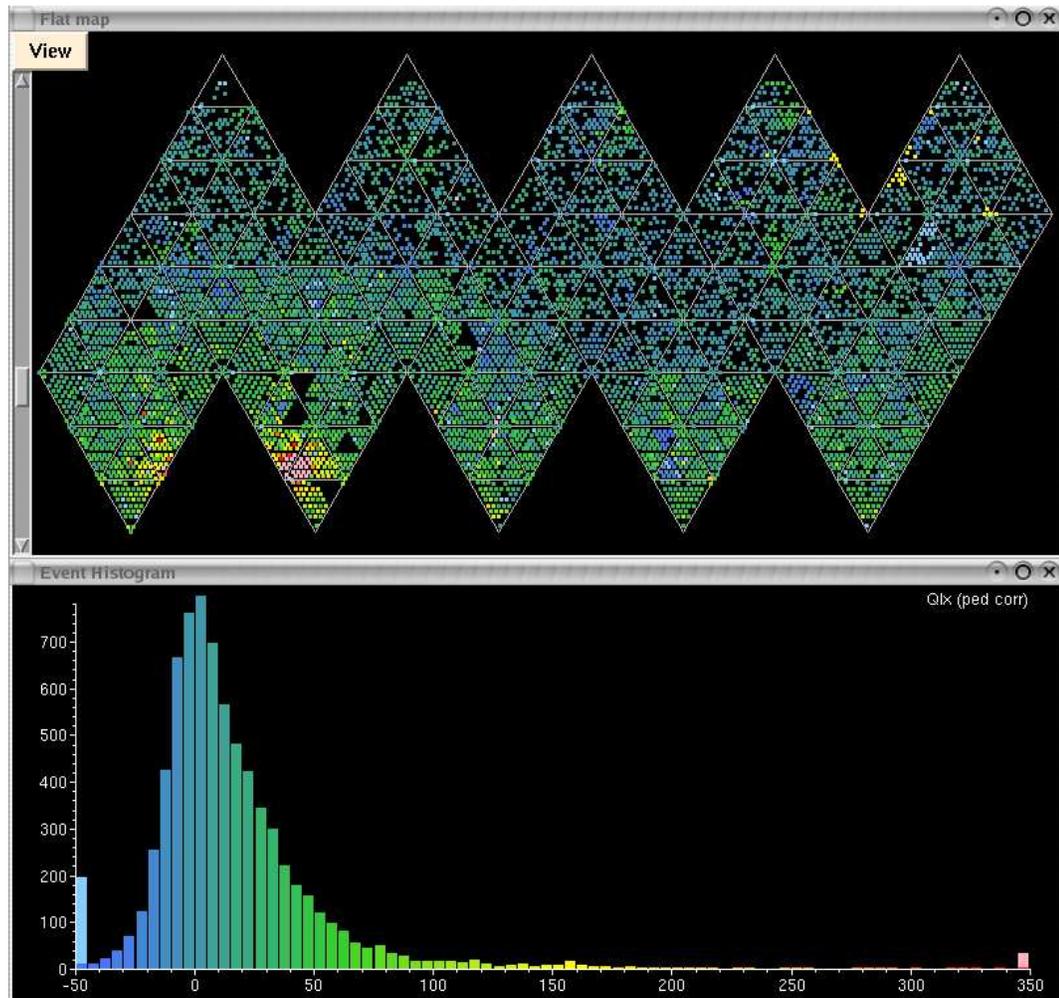
Figure 3.3: A through-going MC Muon XSNOED Flat Charge Map. Each coloured point represents a PMT. The histogram below shows the charge distribution. The high charge group of PMTs near the bottom is the exit point of the muon.

In a perfectly spherical detector, the PMT hit pattern as well as the distribution of PMT charge and time measurements depend only on the impact parameter. This is the minimum distance from a point on the muon track to the centre of the detector. Regardless of the track direction and orientation, muons with the same impact parameter will produce events
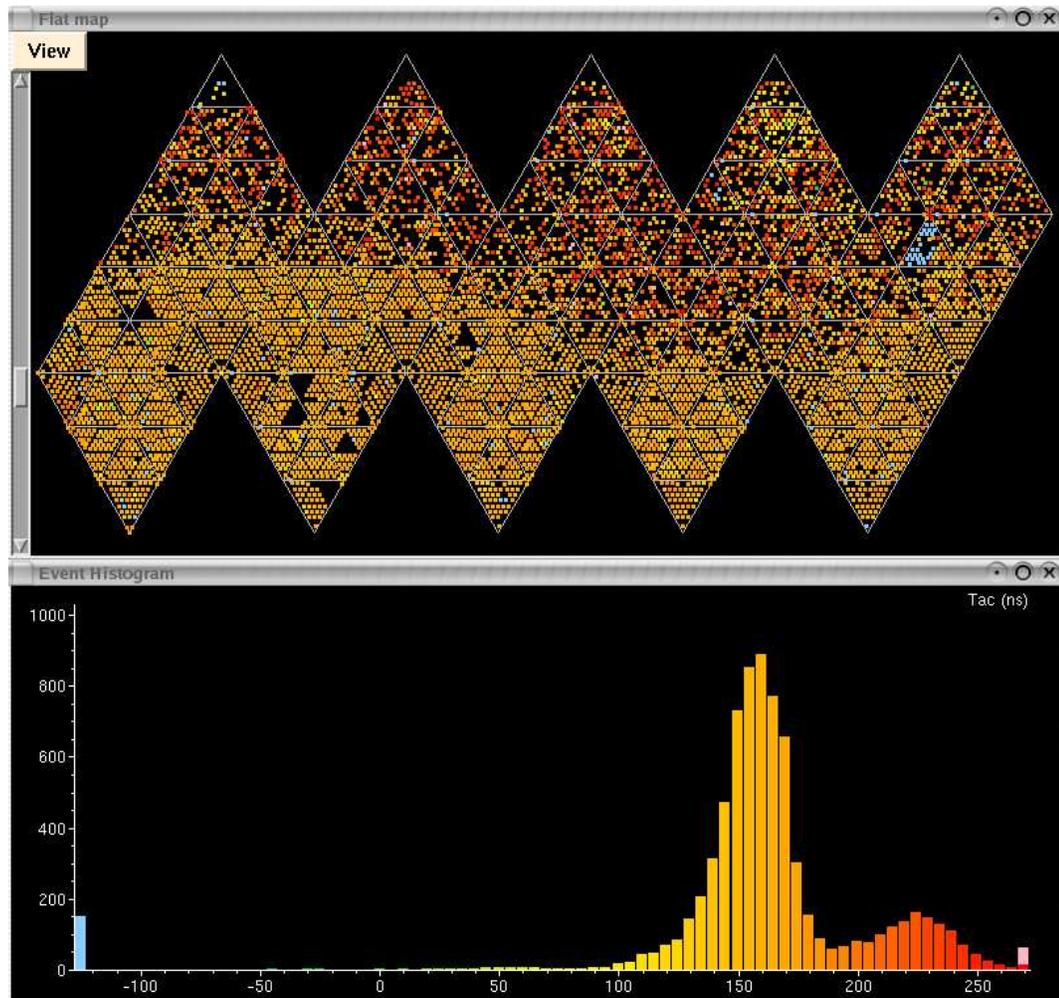
Figure 3.4: A through-going MC Muon XSNOED Flat Timing Map. Each coloured point represents a PMT. The histogram below shows the time distribution. The PMTs hit by late scattered light are shown in red.

with similar characteristics.

There are a number of different parameterizations that describe a line intersecting a sphere of known radius; a description requires four parameters. One example is to specify the entrance and exit positions in polar co-ordinates. The parametrization used in the

49

reconstruction of muon tracks in this thesis defines a special frame in which the muon track is fully described by the impact parameter and then finds the rotation matrix required to take the event into that frame. The equation of motion for a muon entering the detector at time $t = 0$ from above with the positive z direction pointing from the centre of the detector towards the zenith, can be written as

$$\vec{x}(t) = b\hat{x} + (\sqrt{R^2 - b^2} - ct)\hat{z} \tag{3.1}$$

where R is the radius of the detector, b is the impact parameter and the origin is at the centre of the detector.

The following section describes two techniques that utilize the spherical symmetry of SNO to reconstruct through-going muons. The first is a fast and simple algorithm that uses data from only a subset of the hit PMTs in an event to provide a seed for the more detailed second processor. The second routine reconstructs events by maximizing a likelihood function calculated with information from all the PMTs. The routines are referenced by their three-letter SNOMAN processor codes, FTR and FTI.

3.2   Reconstruction - FTR

This fitter reduces the muon event geometry to two unit vectors ($\hat{t}$ and $\hat{q}$) that point roughly to entrance and exit positions from the centre of the detector. The unit vectors are later used to rotate the event into a frame in which FTR has a standard prediction for the muon track parameters. There is a correctable systematic shift in the vectors which allows one to extract a better estimate of the entrance and exit positions. Both vectors have a bias and a statistical spread about the true positions.

The muon entrance position is found by identifying the 80 earliest hit tubes. The routine
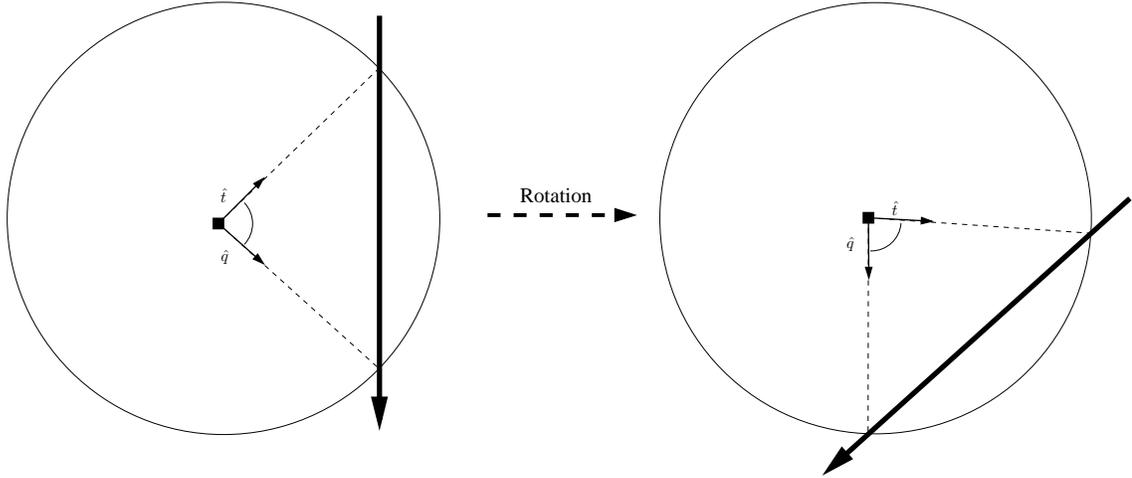
Figure 3.5: This shows the definitions of the charge and time vectors $(\hat{q}, \hat{t})$ and the rotation taken to have the charge vector pointing directly downwards. The angle between $\hat{q}$ and $\hat{t}$ is α_{QT}. The muon track is the thick line. The routine performs this in 3 dimensions.

then finds the median position of these PMTs. From these 80 tubes, the earliest tubes within $26°$ of the median position up to a maximum of 20 tubes were selected. The vector from the centre of the detector to the median of the 20 PMTs is $\hat{t}$. In simulations, the mean angle between $\hat{t}$ and a vector pointing from the detector centre to the true entrance is $6.5°$.

An estimate of the exit position of the muon was found by calculating the charge weighted average position of all the hit PMTs and $\hat{q}$ points from the centre of the detector to that position.

$$\hat{q} = \frac{\sum_{PMT}(q_{PMT} \times \vec{x}_{PMT})}{|| \sum_{PMT}(q_{PMT} \times \vec{x}_{PMT}) ||} \tag{3.2}$$

The mean angle between the true muon exit position and $\hat{q}$ was found to be $\sim 5.2°$ in simulations.

Once $\hat{q}$ and $\hat{t}$ have been calculated, the fitter rotates the vectors such that $\hat{q}$ points in the $-\hat{z}_R$ direction and $\hat{q}$ in the $x_R z_R$ plane with positive x_R. The subscript R refers to the

51
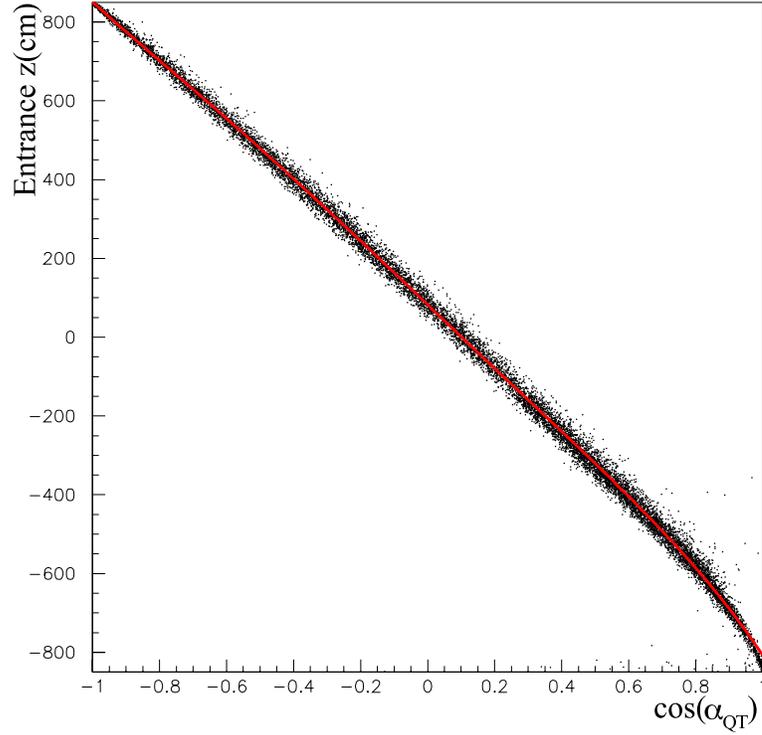
Figure 3.6: This shows the relation between the z coordinate in the rotated frame of the entrance position and $\cos \alpha_{QT}$ for simulated muons. In this frame the exit position is near $z \simeq -850$ cm. A fifth order polynomial fit is shown in red. [48]

rotated frame. In this rotated frame, we can see that the cosine of the angle between $\hat{q}$ and $\hat{t}$ is related to the impact parameter of the muon. If the angle between $\hat{q}$ and $\hat{t}$ (α_{QT}) was zero ($\cos \alpha_{QT} = 1$), and $\hat{q}$ and $\hat{t}$ corresponded exactly to the exit and entrance positions, the muon had an impact parameter of 850 cm and just clipped the edge of the detector. If on the other hand, the angle between the vectors was $180°$ ($\cos \alpha_{QT} = -1$), the muon would have had an impact parameter of zero and gone straight through the centre of the detector.

If $\hat{q}$ and $\hat{t}$ pointed to the true positions, the relationship between the angle and the

impact parameter would be

$$\cos \alpha_{QT} = 2 \left(\frac{b}{R} \right)^2 - 1 \qquad (3.3)$$

where b is the impact parameter and R the radius of the detector.

Simulating muons allowed us to relate $\cos \alpha_{QT}$ and the muon entrance position as well as $\cos \alpha_{QT}$ and the muon direction. The distribution of entrance positions for simulated muons in the FTR frame is show in figure 3.6. A 5th order polynomial fit is applied to obtain a functional form of the relation. A scatter plot showing muon directions in the rotated frame is shown in figure 3.7 with an 11th order polynomial fit. The fits were used to estimate the entrance and direction of the muon from $\cos \alpha_{QT}$.

The values of $\hat{q}$ and $\hat{t}$ may be easily calculated for any event, and with them it is possible to find the rotation matrix to take the event into the frame described above. The FTR routine then calculates $\cos \alpha_{QT}$ and uses it to estimate the expected muon entrance position and direction using the polynomials. By then applying the inverse rotation matrix to the estimated track in the FTR frame, the fitter obtains the entrance position and direction in the SNO detector co-ordinates. The calculated entrance points and directions are passed as a seed for the maximum likelihood technique used in FTI.

3.3   Reconstruction - FTI

The FTI routine relies on all the available information in a muon event to fit for a track. For a given muon direction, entrance position and entrance time, the routine calculates the probability that each PMT would be hit. If a tube is hit, the routine estimates the probability that it would see the charge and time that was recorded. These probabilities were calculated from lookup tables generated from thousands of simulated muons. The
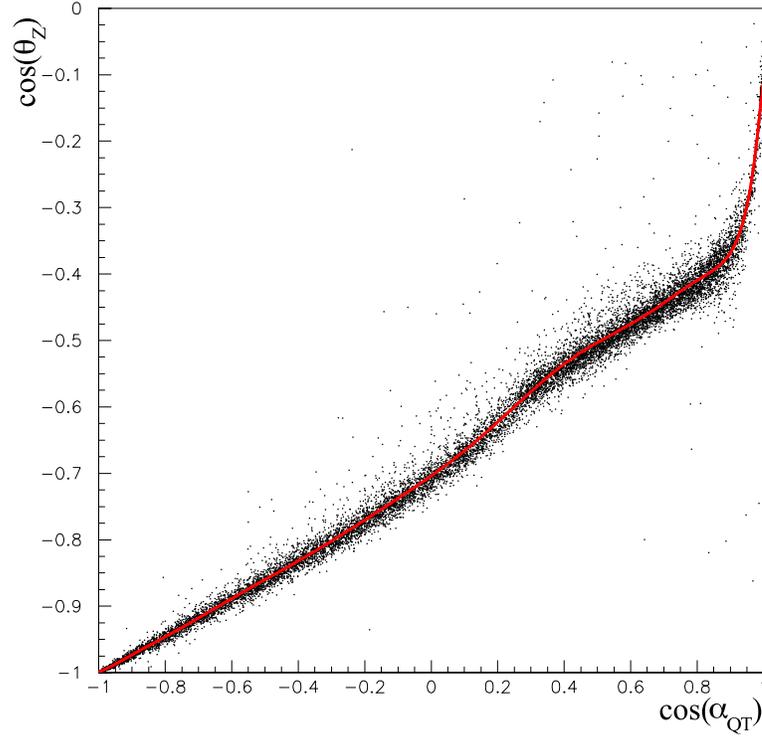
Figure 3.7: This shows the relation between the muon direction and $\cos \alpha_{QT}$. An eleventh order polynomial fit is shown in red.[48]

following subsections present the derivations of the likelihood, the lookup tables created to determine the expected number of detected photoelectrons, and the minimization scheme used in the final fit.

3.3.1 Likelihood Function

We first look at developing the form of the probability a PMT records a hit with the observed charge and time. For any given muon track, the expected number of detected photoelectrons for a specific PMT varies from much less than one up to several hundred. If the expected number of detected photoelectrons is λ then the probability that it would

observe n photoelectrons is a Poisson distribution with

$$P_N(n \mid \lambda) = \frac{e^{-\lambda}}{n!}\lambda^n. \tag{3.4}$$

The probability a PMT records a hit with the observed charge and time is then given as

$$P_N(hit, Q, t \mid \lambda) = \sum_{n=1}^{\infty} P_N(n \mid \lambda) P_Q(Q \mid n) P_T(t \mid n) \tag{3.5}$$

where the probability of a PMT observing a charge Q given n photoelectrons ($P_Q(Q \mid n)$) and the probability of a PMT firing at time t given n photoelectrons ($P_T(t \mid n)$) are defined below. The overall likelihood of the observed PMT hits, times and charges for a given muon track is the product of the probabilities for the individual PMTs.

3.3.2 Probability of Detecting a Charge Given n p.e.

Finding $P_Q(Q \mid n)$ requires simulations, as the charge response of a PMT to a single photon hit is a distribution and not a single value. The distribution has a broad peak and a long tail. The single photoelectron charge spectrum was sampled to produce charge spectra for PMTs hit by n photoelectrons. The charge spectrum for 7 and 20 detected photoelectrons are shown in figure 3.8. The fits allow one to assign probabilities to the hypotheses that a phototube which observes charge q detected n photoelectrons.

The SNOMAN PMT MC was used to generate 100 simulated single photoelectron hits on every online tube. The charges reported by the MC were binned and the resulting histogram was used to produce the cumulative distribution for the probability of observing a charge Q, given a single detected photoelectron. The distribution was sampled to produce the charge histograms for n detected photoelectrons, with $1 \leq n \leq 100$. Above a charge of 10 photoelectrons, the shape of the charge distribution was parameterized by a two-
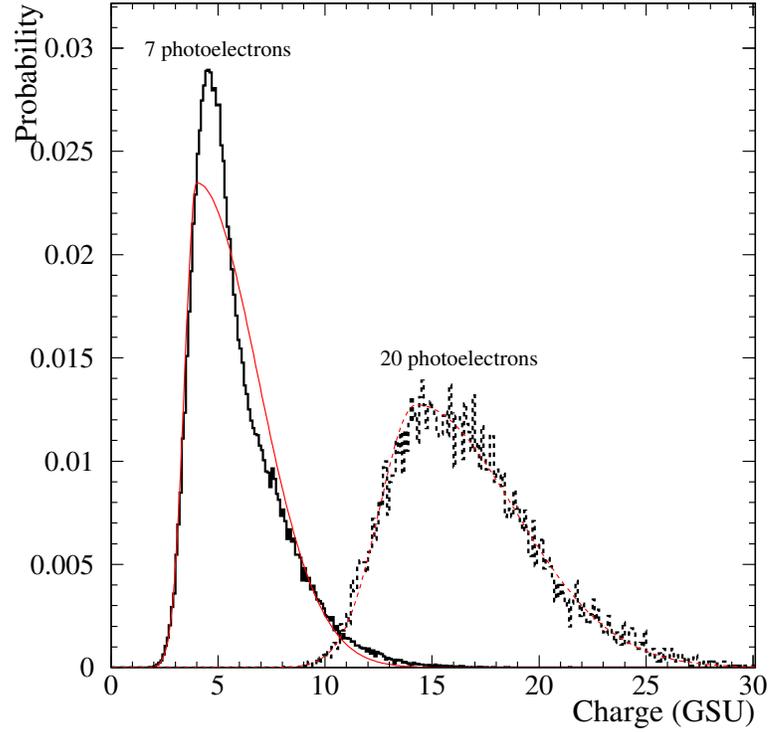
Figure 3.8: These are charge distributions for MC PMTs which detect 7 or 20 photoelectrons. The histograms are shown with a fit to the double-sided Gaussian described by equation 3.6. The 7 p.e. distribution shows the poor fit to the Gaussian and is the reason for the parameterization not being used below 10 photoelectrons. [48]

sided Gaussian function given in equation 3.6. Below a charge of 10 photoelectrons the parameterized fit was poor, so the charge histograms were normalized to produce charge probability distribution functions (PDFs) which provided the required probabilities.

$$P(Q \mid n) = \begin{cases} N \exp\left(\frac{-(Q-\bar{Q})^2}{2\sigma_1^2}\right) & Q < \bar{Q} \\ N \exp\left(\frac{-(Q-\bar{Q})^2}{2\sigma_2^2}\right) & Q > \bar{Q} \end{cases} \tag{3.6}$$

In equation 3.6, N is the normalization, $\bar{Q}$ is the mode of the charge distribution, σ_1

and σ_2 are the left and right RMS of the two-sided Gaussian (equations in 3.7).

$$N = \frac{\sqrt{2}(\text{entries} \times \text{binsize})}{\sqrt{\pi}(\sigma_1 + \sigma_2)}$$

$$\bar{Q} = -2.3531 + 0.83075n$$

$$\sigma_1 = -1.7565 + 0.7807\sqrt{n} \tag{3.7}$$

$$\sigma_2 = 1.1351 + 0.77393\sqrt{n}$$

3.3.3 Probability of Recording a Time Given n p.e.

We now look at the probability of a PMT firing at time t given n photoelectrons ($P_T(t \mid n)$). The distribution of observed PMT times depends on photon propagation from the source to the PMTs and the PMT timing responses to single photoelectrons.

In a simplistic model of PMT timing, the observed time distribution would be a delta function at a time $t = \frac{d}{c} \times n$ after the emission, where d is the distance from source to the PMT, c is the speed of light in a vacuum and n is the index of refraction. However, photons also scatter, increasing the travel distance. In addition, the photon arrival time distribution was dominated by the intrinsic transit time distribution of the SNO PMTs.

For a muon event in SNO the distance from the source to a PMT is different for each PMT. To make comparisons of the PMT times and understand what timing distribution is expected for a muon event a "time residual" must be defined. The time residual used here is the observed PMT time less the propagation time of the photon and muon. The PMT time residual is then,

$$t_{res} = t_{PMT} - t_0 - \frac{d_1}{c} - \frac{d_2}{21.8 \text{ cm/ns}} \tag{3.8}$$
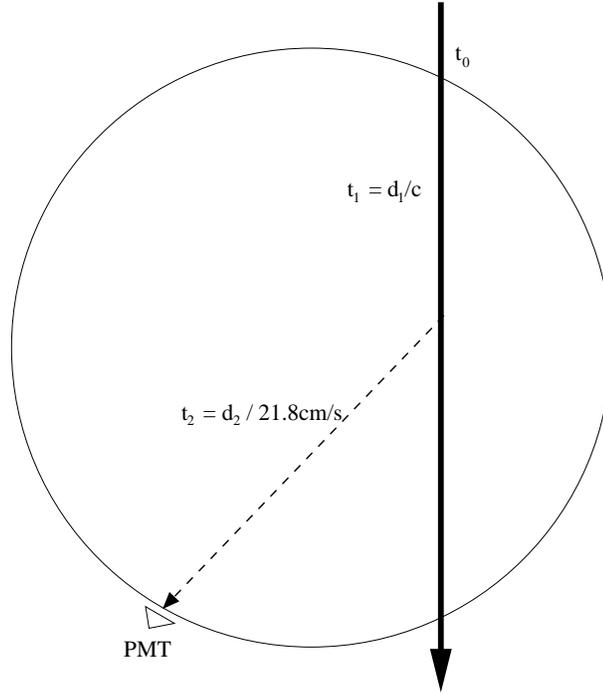
Figure 3.9: This is a schematic defining the elements of the time residual. The muon enters at t_0 and takes t_1 to travel before it produces the Cherenkov light seen by the PMT shown. The time t_2 is the travel time of the light from production to detection.

where t_{PMT} is the PMT hit time, t_0 is the global fit offset time, d_1 is the distance the muon traveled before emitting the Cherenkov photon, c is the speed of light in a vacuum, d_2 is the distance that the Cherenkov photon is expected to have propagated to reach the PMT and 21.8 cm/ns is the speed of light in water. Figure 3.9 shows these values schematically.

With a PMT time residual, we can look at what the probability that a PMT hit by direct light from a muon will be to observe a particular time residual. Figure 3.10 shows the log and linear time residual histograms for MC PMTs. The log plot shows the prepulse and late light effects more clearly. Prepulsing occurs when a photon is not absorbed by photocathode, but ejects an electron directly from the first dynode. The frequency of this effect is roughly
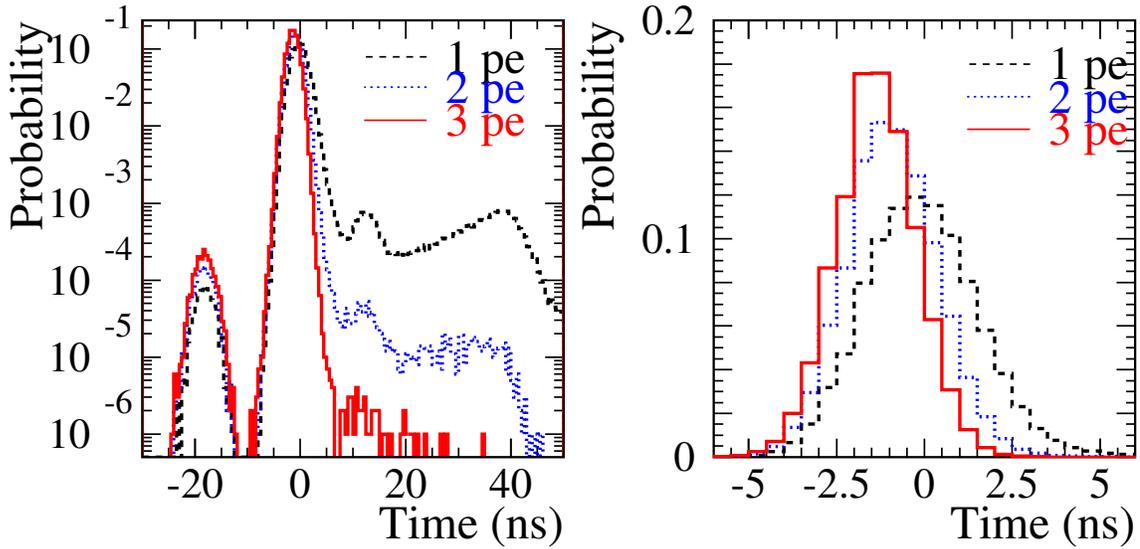
Figure 3.10: The time residual (t_{res}) plot for MC PMTs. The right plot is a linear version of the left. The prepulse peak is seen around -16 ns and the prompt light peak at 0 ns. These plots also show the effect of multiple photoelectrons on the timing distribution [48].

1 in 1500 hits and on average causes a pulse 20 ns earlier than if the electron was ejected from the photocathode. The late light peak is not fully understood. It is most likely reflections occurring either from the light concentrators or the faces of the PMTs. Figure 3.10 also shows the effect of multiple photoelectrons striking a PMT. The distributions for multiple photoelectrons shift to earlier times. Consider the timing distribution for a single photoelectron. To build a distribution for the case of two photoelectrons, one can sample twice from the single photoelectron distribution. PMTs fire on the first signal, so the time of the later photoelectron has no effect on the resulting PMT time residual. This makes multiple photoelectron hits shift the PMT time residual distribution to earlier times. FTI uses three probability distribution functions to describe the different parts of the single photoelectron timing distribution.

Prepulse and prompt light hits are modeled by a Gaussian distributions with mean times

of -15.5 ns and 0.0 ns and widths of 2.2 ns and 1.6 ns respectively [49]. The residual time window for prompt light is from -25 ns to 5 ns.

Early and late hits are modeled with a Heaviside function. The function transitions at 1 ns, is normalized to 1 over -150 ns $< t <$ 200 ns, and the value of the function on either side of the 1 ns is such that the probability of hits before 1 ns corresponds to a dark current of roughly 500 Hz. This roughly accounts for effects such as the PMT dark current scattering, reflections and the late part of the transit time distribution.

For a given PMT hit, a probability needs to be assigned to each of the PDFs. The lookup tables described in the next section provide the probability of whether or not a tube was hit by prompt light as defined above. FTI can also use this information to also determine the probability, α, that a detected photon was a late hit (not prompt). For hit PMTs with a time residual later than 1 ns the probability assigned to being a late hit is α^n, where n is the number of detected photoelectrons.

With these PDFs the probability that a tube hit by a single photoelectron would record time t is

$$
\begin{aligned}
P(t \mid 1 \text{ p.e.}) &= (1-\alpha)\bigg(P(t \mid \text{prepulse}) + P(t \mid \text{prompt})\bigg) \\
&\quad + P(t \mid \text{early or late}) \qquad\qquad (3.9) \\
P(t \mid 1 \text{ p.e.}) &= (1-\alpha)\bigg(\rho \frac{1}{\sqrt{2\pi}\sigma_{pp}} e^{\frac{-(t+15.5)^2}{2\sigma_{pp}^2}} + (1-\rho)\frac{1}{\sqrt{2\pi}\sigma_{pr}} e^{\frac{-t^2}{2\sigma_{pr}^2}}\bigg) \\
&\quad + \begin{cases} \frac{0.675}{9600}/150 \text{ ns} & t < 1 \text{ ns} \\[2mm] \alpha(1-\frac{0.675}{9600})/200 \text{ ns} & t > 1 \text{ ns} \end{cases} \qquad (3.10)
\end{aligned}
$$

where $\sigma_{pp} = 2.2$ ns , $\sigma_{pr} = 1.6$ ns and ρ is the fraction of the prompt hits that result in a prepulse. Matching the parts of the equation to the left plot in figure 3.10, we have the prepulse peak at -15.5 ns, prompt light peak at 0 ns and the late light portion covering

the distribution to the right of the prompt peak.

To generalize the probability to n detected photoelectrons we need to account for the probability of late hits given multiple photoelectron hits and the frequency with which a photon produces a prepulse rather than a prompt pulse. Since the probability of all the hits being late is α^n, the probability of at least one prompt hit is $(1-\alpha^n)$. The prepulse frequency depends on the number of prompt photoelectrons, not the total number of photoelectrons since only the prompt photoelectrons can produce a pulse before the prompt peak. The expected number of prompt photoelectrons is $\tilde{n} = (1 - \alpha)(n)$. A PMT hit by $\tilde{n}$ prompt photoelectrons has a probability $(1 - \rho)^{\tilde{n}}$ of no prompt hits producing a prepulse. The probability that at least one of the prompt photoelectrons produces a prepulse is

$$\chi = 1 - (1 - \rho)^{\tilde{n}}. \tag{3.11}$$

The early (dark current) hit probability is independent of the number of photoelectrons detected, so there is no change to the early part of the Heaviside function.

Combining these arguments we write the probability of recording a time t, given n photoelectrons as

$$
\begin{aligned}
P(t \mid n) \;=\; & (1 - \alpha^n)\left(\chi \frac{1}{\sqrt{2\pi}\sigma_{pp}} e^{\frac{-(t+15.5)^2}{2\sigma_{pp}^2}} + (1 - \chi)\frac{1}{\sqrt{2\pi}\sigma_{pr}} e^{\frac{-t^2}{2\sigma_{pr}^2}}\right) \\
& + \begin{cases} \frac{0.675}{9600}/150 \text{ ns} & t<1 \text{ ns} \\[2mm] \alpha^n(1 - \frac{0.675}{9600})/200 \text{ ns} & t>1 \text{ ns} \end{cases}
\end{aligned}
\tag{3.12}
$$

61

3.3.4 Lookup Tables

A set of lookup tables was produced that characterized muon events as a function of impact parameter. The tables provide the expected number of detected photoelectrons for each PMT for a given muon track. They also include information on the fraction of detected photoelectrons that will be prompt.

The tables store muon information in a frame of reference where equation 3.1 defines the muon position. Each bin in the table is one of 10000 2-D bins in polar angle ($\cos\theta$) and azimuthal angle (ϕ) that cover a hemisphere. In this frame the muon pattern is symmetric under reflection through the x-z plane, so only one hemisphere is needed. The bins are equally spaced in $\cos\theta$ but they are very narrowly spaced near $\phi = 0$ and become larger as $\phi \to \pi$. The ϕ steps were generated through rotations assuming cylindrical symmetry. The uneven spacing accounts for the fact that the pattern of PMT hits from a muon changes most rapidly near the exit point. The exit point is always at $\phi = 0$ but is at a different position in $\cos\theta$ for each impact parameter.

To generate the table for a particular impact parameter, hundreds of muons were simulated. The bin entry includes the fraction of all detected photoelectrons that a PMT at the centre of the bin would observe averaged over all possible muon orientations. A prediction for the number of photoelectrons that a PMT located at the bin centre is expected to detect is found by multiplying the bin entry by the number of detected photoelectrons in an average event. The fraction of prompt hits was also stored for the timing probability distributions described above. Muons were simulated at 108 distinct impact parameter values.

The last steps in predicting the expected number of detected photoelectrons for any PMT are two interpolations. The first is a bilinear interpolation between the grid of PMT positions to obtain table entries for positions of a PMT in the new rotated frame. For each PMT, FTI finds the $\cos\theta$ and ϕ location of the PMT in the reference frame of the tables

and then interpolates between the four bin centroids nearest to the centre of the PMT. The second interpolation is applied to obtain the table values for any impact parameter as only 108 discrete impact parameters are available. While the distribution in the number of detected photoelectrons generally changes slowly with impact parameter, near the exit point of the muon the distribution changes very rapidly. FTI adopts an interpolation scheme where a linear scheme is used if the predicted number of photoelectrons is less than 10, and a cubic spline interpolation if that number is greater than 10. The charge distributions in 3.6 transitioned from poorly parametrized below 10 photoelectrons to well parametrized above and was used here as a point to move to a more complex interpolation in more quickly changing areas (the muon exit point).

3.3.5   Likelihood Maximization

The muon entrance position and direction, entrance time and the number of detected photoelectrons are the 6 parameters in the FTI fit. The routine references the lookup tables to obtain the expected number of detected photoelectrons, λ, which in turn gives the probability the PMT fired and the associated charge distribution. The estimate of the observed charge is calculated by combining λ with the charge distributions for n detected photoelectrons. The tables also provide the prompt fraction, which when combined with the prepulse, prompt and late or early light timing PDFs gives the probability that the PMT would observe the time that it did. The likelihood function used by FTI is simply the product of all these probabilities.

$$\mathcal{L} = \left(\prod_{\mathrm{nothitPMTs}} e^{-\lambda} \right) \left(\prod_{\mathrm{hitPMTs}} \sum_{n=1}^{\infty} \frac{e^{-\lambda}}{n!} \lambda^n P(Q \mid n) P(t \mid \vec{x}_{\mathrm{PMT}}, n) \right) \qquad (3.13)$$

The FTI routine estimates the muon track parameters by finding the combination of track parameters that minimizes the negative logarithm of the likelihood function. This

minimization routine is a simulated annealing method using a downhill simplex method [48].

The downhill simplex method modifies a set of function arguments until the function value for each set of arguments agrees within some tolerance. As each function value is being modified to prefer lower values, their convergence suggests a minimum. An $n + 1$ by n array of function arguments is the simplex, where n is the number of function arguments. Its simplest form has the routine selecting and replacing the point with the largest (worst) functional value. The new point is chosen by reflecting, extending or contracting the position in argument space through the n-D plane formed by the other simplex entries. If the new value of the function is still the worst compared to the others in the simplex, the routine tries again. If the new value is better, the new value replaces the old point and the process continues with the current worst point. There are known problems with this simple method as the algorithm tends to get stuck in local minima.

Simulated annealing is a modification of the downhill simplex method that is better at finding global minima in function spaces with many local minima. Instead of always discarding the worst point in the simplex, the routine may instead, with some probability, replace one of the other points. The probability of discarding another point is determined by user-defined "temperature" stages. The routine starts at "high temperature" where the discard probability is high and steps to "zero temperature" where the probability becomes zero. The ability to search far from the starting seed or initial entries in the high temperature phases is the strength of the algorithm. This is what allows the routine to find global minima in rough function spaces.

The first step in this method is to seed the simplex. The starting simplex is the seed fit and perturbations of that fit in each of the parameters. In our 6 parameter fit, we require 7 entries in the simplex. The first entry consists of the fitted values from FTR as the starting

seed for the entrance position, direction and time. The initial value for the expected total number of detected photoelectrons is estimated to be 1.16 times the total charge measured in the event in gain scaled units. The following 6 entries are then generated by making minor perturbations of this first entry. The perturbations take place in the special frame where the equation of motion is given by equation 3.1. In this frame, the entries are generated by varying the impact parameter and track rotations about each of the 3 axes.

The routine progresses linearly through 8 temperature steps with 30 new fit attempts at each temperature. After this initial phase, the temperature is raised back up, and decreased in 20 steps of 15 fit attempts. If the simplex points have a negative log likelihood within a factor of 10^{-4} of each other the routine ends. Once the first pass has finished the entire procedure is repeated using the results from the as a seed for the second pass, rather than the FTR fit parameters.

The fitter was developed in the D_2O phase of SNO. While the addition of salt to the D_2O volume was not expected to have an impact on this reconstruction algorithm, the NCDs in the last phase of SNO were expected to block and scatter a small but detectable amount of light. The performance of the fitter through all of the phases is described in the next section.

Chapter 4

Simulated Muon Events

The development of this analysis relied on Monte Carlo simulations of muon events. The MC was used to build the event fitter mentioned in the previous chapter and define cuts used to distinguish background events from signal events. Here, the simulated datasets and cuts used in the final analysis are described.

4.1 Monte Carlo Datasets

Simulated muons allowed us to develop and check analysis tools like the fitter mentioned above, provided estimates of expected muon events that we used to build probability distributions and were our best calibrating source. Through most the development of this analysis there was no physical external check for muon events in SNO. In the last year of running, wire chambers were installed above the detector to act as an external calibration system for muon events. While not explicitly used in this thesis, the results from that analysis are given in [50]. Two MC datasets were generated, one for the upward-going muons and one for the direct cosmic ray muons. Here we provide the setup parameters and resulting event distributions for the MC datasets.

The upward-going muon set was combined from two separate simulations. In the first, simulated muons were produced in the rock surrounding SNO. The second set had the muons originate in the water in SNO. The rock simulation began with the flux of neutrinos given in

$\hat{q}$
$\hat{t}$
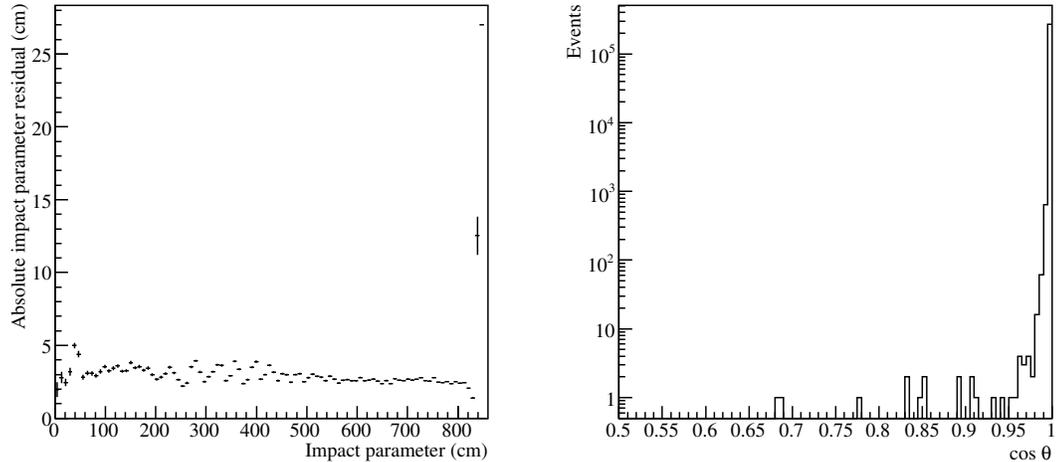
Figure 4.1: FTI performance on MC muons. The impact parameter residuals (absolute difference between the known and calculated values) are shown on the left in a profile histogram and cosine of the angle between the fitted and true directions are on the right. The points in the profile histogram show the mean residual in each impact parameter bin while the size of vertical error bars show RMS deviation within each bin. The fitter consistently performs well up to 840 cm. This is the main reason for the impact parameter cut. The mean of the overall impact parameter residual distribution is 0.1 cm with an RMS deviation of 4.8 cm. A $\cos\theta$ value of 1 occurs for fits where the angle between the fitted and true directions is zero.

equation 1.12 with primary energy 600 MeV to 8.9 TeV and the Sun at "solar maximum". The flux parametrization was taken from the Bartol group [42]. The D_2O and salt phases occurred through solar maximum, while the solar cycle transitioned to a minimum in the NCD phase. The effect of the solar cycle on this analysis is detailed in the systematic errors discussion (section 6.1.2). These neutrinos were run at $500\times$ normal statistics and propagated through SNO-specific rock to the detector using `nuance`. Neutrinos interacting with the rock produced muons which were then tracked to the detector and passed to SNOMAN. This was the main source of atmospheric neutrino-induced muons (upward-going muons). Figure 4.1 shows the performance of the muon track fitter with the data

		Events	Events $\cos\theta > 0.4$	Expected events	Expected rate (yr^{-1})
D_2O	rock	75200	56967	113.9±6.1	123.4±6.7
	water	3984	2924	14.6±0.9	15.8±1.0
Salt	rock	109838	82865	165.7±8.8	121.2±6.5
	water	6381	4678	23.4±0.9	17.1±0.8
NCD	rock	85550	64257	128.5±6.9	119.6±6.5
	water	4912	3673	18.4±0.9	17.1±0.8
Totals		285865	215364	464.5±24.6	138.0±7.3

Table 4.1: A summary of upward-going MC muon events passing the analysis cuts. The final rate in the last row is not a sum of the column, but an expected rate calculated from the total number of expected events considering the entire livetime. The errors include systematic effects discussed in the signal extraction chapter.

cleaning cuts (described below) applied. The reconstructed impact parameter residuals are shown on the left and the cosine of the angle between the fitted and true muon directions are on the right. Perfect alignment of the fitted and true directions give a $\cos\theta$ value of 1. Both distributions show the fitter reconstructing events very well.

The second subcomponent of the upward-going muon MC was produced from neutrinos interacting with any of the water in the detector to produce muons. This includes the H_2O inside and outside the PSUP as well as the D_2O in the acrylic vessel. All three types of neutrinos were simulated (ν_e, ν_μ, ν_τ) and electron and muon events that passed the analysis cuts were included in the angular PDFs. The taus were generated an artificially high rate compared to the electrons and muons for a separate SNO analysis and therefore tau events in the detector were removed. The electrons and muons were generated at 200× normal statistics. Combining the muons from these two MC simulations provided the simulated atmospheric neutrino-induced muon signal. The total event numbers and rates for each simulated phase are given in table 4.1. The errors are all statistical.

The direct cosmic ray muons are produced in the atmosphere. The muons have a limited

	Events	Rate (day^{-1})
D$_2$O	20056	59.5±0.4
Salt	30312	60.7±0.3
NCD	23744	60.5±0.4
Totals	74112	60.3±0.6

Table 4.2: A summary of direct cosmic ray MC muon events passing the analysis cuts. The final rate in the last row is not a sum of the column, but an expected rate calculated from the total number of expected events considering the entire livetime.

range in rock, so their detected numbers are reduced significantly as a function of zenith angle and almost none are seen with a cosine of the incoming zenith angle less than 0.4. These were a background to the neutrino-induced muon dataset and are the reason for the angular cut. The expected number of events from this source that pass the angular cut was non-zero and accounted for in the analysis as a background. The natural rate of direct cosmic ray muons was high enough that running the simulations at some increased rate would not have reduced the uncertainties significantly. The simulation for direct cosmic ray muons were run at its natural rate. Table 4.2 shows a summary of the MC events from the direct cosmic ray muon simulation. Figure 4.2 shows the combined MC zenith angle histogram.

Comment on Stopping Muons

These MC simulations included muons that stopped in the detector as well as through-going muons. All of the components of the muon fitter and the muon event cuts that were developed were done so with through-going muons as the targeted events. While flagging muons that stop in the centre of the detector is a relatively easy task, the difficulty increases quickly as the tracks stop closer and closer to the PSUP since the events look more and more like through-going muons. Characterizing these events with the current muon fitter
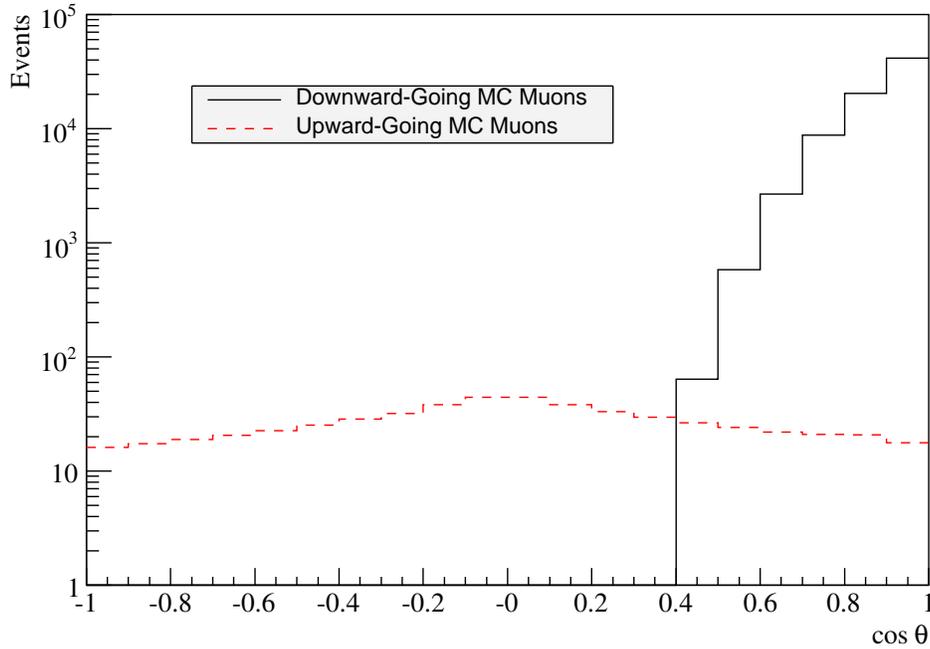
Figure 4.2: The combined zenith-angle histogram of all muon MC sets. The histogram is scaled to the livetime. The distribution of upward-going muons shown in dashed red is roughly symmetric about $\cos\theta = 0$ and the large number of downward-going muons shown in solid black are seen with $\cos\theta > 0.4$.

would be a very difficult task.

Labeling events as signal or background was generally a well-defined operation, but in the case of stopping muons, those that pass the cuts were labeled signal events. This was possible considering the final flux results are quoted in reference to the fully simulated MC flux. Those stopping muons that pass the cuts in the MC were expected to pass the cuts in the data. Essentially, the effect of having stopping muons in our through-going sample cancels itself out in the final flux result. Stopping muons will be discussed further in the description of muon cuts below and in the final results chapter.

Cut Name	Cut Description	Pass Value
PMTHIT	# of Calibrated PMTs	> 500
NNECK	# of Firing Neck PMTs	< 4
BURST	Pass/Fail	Pass
RAWTRMS	Raw tube time RMS	< 38 ns
RAWQRMS	Raw tube charge RMS	> 4.5 GSU
FIB	Impact Parameter	< 830 cm
FNP	Fit number of p.e.	> 2000 p.e.
CITR	Cone In-Time Ratio	> 0.85
RCHT	Ratio of In-Cone Hits	> 0.7
LDISC	Linear discriminant	> 0.6
DEDX	dE/dx	> 200 MeV/m

Table 4.3: A through-going muon analysis cut summary showing the cut name, description and pass value. The cuts are divided into two groups. The top group has cuts which do not depend on the fitted muon track, while the bottom group of cuts are dependent on the fitted track.

4.2   Event Cuts

Muons traveling through SNO are quite distinguishable in comparison to electron events and many backgrounds, since a large amount of energy is deposited in the detector in a very structured manner. The time and charge distributions of the hit PMTs, as shown in the event reconstruction are distinct. This made it possible to select muons from the dataset and remove backgrounds with a small number of cuts. The charge and timing information of muons was the basis for many of the cuts.

The cuts were split into two branches. Low-level cuts were placed into the first branch and high-level cuts into the second. The division was made to set the basis for the bifurcated analysis described in the following chapter. A bifurcated analysis uses two cuts or cut branches to estimate backgrounds that are not modeled. Table 4.3 summarizes the cuts.
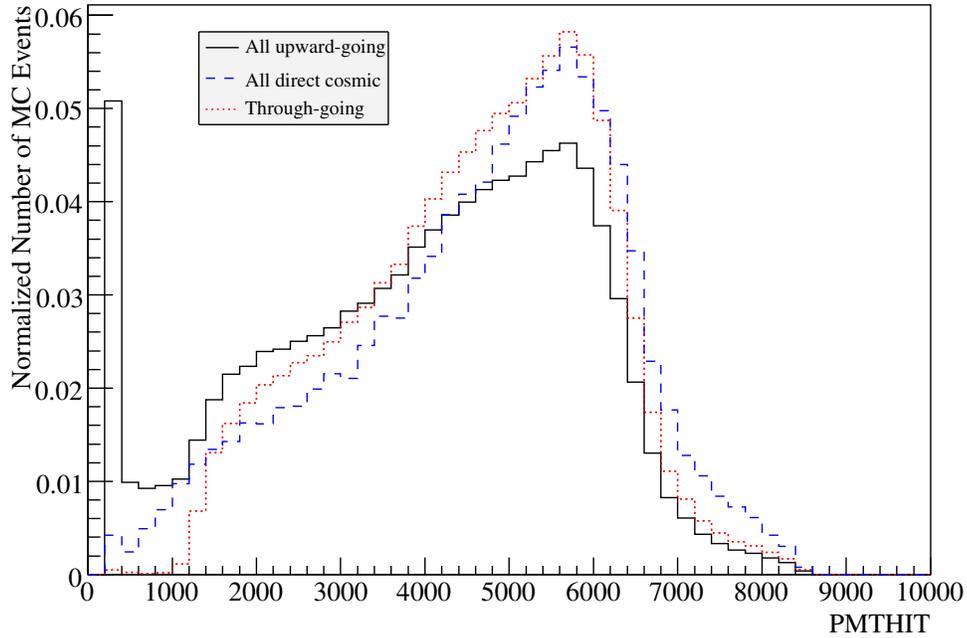
Figure 4.3: The distribution of the number of hit PMTs for the Monte Carlo data set. The solid line is the upward-going set, and the dashed line shows the direct cosmic ray muons and the dotted line shows only through-going muons. The upward-going muons contain relatively more stopping muons than the direct cosmic ray muons. No other cuts are applied.

4.2.1   Low-level Cuts

These cuts include the parameters directly recorded from the event and independent of any fits to the events. The first cut, PMTHIT required that a minimum of 500 calibrated tubes fired in the event. Figure 4.3 shows the distribution of the number of calibrated tubes that fired in a set of simulated muons. All events with fewer than 1000 hit tubes either stopped in the detector or had a track with an impact parameter greater than 830 cm. A cut at 500 was a conservative value giving zero sacrifice.

The neck of the detector contains piping and allows access to D_2O from the deck. Four

72

PMTs were installed in the neck region to detect light being produced in that region from discharges or other instrumental sources. If all the neck tubes fired the event was discarded as a background. Since a number of cuts were used it was not necessary, nor possible for each individual cut to be perfectly efficient at removing backgrounds. While events which fire 1-3 neck tubes could still be a neck event, the higher-level cuts would be able to separate them out. The cut was placed at four to ensure that none of the actual signal events was being removed.

Many of the instrumental backgrounds occurred in bursts. If four events with greater than 250 PMTs hit occur within 2 s, the event was flagged as a burst event. These values were obtained by studying bursts in the D_2O dataset. Burst events were removed along with the time window they occupy. This simply reduces our livetime by roughly 0.025%.

The raw timing and charge information from the PMTs is also used to remove background events. Backgrounds like burst events can have flat PMT hit time distributions. The root mean squared (RMS) of the PMT hit times is a measure of the flatness and based on MC events, the RMS of the raw time had to be less than 38 ns. Point-source background events can sometimes have a narrow charge distribution. The RMS charge distribution is larger for a muon event and therefore events with a RMS charge of less than 4.5 GSU were removed. While not all background sources were point-source, it is shown in the next section that this cut was effective at removing those that were.

4.2.2 High-level Cuts

These values are only available after a fit to the muon event has been made. The first of these cuts define our fiducial area. The ability of the muon event fitter to accurately determine a muon track is consistently good over most impact parameters, but due to the geometry of the detector and PMT housings, events with an impact parameter greater than

roughly 840 cm are extremely difficult to reconstruct. This puts an upper bound in impact parameter on the muon events that we can confidently reconstruct and restricts the area with which we include muon events in our analysis. Based on simulated events, events which reconstructed with an impact parameter greater than 830 cm were removed. This gives a fiducial area of 216.4 m^2, which is roughly 95% of the total detector area.

The FTI routine fits for the number of photoelectrons the detector would observe in an average event at a given impact parameter. This value is used as a cut instead of using a total charge sum because it is less sensitive to how many PMTs are online at any given time. For an event, if a number of PMTs are offline near the exit point of the muon, where the detected charge is highest, a simple sum would be significantly different than if that same event exited when all the PMTs were online. Simulated through-going muons never resulted in a fit number of photoelectrons less than 3000. A conservative cut of 2000 photoelectrons was used.

The next two high-level cuts rely on the pattern and timing of PMTs hit by the Cherenkov cone. For a through-going muon, nearly all the PMTs that are hit are inside the cone. Taking the ratio of tubes within the cone that should have been hit and were hit to the number of tubes that should have been hit provides a strong test of the fitted track. The value of the Ratio of Cone Hit Tubes (RCHT) was set at 0.7 based again on MC muon events. Stopping muon events have PMT hit patterns that are not filled in which gives a small RCHT value. The Cone In Time Ratio (CITR) is the number of PMTs within the Cherenkov cone that fire within a 5 ns window around the fit event time. Background events where light was produced at a point instead of a track have a low CITR value since the timing distribution for a point source is relatively flat. Events with CITR greater than 0.85 passed the cut.

A linear discriminant was also used as a higher-level cut. This method cuts a line with
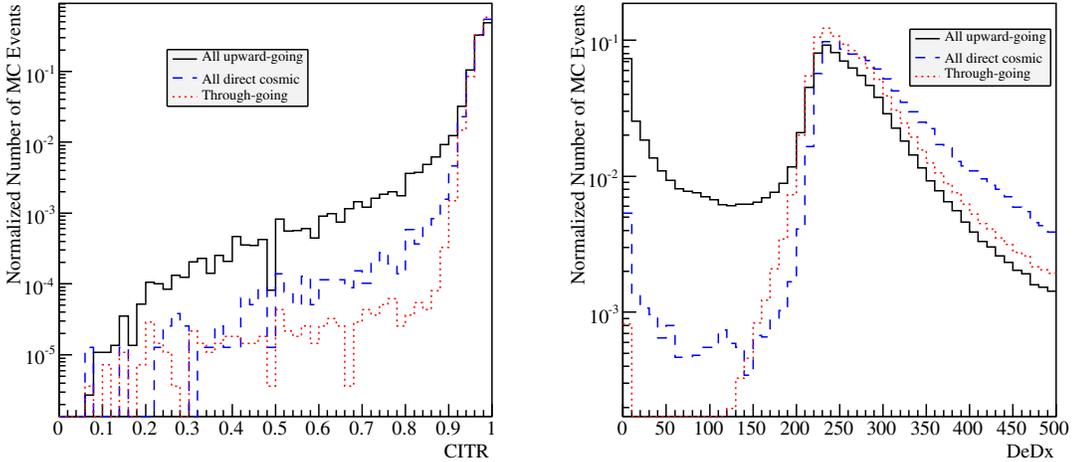
Figure 4.4: The distributions of CITR and DEDX for MC muons. The solid line is the upward-going set, and the dashed line shows the direct cosmic ray muons and the dotted line shows only through-going muons. Stopping muons that do not fill in the Cherenkov ring pattern have low values of CITR. Since the calculation of dE/dx assumes the muons travel all the way through the detector, events which stop deposit a smaller amount of energy than expected. No other cuts are applied.

some slope through a plane made by two values. Here a linear discriminant was formed with a combination of CITR and CRMS. CRMS is the in-cone time RMS, or the RMS of the PMT time residuals within the Cherenkov cone. For a muon traveling through the detector, each PMT should be hit by direct light at a time established by the PMT position in relation to the muon track. Background events have a high CRMS value. While the effectiveness of CRMS as a cut alone is low, CRMS in linear combination with CITR provides a strong means to remove background events.

Muons that stop in the detector are a significant background in the through-going analysis. Muons produce light and they lose energy as they propagate through the detector. For a given impact parameter, which translates into track length, a through-going muon

	Direct cosmic ray μ			Upward-going μ		
	D_2O	Salt	NCD	D_2O	Salt	NCD
Fib	94.9±0.9	95.4±0.7	95.3±0.9	94.6±0.4	94.7±0.3	94.8±0.4
PMT Hit	98.3±0.9	98.5±0.7	98.5±0.9	90.7±0.3	91.5±0.3	92.1±0.4
Nneck	98.9±0.9	99.0±0.7	99.0±0.9	96.9±0.4	97.0±0.3	96.6±0.4
FnP	98.0±0.9	98.1±0.7	98.2±0.9	87.9±0.3	88.5±0.3	89.6±0.4
RawT	98.8±0.9	98.9±0.7	98.9±0.9	96.3±0.4	96.5±0.3	96.2±0.4
RawQ	97.5±0.9	97.5±0.7	97.4±0.9	76.1±0.3	76.3±0.2	75.6±0.3
CITR	99.2±0.9	99.3±0.7	99.4±0.9	96.2±0.4	96.4±0.3	96.8±0.4
RCHT	97.7±0.9	97.7±0.7	97.8±0.9	81.1±0.3	81.5±0.3	82.6±0.4
LinD	97.5±0.9	97.6±0.7	97.8±0.9	87.3±0.3	87.5±0.3	88.6±0.4
DEDX	97.2±0.9	97.0±0.7	97.0±0.9	72.8±0.3	73.0±0.2	73.1±0.3

Table 4.4: The MC cut efficiencies for the three phases of SNO. The values are the number of events that passed the cut divided by the total number of simulated events. Stopping muons were not removed from the sample for the calculation of these numbers. Since the each of the cut efficiencies remain consistent across phases, we can conclude that the properties of the simulated muon events were consistent.

deposits a defined amount of energy. A description of muon propagation through matter and energy loss was given in section 1.2.2. If a muon stops in the detector the amount of energy not deposited is related to how much track length the muon had left to travel. An estimate of the energy loss per unit distance (dE/dx) of the muon was developed to track this. FTI fits for the number of detected photoelectrons and this provides an estimate of the charge deposited in the detector. Simulated through-going events were run to calibrate and convert this number into an energy per unit length. Events with DEDX greater than 200 MeV/m were included as through-going muons. The cut label "DEDX" is differentiated from the actual physical quantity "dE/dx" as the former was only an estimated value proportional to the physical quantity. Figure 4.4 shows the CITR and DEDX distributions for MC muons. CITR provides information about the PMT timing and hit pattern while DEDX reveals the charge per unit length that an event deposits into the detector.
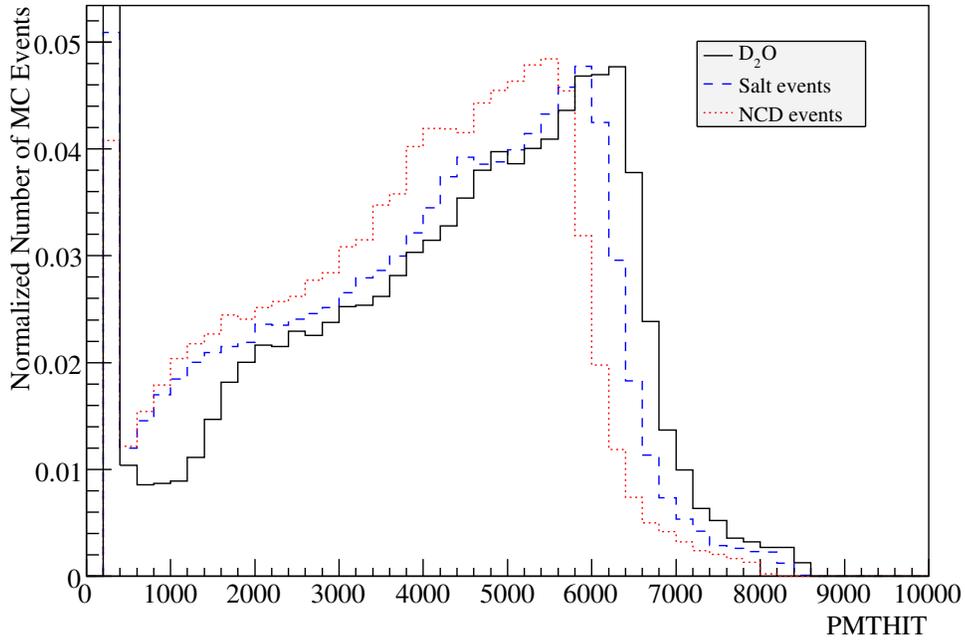
Figure 4.5: The PMTHIT cut distribution for each of the three phases. While the shapes are mutually consistent, the peak of the distribution shifts to smaller values in each successive phase. This is attributed to PMT failures through the lifetime of the detector.

4.2.3   Monte Carlo Three-Phase Consistency

The cuts and analysis were originally developed for the D_2O phase of SNO. The addition of salt in the second phase of the detector and NCDs in the third phase should not have a significant effect on the analysis. Consistency checks were performed to ensure the data and MC in all phases were understood. Table 4.4 compares cut efficiencies across the three phases. The efficiency is the number of simulated muon events that passed the cut divided by the total number of simulated muon events. For each of the MC datasets (direct cosmic ray and upward-going) most the cut efficiencies were consistent over the three phases of the experiment. The PMTHIT and FnP distributions show a shift to fewer hit PMTs and
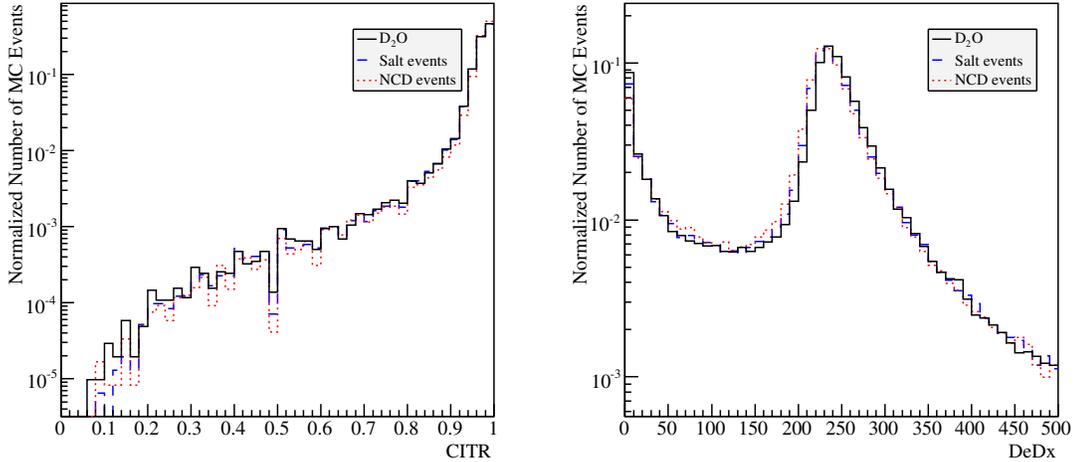
Figure 4.6: The CITR and DEDX distributions for each of the three phases. The agreement between the distributions suggest the timing, charge and hit pattern of events remain consistent across phases.

lower charge which is expected as PMTs fail over the lifetime of the detector. Figures 4.5 and 4.6 show the cut distributions for PMTHIT, CITR and DEDX for each phase of SNO. The CITR and DEDX distributions are similar.

It is notable that even through the NCD phase, with 40 proportional counters installed in the D_2O volume of the detector, that more significant changes to the cut efficiencies and fitter performance were not seen. The drop in PMTHIT was probably also due to the NCDs occulting and scattering light. The main reason for the stability of muon events across phases is the large amount of light that is produced by an average muon. Many of the smaller scale effects were most likely undetectable given the thousands of PMTs that were hit. The NCDs have a diameter of roughly 5 cm and were deployed in the central region of SNO. While the amount of light that is blocked or scattered depends on the geometry of the muon, it would appear on average that the impact on ability of the fitter to reconstruct

the muon tracks was low.

The NCDs would not have an effect on the neck cut as the cut was targeted specifically for neck events. There would also not be a significant reason for the RAWTRMS and RAWQRMS distributions to change since it would require a large systematic spreading of the timing and charge distributions. RCHT and CITR are the ratios of tubes that should have been hit. For an average muon that travels roughly through the centre of the detector, more than 4000 PMTs should be hit by direct light. The shadowing effect of the NCDs was not significant enough to change the ratio. The consistency through the phases of the DEDX distribution was expected as it was tuned to simulated muons in each phase.

Consistency in the cut efficiencies and sample cut distributions suggest the MC datasets are understood across phases and are consistent with each other. Data comparisons are presented in the next chapter.

Chapter 5

Data Selection

Here we will discuss the steps taken in using the cuts developed from the MC to select the data events for this analysis. The cuts served to separate signal from background events, provided measures to check the consistency of the data through the D_2O and salt phases and were applied in a bifurcated analysis to estimate the contamination of our sample by non-physical backgrounds. Descriptions of the data and livetime are presented first, followed by a listing of both physical and non-physical backgrounds the cuts were intended to remove. We then detail the use of these cuts to perform consistency checks.

5.1   The Data

In an experiment, it is important to control and understand the state and environment of the detector in order to properly characterize the signal and backgrounds. The condition of the detector was flagged by a detector operator. This allowed for the division of the detector livetime into "runs". There were generally three types of runs the detector could have been in. The first was the optimal data taking mode ("neutrino run") where all or most systems were functioning properly. The second was a "calibration run", in which any of the numerous calibrations to the electronics or water systems were being performed. The last mode was a "maintenance run" which occurred when the detector was in a transitional state either because of a malfunction or in preparation of calibrations or electronics repairs.

All SNO analyses use data from neutrino runs. A series of general checks was performed before the detector could be placed in neutrino mode. These checks included proper voltage thresholds on the PMTs, no calibration or maintenance activity around the detector and anything else that would compromise the quality of the data. Flagging neutrino runs was the first step in selecting data for any of the analyses at SNO. A run selection committee applied further restrictions to the neutrino runs in order to select runs acceptable for the neutrino analysis. The committee checked for smaller-scale hardware problems like offline crates or out-of-sync timing systems and any other factors that may affect the data before a run could be switched out of a neutrino run. Run selection was meant to maximize livetime, while removing runs with instrumental effects that would bias or contaminate the data. In addition to the standard neutrino analysis runs, certain runs were included with relaxed constraints. A full account of the neutrino run selection criteria and list of runs is given in [51]. Here we will describe a few of the criteria that were relaxed in the muon run selection that recovered runs removed by the neutrino analysis.

A muon event sees thousands of PMTs fire and therefore our analysis was less sensitive to many of the backgrounds that are problematic for the other neutrino analyses at SNO. Fourteen magnetic coils in the detector cavity compensate for the Earth's magnetic field. PMTs lose gain and efficiency in a magnetic field which results in a change in the energy response of the detector. Runs where one or two of the coils were offline were rejected from the solar neutrino run list since a small number of affected tubes might have been a significant number of the hit PMTs in an event.

The solar neutrino analysis was also very sensitive to many low energy radioactive backgrounds. Runs were removed from the neutrino list if radioactive backgrounds were present either accidentally or as result of calibrations. At times, certain crates were taken offline either for maintenance or other hardware problems. With one crate offline, up to 512

	D$_2$OPhase	Salt Phase	NCD Phase	Totals
Raw livetime ($\times$ 10^7 s)	2.9146883	4.3153939	3.3945181	10.6246003
Retrigger cut (sec)	25	62	56	143
Muon Nhit burst cut (sec)	7405	3552	4924	15881
Subrun Boundary (sec)	0	0	23032	23032
Corrected livetime (days)	337.25$\pm$0.02	499.45$\pm$0.02	392.56$\pm$0.05	1229.26$\pm$0.03

Table 5.1: The livetime summary for the three phases. Shown are the raw livetimes as well as the corrections made to account for retrigger dead time, the burst event cut and subrun boundary effects.

channels or PMTs could be offline. This could a serious problem for the neutrino analysis as solar neutrino events only hit up to 150 PMTs. These runs were also removed from the solar neutrino run list.

With the runs selected, a livetime was calculated using the 10 MHz clock. The raw livetime recorded was corrected for retrigger dead-time and the muon NHIT burst cut. The retrigger dead-time accounts for the 5 μs after an event was recorded when the detector could not record another event. The muon NHIT burst cut, removes a set of events if 4 events occur within 2 s and each hit over 250 PMTs. Larger runs were divided into smaller subruns. The subrun boundary cut was a 60 s cut at the beginning of each subrun. With the installation of the NCD and supporting data acquisition hardware and software, the processor which deals with the temporal correlations of events could possibly miss events in that window. Table 5.1 summarizes the livetime values.

5.2 Blindness

Over the development of the analysis, a blindness scheme for data was in place to ensure an unbiased approach in creating the fitter, defining the cut parameters and any other parts involved with developing the analysis. It is meant to prevent any tuning of the analysis to

	Downward-going μ		Upward-going μ	
	Events	Rate (μ/day)	Events	Rate (μ/yr)
D$_2$O	20985	62.2 ± 0.4	149	161 ± 15
Salt	31287	62.6 ± 0.4	192	140 ± 10
NCD	24477	62.4 ± 0.4	173	161 ± 15
Totals	76749	62.4 ± 0.2	514	153 ± 7

Table 5.2: Event numbers, rates and totals for the full open dataset.

produce "expected" results. A blindness scheme can either limit the amount of available data, making conclusions statistically insignificant or mask an observable quantity of the data so a key result cannot be obtained or in many cases, both.

The full D$_2$O dataset was previously open for the analysis in 2001 [33]. The current analysis on the D$_2$O data cannot strictly be called blind but the scheme used here covered the data over all three phases and was implemented at the beginning of this analysis. Since the analysis was already statistics limited, reducing the amount of available data was a significant part of the blindness scheme.

For the D$_2$O and salt phase data, runs 10000 to 12168 and 20684 to 22450 were designated open data. This resulted in roughly a 65% data reduction. All other runs in those phases were not looked at during the development of the analysis. The data in the NCD phase were blinded by a standard SNO data division routine in SNOMAN. The routine tags events as blind based on a user defined blindness fraction. The analysis here used a 20% open fraction for the NCD phase. Since the final atmospheric neutrino oscillation analysis was dependent on the muon zenith angle distribution, muon events were removed from the open dataset as some function of zenith angle. Six functions were written into SNOMAN and one of those functions was chosen by two people at SNO, unrelated to the muon analysis.
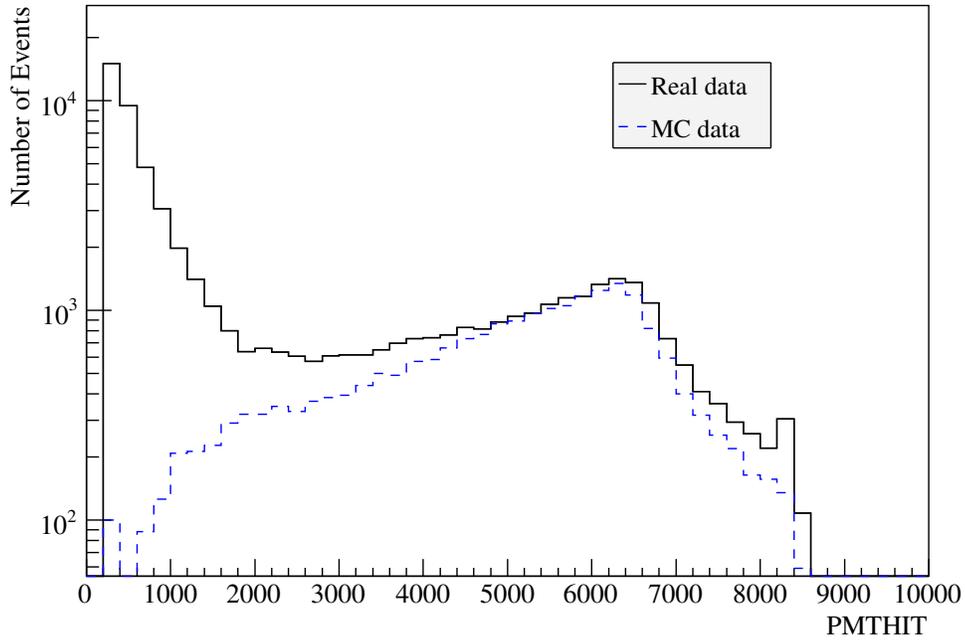
Figure 5.1: The PMTHIT distributions for the entire dataset and MC set with no cuts applied. The peaks of the distributions around 6000 hit PMTs agree. Instrumental backgrounds are not simulated and they are seen as excesses in PMTHIT bins below roughly 3000 and above 7000.

The blindness scheme was lifted after the fitter, cuts, background analysis and error analyses were finalized. Once all of the data were open, the overall event numbers and rates were checked before the analysis and signal extraction code were run on the complete dataset. The checks were completed independently by three members of the muon group. Table 5.2 summarizes the event numbers and rates across phases and the totals.
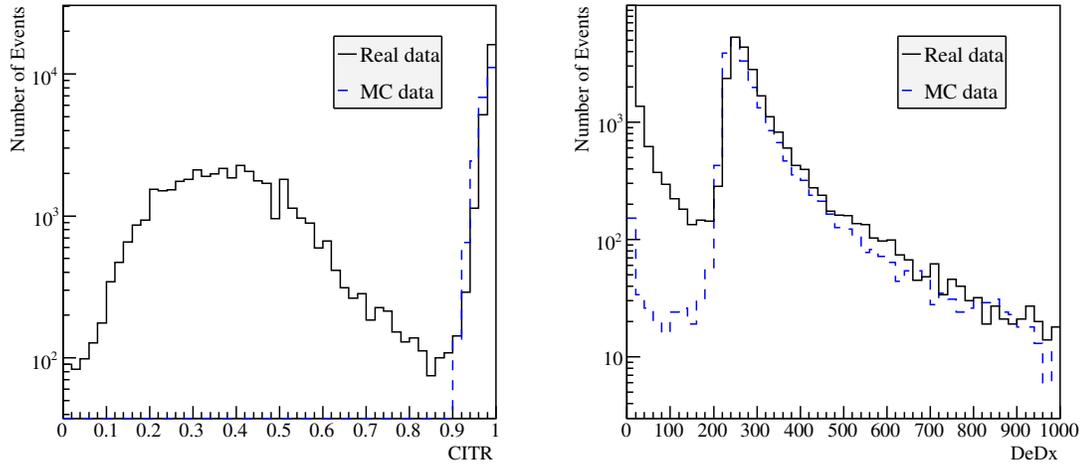
Figure 5.2: Two high-level cut distributions for the data and MC. CITR is shown on the left and DeDx on the right. The general agreement of the signal in both cases also shows the distribution of backgrounds in each of the cut parameters.

5.3   Dataset Consistency

Before analyzing the data, comparisons between the MC events and real events were made to ensure that the MC data were a valid simulation of the real data. In addition to the comparisons between data and MC, the consistency of the data between each of the three phases was checked. Here the cut distributions and efficiencies are used to make these comparisons.

5.3.1   Data and MC Comparisons

In comparing the MC datasets, cut efficiencies and the cut distributions were compared. Comparing the data to MC with cut efficiencies does not provide much information since in the data the signal and background fractions are unknown. Any quantitative check would be difficult because of this. However, inspecting cut distributions could still provide
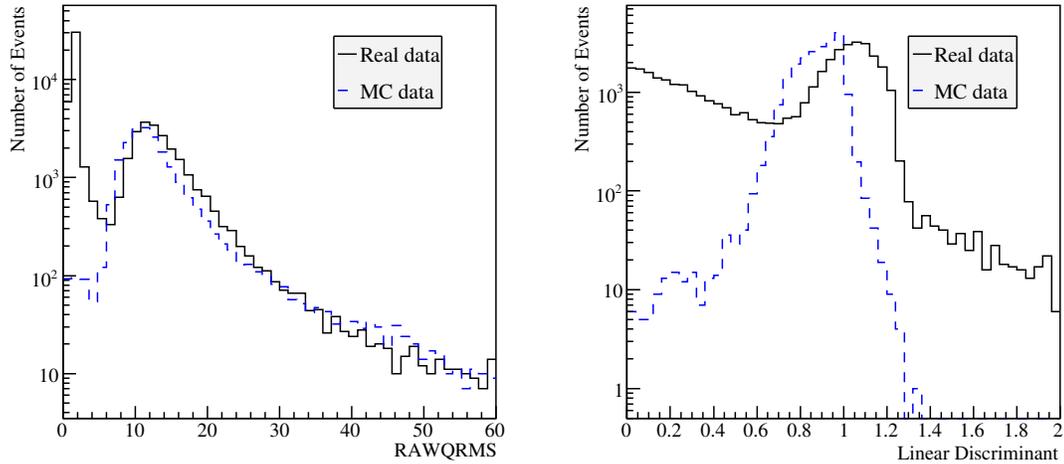
Figure 5.3: RAWQRMS and the linear discriminant cut distributions for the data and MC. The disagreement of the two distributions are both discussed in the systematics section below.

a method to compare the data to the MC. Figures 5.1, 5.2 and 5.3 compare a number of cut distributions for the data and MC with no cuts applied. The number of hit PMTs is a very simple low-level cut while CITR and DEDX give time and charge information of the event based on the fitted muon. The RAWQRMS and linear discriminant are shown as well since their differences were large enough to be treated as systematic errors described in the next chapter. In the signal region of each of the plots, the data and MC generally agree for all except the linear discriminant. Since the data distribution of real events was shifted away from the linear discriminant cut value there was no concern that real events were being removed by this cut. As with figure 5.1, the MC set only contains muons and does not include any of the instrumental backgrounds. While each cut may have instrumental backgrounds which leak into the signal region, we will show in the following section that with the full set of cuts used the fraction instrumental backgrounds that leak into our
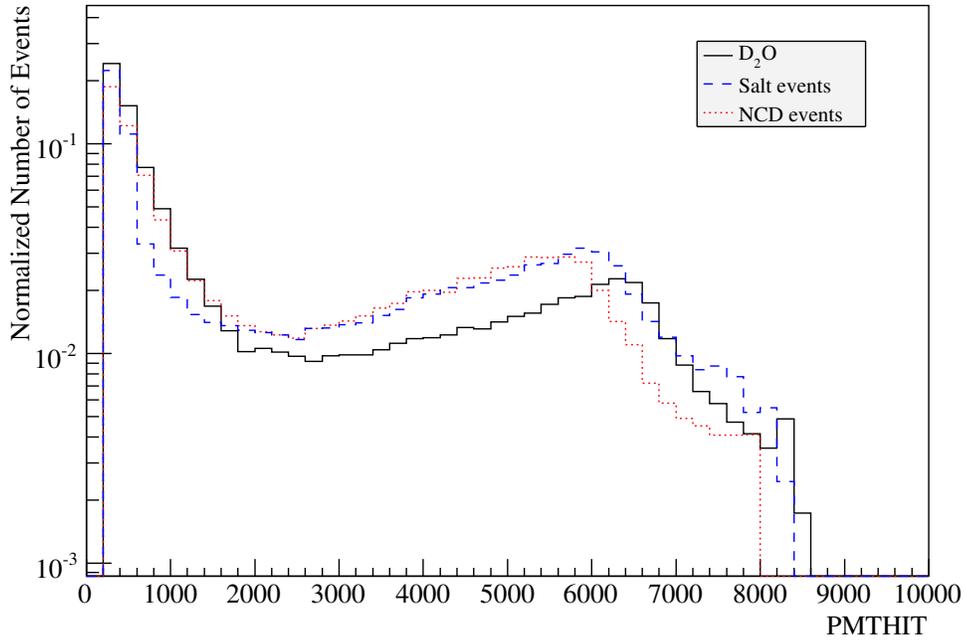
Figure 5.4: A comparison of the PMTHIT data distributions for each phase. The shift of the peak to smaller PMTHIT values through each successive phase is seen here as in the MC (figure 4.5). Since each phase had a different livetime, area under each of these curves was normalized to one.

sample was extremely low and that fraction was insensitive to variations in individual cut placement.

5.3.2 Three-Phase Comparisons

Each phase introduced a new aspect to the detector which provided different measurements for solar neutrino analysis. These changes were expected to have a negligible impact on the muon analysis. Checks were necessary to ensure these changes had little or no effect on our understanding of muon events or the ability of the cuts to distinguish background from signal. Even though a number of cuts were applied to the data to accomplish this, looking
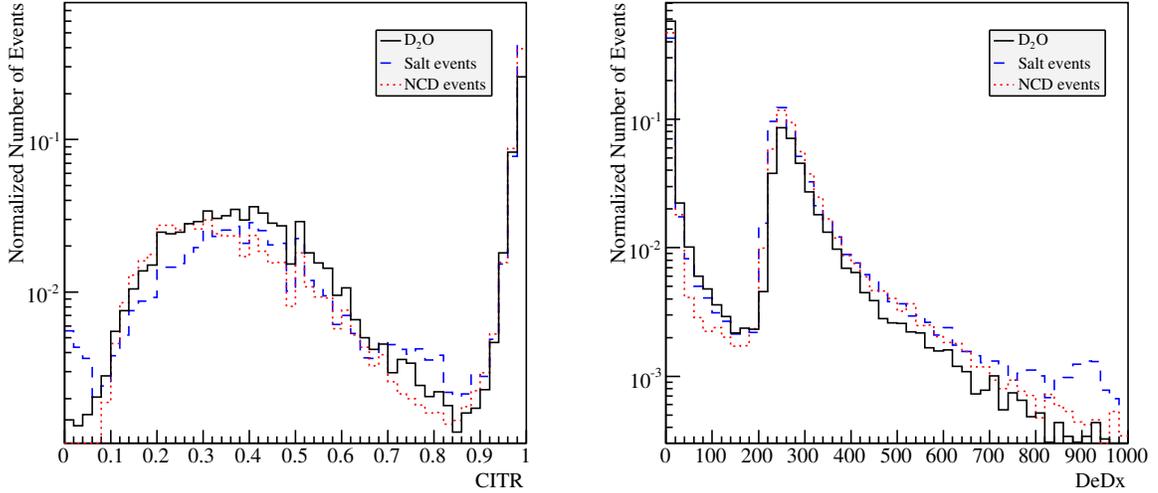
Figure 5.5: Comparisons of CITR and DeDx through each phase. As with the MC comparisons, there is a general agreement for each cut. The small variations are attributed to the properties of the backgrounds changing through the phases. Since the area under the plots are normalized to one, the signal regions in the plots show small variations as well.

at individual cut distributions over each phase helped to check that the characteristics of a muon event remains relatively consistent. It was more important to the final analysis that the cut distributions of the MC and real data agreed, since the signal events were simulated with a full MC of the expected muon flux with the appropriate run and phase conditions, The checks between phases were intended to search for more drastic differences that might have been overlooked.

Figures 5.4 and 5.5 show the PMTHIT, CITR and DEDX normalized cut distributions for the data over each of the phases. The shift of the peak to smaller PMTHIT values through each successive phase was also evident in the MC and caused by PMT failures as the detector aged. The general agreement of the higher-level cuts shown in figure 5.2 suggest no significant changes to muon events over the phases of SNO.

5.4 Background Estimates

A background was defined to be any event that could be mistaken for a through-going muon event. The fraction of background events that made it into the signal regions of cut parameter space was called the contamination. While muon events were fully simulated, there existed no such MC of the instrumental backgrounds. However, it was still possible to understand the different sources of contamination and estimate their impact on the analysis. We will first describe some of the background sources then present the bifurcated analysis used to calculate estimates of contamination.

5.4.1 Background Sources

Backgrounds can be classified as physical or instrumental. Physical backgrounds for the through-going analysis included any nuclear or particle event that was produced inside the detector, stopped inside the detector or both. These events were real physics events that were excluded from our analysis because the current muon fitter could not fit them well. Since all of the significant physics interactions and particles were simulated, we defined any event passing all the cuts as a through-going muon. This definition accounts for the loss of physics events because of the cuts (known as "sacrifice") and the contamination of background physics events since the efficiencies were known from the MC simulations.

On the other hand, with no simulations, the problem of estimating instrumental backgrounds was more difficult. Instrumental backgrounds include events which produce a lot of light. Such events are PMT connector breakdowns, neck events and bubblers. The breakdown of the wet-end PMT connector could produce a large amount of light in a hit pattern resembling a muon event. A closer look at the timing structure of the hit PMTs would separate these from real muon events since the light originated from a point rather than a track. These breakdowns occurred approximately 3 times per week. Figures 5.6 and 5.7

89
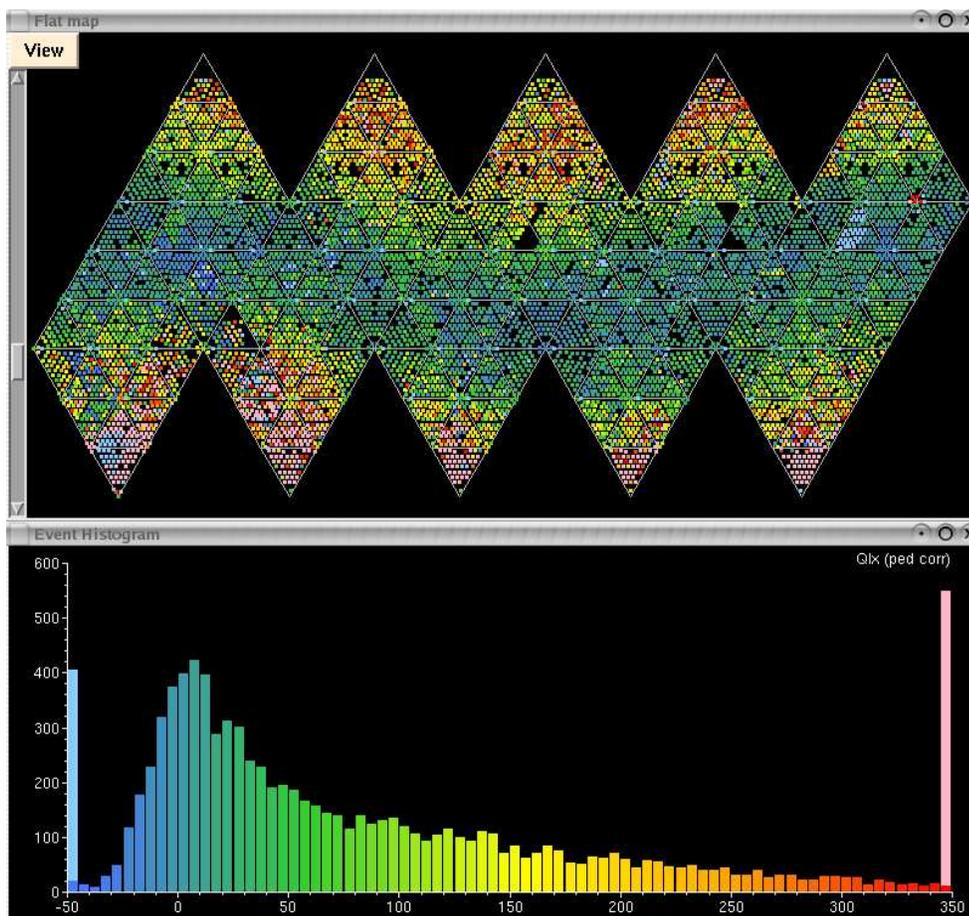
Figure 5.6: PMT connector breakdown XSNOED Flat Charge Map. A single through-going muon would not produce two high charge PMT groups as shown with one at the top and one at the bottom.

show flat maps of the charge and time distribution of the same PMT breakdown event. Thousands of PMTs were hit as in a muon event, but there were two high-charge PMT groups at the top and bottom. A muon event would only have one at the exit point. The PMT time distribution shows a large group of early PMTs on the left hand side. The timing of the other tubes were more suggestive of a point source of light than a muon producing light along a track.
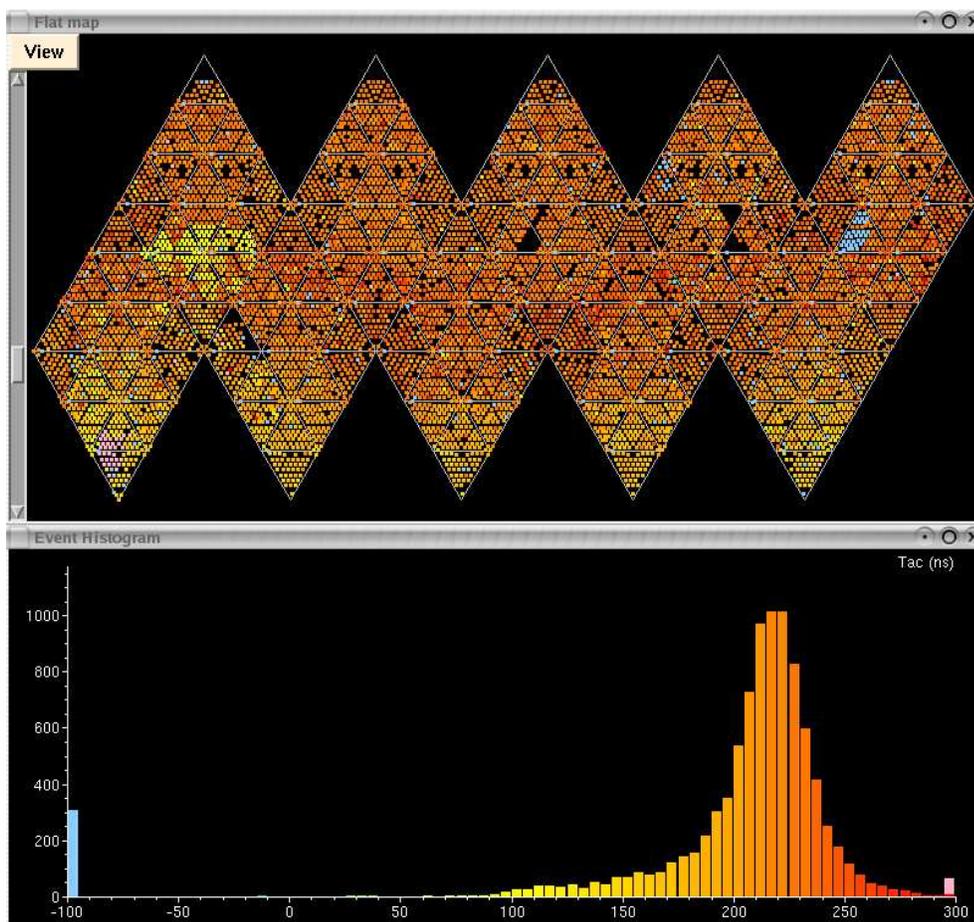
Figure 5.7: PMT connector breakdown XSNOED flat time map. The yellow group of PMTs on the left hand side were probably the cause of the event. A PMT breaking down in that group would trigger PMTs to fire early in that area, then flash light across the detector.

With the many cables and calibration equipment in the neck, electrical discharges would occur. These discharges would produce light down into the detector much like muon. The neck tube cut effectively removed all of these types of events. Over the livetime in this analysis neck events occurred roughly 20 times per day. Figures 5.8 and 5.9 show flat maps of the charge and time distribution of the same neck event. All four of the neck tubes fired. An upward-going muon could produce the charge map show, but timing of PMT hits cannot
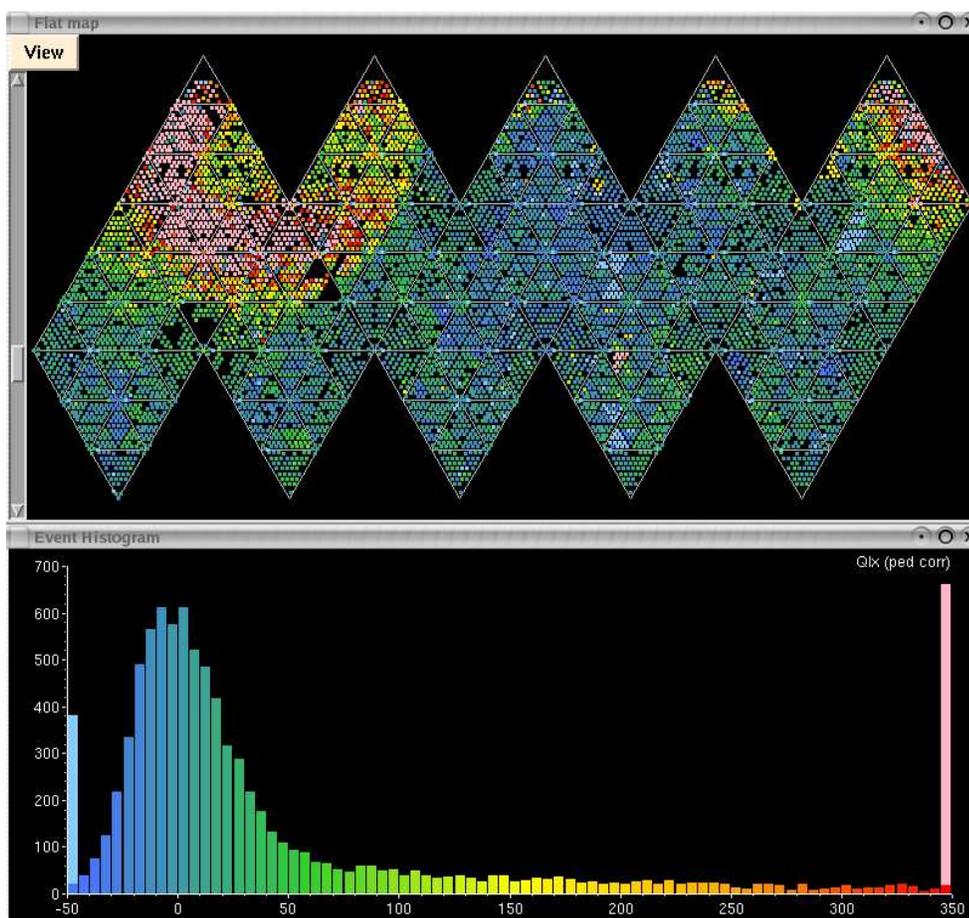
Figure 5.8: Neck event XSNOED flat charge map. All four neck tubes fired. The group of high charge PMTs could be from an upward-going muon, but the time distribution of PMTs in figure 5.9 could not have been from a muon.

be from a muon track. Light from the neck could have flashed down one side of detector (the yellow strip of PMTs). Neck events do not usually have the high charge group of PMTs so this may have been caused by a bubble. The bubbles were produced by the system that monitored the water levels in SNO. Air bubbles float up through the detector to the neck producing light along the way. Runs in which the bubbler system was on were flagged as maintenance runs and not included in the final run list, but even when the system was off,
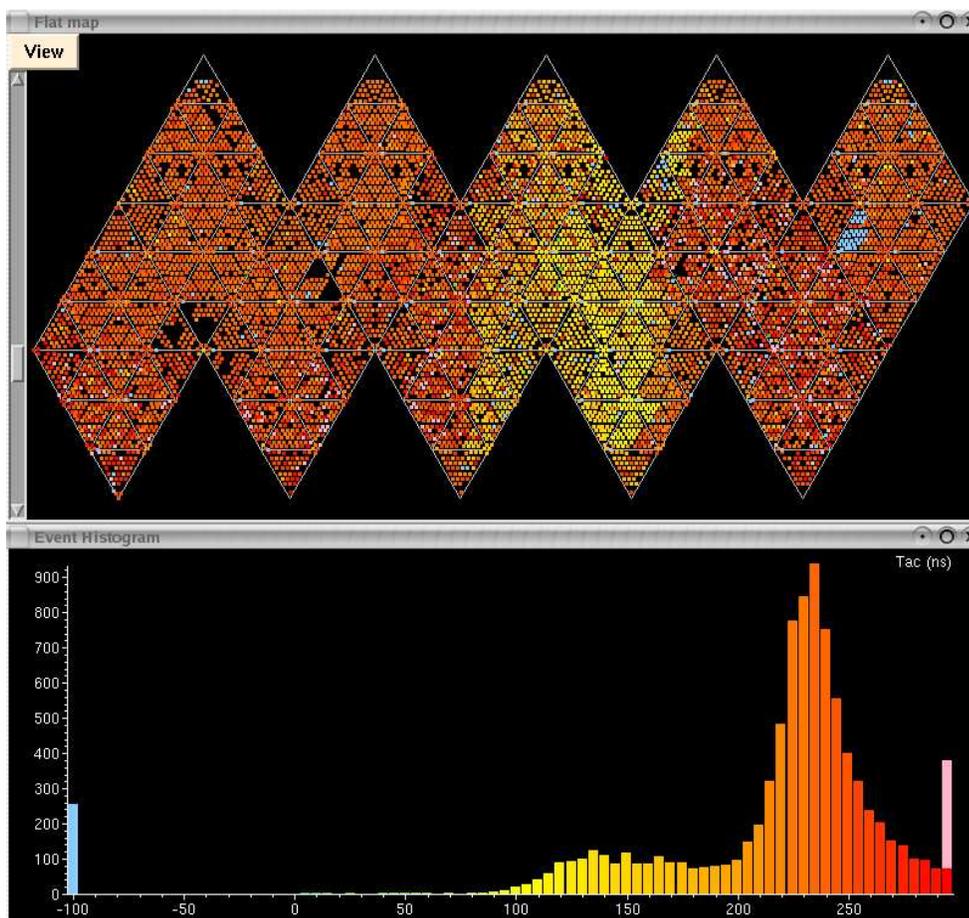
Figure 5.9: Neck event XSNOED flat timing map. All four neck tubes fired. The early tubes were down one side of the detector in a strip. with all other tubes hit at the same time (red PMTs).

random bubbles were still produced.

5.4.2 The Bifurcated Analysis

The analysis here relies on information present when two cuts or cut branches are compared against each other. A cut branch is a number of cuts grouped together. Given a dataset with signal and background distributions in each of the cuts, estimates of the contamination of

93

	Cut 2 Pass	Cut 2 Fail
Cut 1 Pass	a	d
Cut 1 Fail	b	c

Table 5.3: A schematic of the bifurcated analysis. The values a, b, c and d are the number of events in each of the boxes.

the signal region by backgrounds and the loss of signal into the background regions (sacrifice) can be obtained. In table 5.3, a, b, c and d are the observed number of events that fall in each box.

We can write equations for a, b, c and d (equations 5.1) in terms of a number of signal events, ν, and a number of background events β. Since the signal and background events have some distribution in cut space, we also need the cut efficiencies of both branches (ϵ_1, ϵ_2) and the leakage fractions for each cut (λ_1, λ_2). The efficiency is the fraction of signal events that pass the cut and the leakage is the fraction of background events that pass the cut. As an example we can look at the equation for a, the observed number of signal events. The first term, $\epsilon_1 \epsilon_2 \nu$, is the efficiency of both branches multiplied by the true signal and the second term is the leakage fraction of background events into our signal region ($\lambda_1 \lambda_2 \beta$).

$$
\begin{aligned}
a &= \epsilon_1 \epsilon_2 \nu + \lambda_1 \lambda_2 \beta \\
b &= (1 - \epsilon_1) \epsilon_2 \nu + (1 - \lambda_1) \lambda_2 \beta \\
c &= (1 - \epsilon_1)(1 - \epsilon_2) \nu + (1 - \lambda_1)(1 - \lambda_2) \beta \\
d &= \epsilon_1 (1 - \epsilon_2) \nu + \lambda_1 (1 - \lambda_2) \beta
\end{aligned}
\tag{5.1}
$$

The efficiencies of the cut branches were estimated from MC simulations. We now have a system of equations with four observables (a, b, c, d) in terms of four unknowns $(\nu, \beta, \lambda_1, \lambda_2)$. Once the cut branches were applied and the number of events in each region

94

	a	b	c	d
Down-going μ	77376	148	57152	1174
Up-going μ	514	4	38875	910

Table 5.4: The number of events passing and failing the cut branches for both datasets. The a, b, c and d labels were defined in table 5.3.

were tabulated, the signal, background and leakage fractions were calculated by inverting the above system of equations. The inverted system was complex, but only because of the cut efficiencies. In the limit, $\epsilon_1 = \epsilon_2 = 1$, the inverted equations become much simpler and the key contributions are apparent. The equations in 5.2 show the estimates for cut leakage and contamination in this limit.

$$
\begin{aligned}
\lambda_1 &= \frac{d}{c+d} \\
\lambda_2 &= \frac{b}{b+c} \\
\beta &= \frac{(c+d)(b+c)}{c} \\
\nu &= a - \lambda_1\lambda_2\beta = a - \frac{bd}{c}
\end{aligned}
\tag{5.2}
$$

The leakage for the first cut branch, λ_1, is the ratio of events that passed the first branch to the total number of events that failed the second branch. The other leakage fraction is obtained in the same manner, with the branches exchanged. The number of signal events is the number of observed events in the signal less the number of background events that contaminate the signal region ($\lambda_1\lambda_2\beta$). The analysis was completely dependent on the two regions where the events pass one cut but fail the other and the assumption that the two cut branches are independent (orthogonal). Correlations between the cut branches affected the estimates, so the orthogonality of the cut branches has to be rigorously tested.

	Down-going μ	Up-going μ
Signal estimate	77505 ± 280	514 ± 23
Leakage 1 (λ_1)	0.019 ± 0.001	0.023 ± 0.001
Leakage 2 (λ_2) ($\times 10^{-5}$)	177 ± 28	9 ± 5
Background estimate	58344 ± 485	39788 ± 400
Contamination	1.9 ± 0.3	0.09 ± 0.05
Fractional Contamination ($\times 10^{-5}$)	2.5 ± 0.4	17 ± 10

Table 5.5: A summary of the bifurcated analysis results with statistical errors from counting (a, b, c, d) propagated through the analysis. In both cases the contamination was a small fraction of a percent with an estimated 3 events in the downward-going sample and 0.08 events in the upward-going sample.

The two cut branches, summarized in table 4.3, were applied to the MC events to obtain the cut efficiencies then applied to the data. The efficiencies were 0.9994±0.0002 and 0.9989±0.0002 for branch 1 and 2 respectively. The results from the full bifurcated analysis are shown in tables 5.4 and 5.5. The leakage of branch two was smaller than branch one in both samples supporting that the high-level cuts were more effective at removing instrumental backgrounds than the low-level cuts. In both cases the contamination was a very small percentage with an estimated 3 contamination events in the direct cosmic ray muon sample and 0.1 events in the upward-going muon sample. This agreed well with previous studies showing that the cuts used were very efficient at removing instrumental backgrounds.

5.4.3 Orthogonality Check

The main method for testing the robustness of the method was to relax and tighten cuts, allowing backgrounds into or out of the signal region in order to check that the resulting estimates still held. If the cuts and distributions of signal and background were well behaved, the analysis should have identified the new events in the signal region as contamination events while the estimate for the number of signal events remained the same. The

Pair Number	Cut 1	Cut 2	Signal (events)
1	Bursts	CITR	6675 ± 82
2	FNP	CITR	6605 ± 82
3	RAWTRMS	CITR	6677 ± 82
4	RAWQRMS	CITR	6598 ± 82
5	Bursts	RCHT	6517 ± 82
6	FNP	RCHT	7096 ± 85
7	RAWTRMS	RCHT	6737 ± 93
8	RAWQRMS	RCHT	7008 ± 84
9	ALL	ALL	6545 ± 81

Table 5.6: A summary of results from cut pairs for the orthogonality check of the bifurcated analysis. The result from the full bifurcated analysis is included for comparison.

main concern were correlations between the cuts. The method and equations in 5.2 were developed on the assumption that the cuts were independent. When dealing with multiple cuts in a branch, the results become harder to interpret. We first look at how cuts can be correlated and the effects correlations can have on the signal estimates. We then look at the entire analysis and check the reliability of the contamination estimates.

Paired Cut Correlations and Signal Estimates

The bifurcated analysis was tested first by taking the cut branches apart and looking at individual relationships between one cut from the first branch and one from the second. The bifurcated analysis was reapplied to the data with only two cuts per trial. Table 5.6 summarizes the 8 cut pairs that were chosen and the resulting signal estimates from those pairs as well as the results from the full analysis. The data used here were D_2O events with the blindness scheme in place so the absolute signal estimates from this check were some fraction of signal estimates from the final analysis. This test was completed before the blindness scheme was lifted and provided enough statistics for the checks described below.
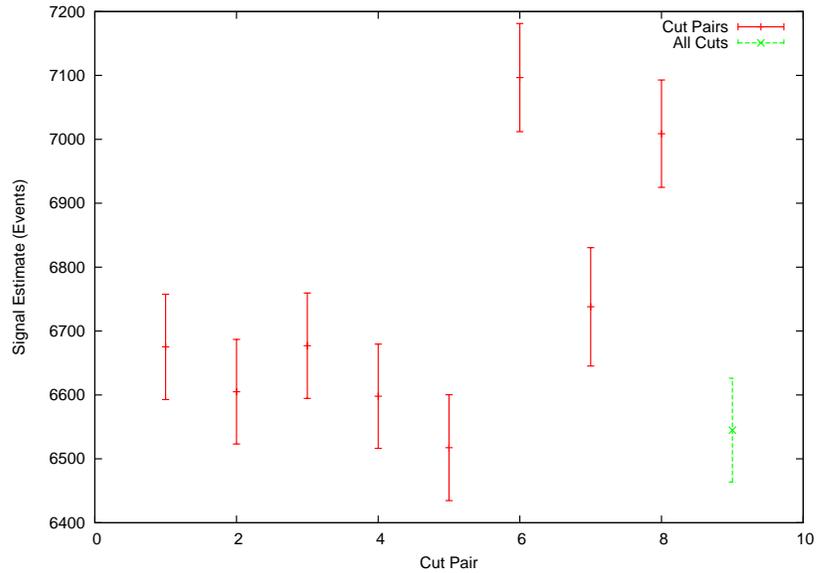
Figure 5.10: Signal estimate results from cut pairs for the orthogonality check of the bifurcated analysis. The result from the full bifurcated analysis is included for comparison

.

The burst, FNP, RAWTRMS and RAWQRMS cuts were paired with either the CITR or RCHT cut. Figure 5.10 shows the data from table 5.6.

All of the cuts were designed to cut as few signal events as possible. In deciding on individual cut values, high signal acceptance was favoured over efficient background removal. A series of cuts were used to ensure that if one cut could not efficiently remove a particular background while saving all of the signal, another cut would. The results in table 5.6 show that certain pairs of cuts (6 - 8) have a higher signal estimate than the others as well as the estimate from the full analysis. Higher signal estimates could result from either background events being counted as signal events or from actual signal events which some cuts were removing. Since all of the cuts were tuned to have nearly zero sacrifice, the higher signal estimates must have been a result of backgrounds leaking into the signal. Cut pairs 1 - 5 remove background events more effectively than cut pairs 6 - 8. The background leakage

Figure 5.11: The raw tube time RMS (RAWTRMS) is plotted against the ratio of in-cone hits (RCHT). The lines show the cut values and the signal region.

was still thought to be extremely low final analysis since all of the cuts were applied.

A more thorough investigation of 2 pairs (7 and 8) was performed. Both cut pairs relied on the RCHT cut from the higher-level branch. The results from cut pair 7 were closer to those from the final analysis than cut pair 8. For each pair, one of the cuts was held at the normal value while the other cut was both relaxed and tightened around the nominal cut value. Figure 5.11 is a scatter plot of the values for RAWTRMS against RCHT. The signal region is in the lower right. There do not appear to be any significant correlations between the two cuts in the region around the cut values.

Figure 5.12: These plots show signal estimates from the bifurcated analysis when pairing RAWTRMS and RCHT. RCHT was held at 0.7 and the RAWTRMS cut was moved in the upper plot. The data in the lower plot are from the bifurcated analysis when the reverse was done. RAWTRMS was held at 40, while RCHT was moved.

Figure 5.12 shows the results from the cut relaxing and tightening for cut pair 7 (RAWTRMS-RCHT). In the top plot, signal estimates from the bifurcated analysis are shown with RCHT held at 0.7 while varying RAWTRMS. The data in the lower plot are from the bifurcated analysis with the RAWTRMS cut value held at 40, while RCHT was moved. The flatness of both plots near the nominal cut values suggest that the bifurcated analysis applied with these two cuts was robust and that the cuts were uncorrelated. This agrees with what was expected from looking at the scatter plot in figure 5.11. It should be noted that signal estimates were used as a measure of robustness for the scheme because it was clear what those estimates represented. The key estimate used in the final analysis was of the contamination level.

We now look at cut pair 8. Figure 5.13 is a scatter plot of the cut values for RAWQRMS against RCHT. The signal region is in the upper right. In the area where the two cuts intersect each other, there is a distribution of events which turn upwards from a RCHT value of 0.6 to 0.8.

Figure 5.14 shows the results from the cut relaxing and tightening for RAWQRMS and RCHT. In the top plot, signal estimates from the bifurcated analysis are shown with RCHT held at 0.7 while varying RAWTRMS. The data in the lower plot are from the bifurcated analysis with the RAWTRMS cut value held at 40, while RCHT was moved. In both cases, the signal estimates from bifurcated analysis varied as the cut did. The sample of events near the intersection of the cuts which were correlated made it difficult for the analysis to produce robust estimates.

Breaking the analysis down into cut pairs, showed the differences in signal estimate and how those differences translate into an understanding of the robustness of the scheme and the cut orthogonality. Since the full analysis was in agreement with a pair of cuts shown to be be orthogonal, it would suggest that strong, orthogonal cuts provide the majority of the
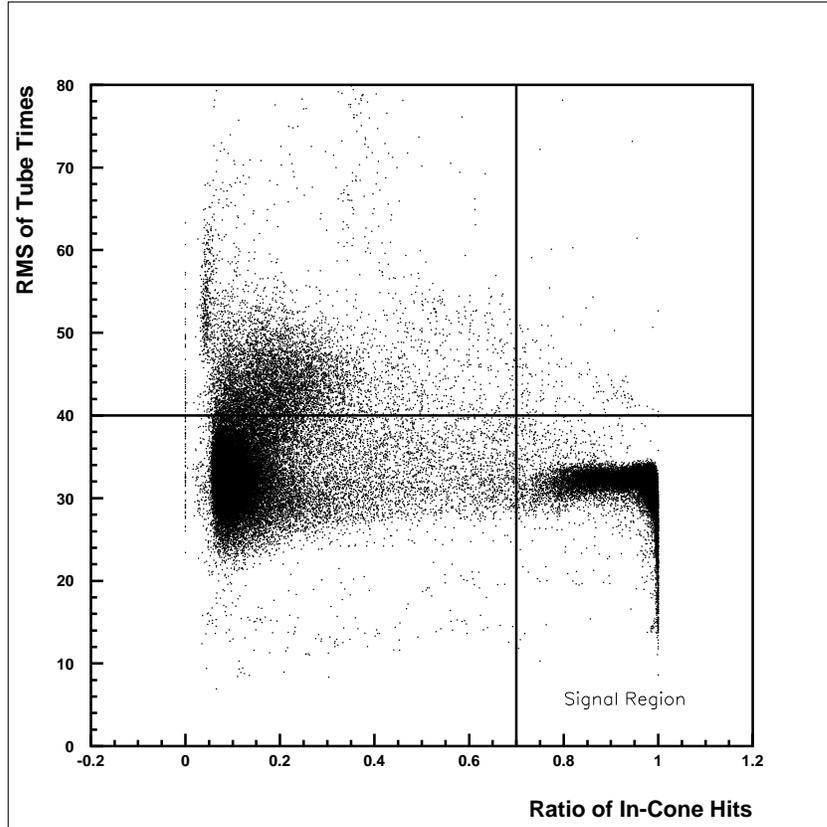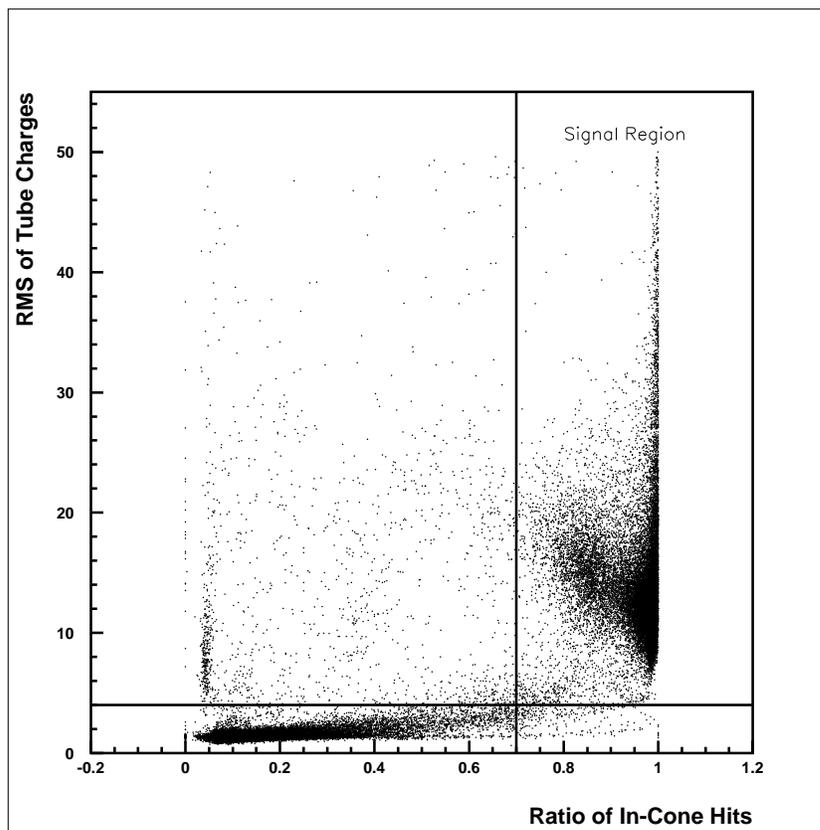
Figure 5.13: The raw tube charge RMS (RAWQRMS) is plotted against the ratio of in-cone hits (RCHT). The lines show the cut values and the signal region.

robustness in the analysis. While a number of other correlations and dependencies can exist in weaker cuts, a reliable result was still possible in the overall analysis. The goal of the entire cut scheme was to remove backgrounds events efficiently while preserving the signal events. The main reason for not reducing the entire scheme to 1 pair of strong robust cuts was the variety of background events in the data. Even the strongest pair of cuts were not sensitive to all types of background events.
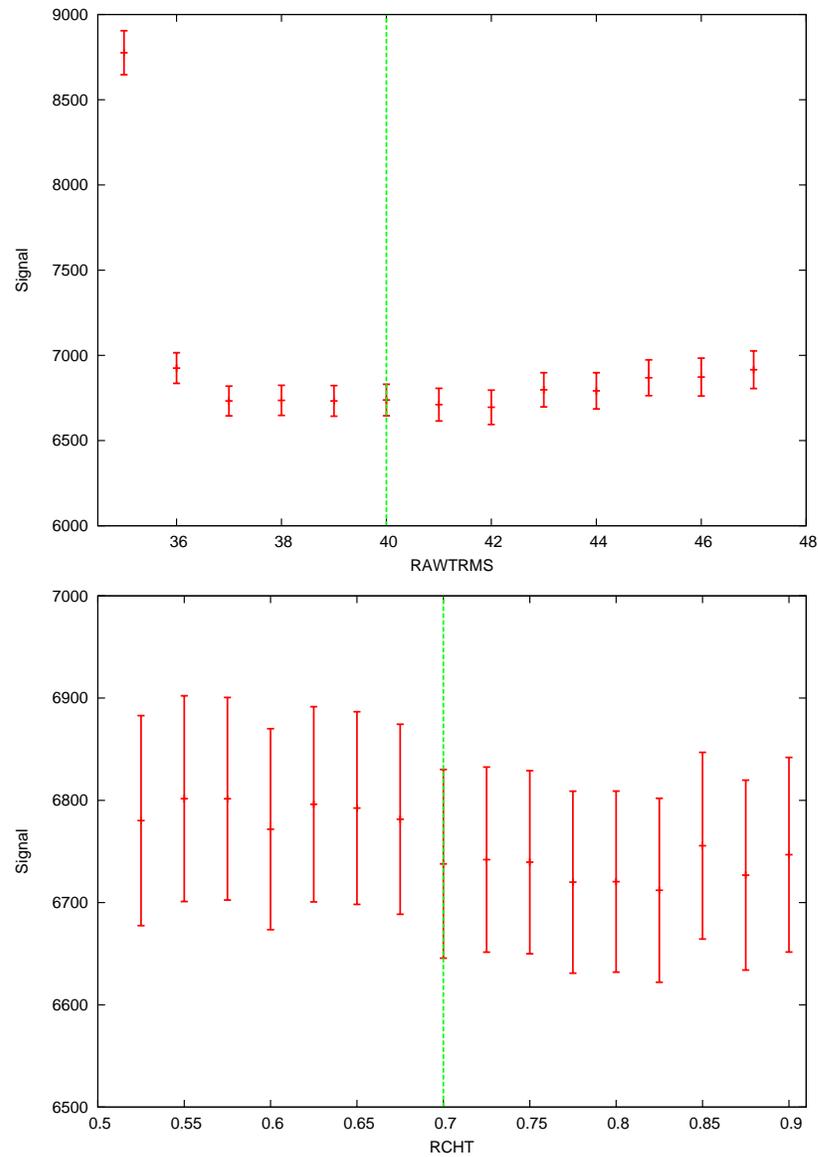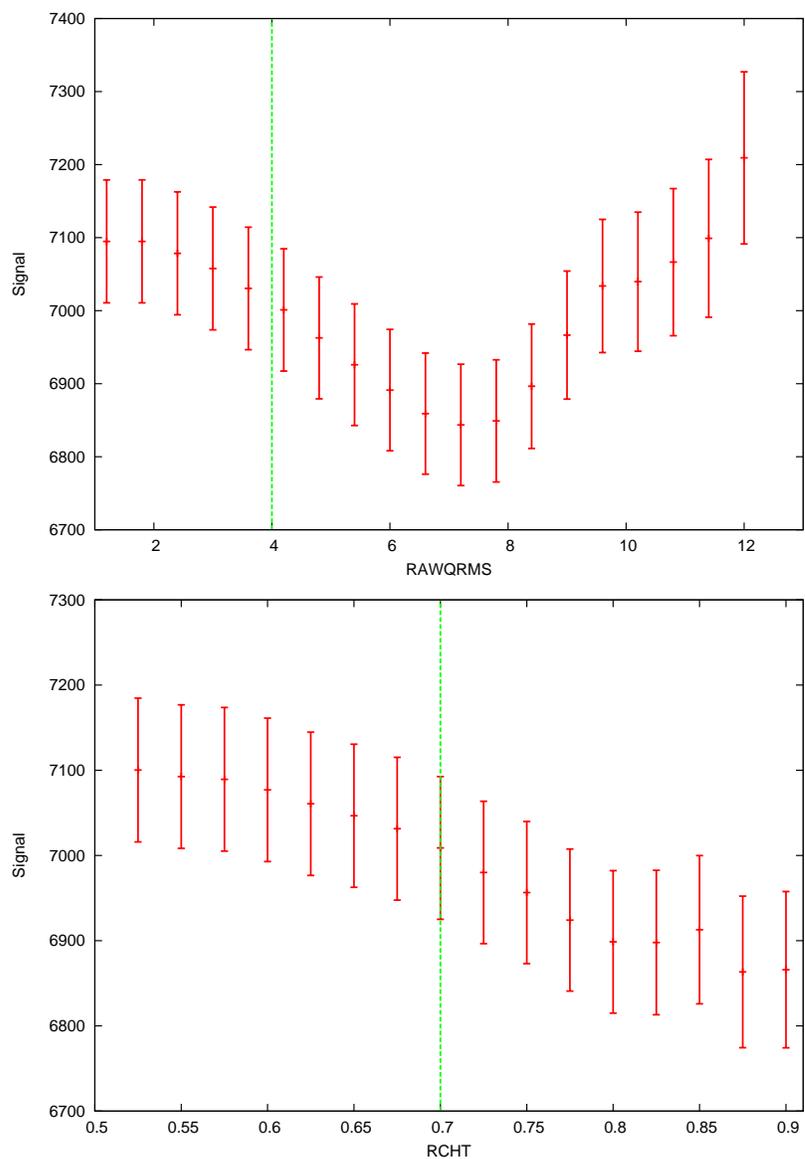
Figure 5.14: These plots show signal estimates from the bifurcated analysis when pairing RAWQRMS and RCHT. RCHT was held at 0.7 and the RAWTRMS cut was moved in the upper plot. The data in the lower plot are from the bifurcated analysis when the reverse was done. RAWTRMS was held at 40, while RCHT was moved.

Cut Name	Original Value	Relaxed Value
PMTHIT	> 500	> 100
NNECK	< 4	< 4
RAWTRMS	< 38 GSU	< 65 GSU
RAWQRMS	> 4.5 ns	> 2.0 ns
FNP	> 2000 p.e.	> 200 p.e.
CITR	> 0.85	> 0.6
RCHT	> 0.7	> 0.5
LDISC	> 0.6	> 0.2
DEDX	> 200 MeV/m	> 50 MeV/m

Table 5.7: A comparison of the relaxed cut values to the original values for the bifurcated analysis.

Full Analysis and Contamination Estimates

With an understanding of paired cut correlations we can turn back to the full analysis. All of the cuts were relaxed in one step to the values given in table 5.7 and the bifurcated analysis reapplied. The neck cut was not relaxed since this cut specifically removes one type of instrumental background (neck events). Since it was the only targeted cut in the scheme and because it has been previously studied it was left at its nominal value for this check. Table 5.8 shows the results from the bifurcated analysis with the relaxed cuts. The contamination estimate remained the same within statistical errors. The distribution of the instrumental backgrounds within cut space was distinct from the distribution of signal events (e.g. CITR and DeDx figure 5.2). This means that a conservative relaxing of the cuts, as was done here, would provide results consistent with the original analysis. It would be possible to be more aggressive with the changes to the cut values in order to allow in a larger number of instrumental backgrounds, but since the estimates were well behaved around the original cut values they were believed to be robust.

It is notable that in relaxing all of the cuts in the full analysis, the signal estimate increased roughly by the number of new events in a. Since a number of the cuts were

	Original Values	Relaxed Values
Pass 1 and Pass 2 (*a*)	21023	21434
Fail 1 and Pass 2 (*b*)	29	22
Fail 2 and Fail 2 (*c*)	23661	23124
Fail 2 and Pass 1 (*d*)	474	607
Signal events	21058 ± 145	21470 ± 147
Contamination	0.31 ± 0.12	0.23 ± 0.16

Table 5.8: Bifurcated results with relaxed cut values. The estimated signal events and contamination agree with the original analysis suggesting that the contamination was insensitive to variations in individual cut values.

relaxed at the same time, these results cannot be directly compared with those in the previous section. In the above cut-pair analysis, the signal estimates varied with varying cut values when a correlation was present. In the full analysis this shift in signal estimate can be attributed to one particular source of events, stopping muons. As described above, the data consist of signal events, instrumental backgrounds as well as physical backgrounds. Since the signal events were defined to be any physics events that pass the cuts, relaxing cuts to accept more physics events would only increase the signal.

While the cuts were developed to have almost zero sacrifice for through-going muons, stopping muons populated cut space often on both the pass side and fail side of the defined cut values. The final estimate of the muon flux at SNO will be compared to the flux seen in the fully simulated dataset. With the consistency of data and MC, stopping muons that fail a cut will fail both in the data and MC. This will not skew the flux comparison.

The aim of the bifurcated analysis was to estimate the number of instrumental backgrounds which contaminated our sample of signal events. After checking the correlations between cut pairs in the analysis as well as relaxing the entire scheme, we have found the contamination estimates to be reliable and robust.

Chapter 6

Signal Extraction

With the datasets and analysis finalized, the next step was to extract physical values and interpret the results. Here we will describe the tools used and results in the direct cosmic ray muon flux and the neutrino-induced muon analysis.

6.1 Neutrino-Induced Muons

Neutrinos produced in the atmosphere by cosmic rays travel generally unimpeded through the Earth. If a muon neutrino, traveling towards SNO, interacts with the rock around the detector and produces a muon, the muon can be observed. The observed muon reveals the existence of a neutrino up until the interaction and its incoming direction. A neutrino coming horizontally towards SNO travels through roughly 160 km of rock whereas one coming directly upwards traveled through 12 700 km. Since neutrino oscillation is dependent on the distance traveled as shown in equation 6.1, we can measure the number of muons detected in a given zenith angle bin, and compare that to the MC to select the best fitting pair of oscillation parameters (Δm^2, $\sin^2 2\theta$) and overall flux normalization, Φ_0. The majority of the initialization and generation of the probability distribution function (PDF) was based on a previous fitter [52], but the method of extracting oscillation parameters and an overall flux normalization was reworked.
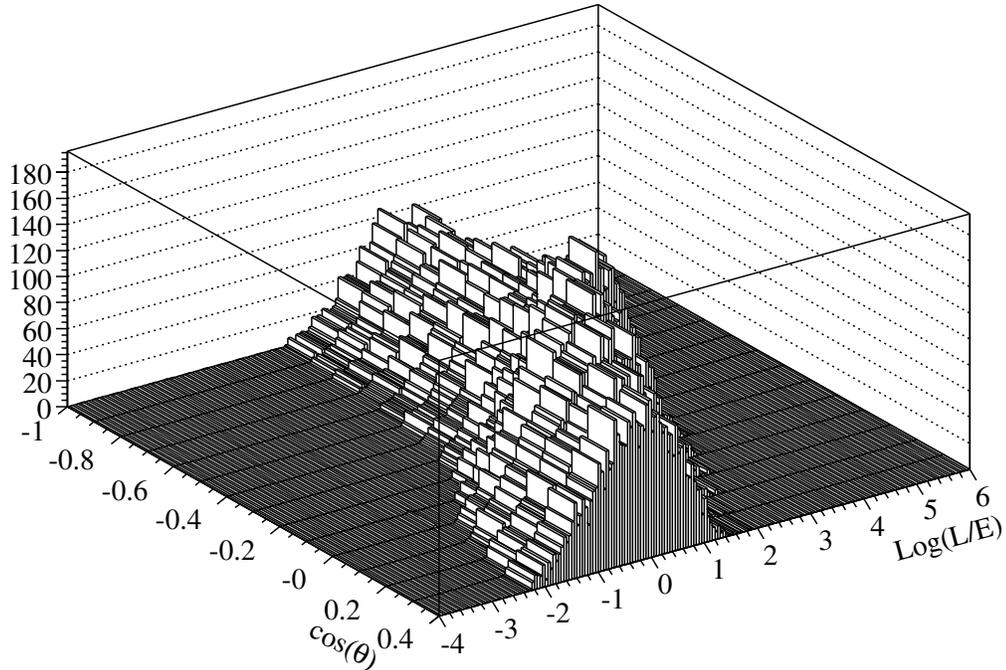
Figure 6.1: 2-D histogram used to generate the zenith angle PDFs for the oscillation analysis. For each $\cos\theta$ bin the spread in L/E was due to the original neutrino energy distribution and the spread in L from the width of the bin.

6.1.1 Method

The first task was to generate the zenith angle distributions from the MC data. These zenith angle PDFs were used as the reference for the log-likelihood calculation done in the fit. A set of PDFs was built by combining MC muon events in a 2-D histogram, binned in $\cos\theta$ and $\log(L/E)$ (figure 6.1), where $\cos\theta$ is the cosine of the incoming muon polar angle and L/E is the path length of the neutrino divided by its energy. Equation 6.1 gives the probability a muon-neutrino of energy E_ν is detected as a muon-neutrino after traveling a distance L, given a pair of mixing parameters $(\Delta m^2, \sin^2 2\theta)$. The 2-D histogram was

used along with equation 6.1 to generate simulated zenith angle distributions. Equation 6.1 shows the probability of a muon type neutrino traveling a distance L and be detected as a muon type neutrino. The Φ_0 factor accounts for an overall normalization of the theoretical flux. The units for L/E are km/GeV and those for Δm^2 are eV^2.

$$P(L/E_\nu, \theta, \Delta m^2)_{\mu\mu} = \Phi_0 \cdot \left\{ 1 - \sin^2 2\theta \cdot \sin^2 \left(\frac{1.27 \Delta m^2 L}{E_\nu} \right) \right\} \tag{6.1}$$

One million zenith angle distributions were made, with 100 sample points in each parameter, spread over the following ranges, $0 < \Delta m^2 < 0.03$ eV^2, $0 < \sin^2 2\theta < 1.0$, and $0. < \Phi_0 < 2.0$. The range boundaries were tested below, but these values were used in the final fit. For each zenith angle distribution, the data were binned into 14 bins in $\cos\theta$. Each of the 14 bin values was calculated by a weighted sum over L/E as shown in equation 6.2, where $H(L/E, \cos\theta)$ is the value of the 2-D histogram at a given $(L/E, \cos\theta)$.

$$BinVal = \frac{\displaystyle\sum_{L/E=0}^{10000} P(L/E_\nu, \theta, \Delta m^2) \times H(L/E, \cos\theta)}{\displaystyle\sum_{L/E=0}^{10000} H(L/E, \cos\theta)} \tag{6.2}$$

The resulting zenith angle distributions are compared to the data by summing the log-likelihood value over the bins in the distribution (equation 6.3). The number of events in a given bin are given by N_{data} and N_{mc} for the data and MC respectively.

$$-2\ln\mathcal{L} = 2\sum N_{data} \ln \frac{N_{data}}{N_{mc}} - (N_{mc} - N_{data}) \tag{6.3}$$

The likelihood value for each parameter triplet (Δm^2, $\sin^2 2\theta$, Φ_0) was stored and a minimum was scanned for, after all the values were calculated. With this array, it was possible to plot 1-D projections of each parameter and 2-D contour plots showing the levels

Δ log-like.	Confidence
2 DOF	Level
2.30	68.3%
4.61	90%
11.83	99.99%

Table 6.1: Δ log-likelihood values used in all contour plots.

of uncertainty around the best fitting point. Table 6.1 shows the values and contour levels used in results plots given below.

We have less statistics than Super-K and MINOS, but our ability to measure the un-oscillated atmospheric neutrino flux is unique and allows the SNO data to be normalized in a model-independent manner. We quote the best fitting oscillation parameters given SNO data and we also use Super-K and MINOS data to help constrain those parameters and reapply the analysis. This reduces the uncertainty on our flux, providing the best flux measurement. This was done by adding 3 extra terms to the likelihood function, one for each constraint ($\Delta m^2_{SK} = 2.1^{+0.6}_{-0.4} \times 10^{-3}$ eV^2, $\sin^2 2\theta_{SK} = 1.000 \pm 0.032$ [53], and $\Delta m^2_{MINOS} = (2.43 \pm 0.13) \times 10^{-3}$ eV^2 [54]). The constraints were assumed to be uncorrelated with each other. MINOS and Super-K were independent experiments. The two constraints from Super-K are shown in figure 6.2 and do not show any strong correlations in the contours. The form of each constraint is generally the difference between the central value quoted from each of the sources, $\hat{x}$, and our trial value, x, divided by the uncertainty, σ, in that central value. This is shown in the following equation.

$$\Delta L = \left(\frac{\hat{x} - x}{\sigma} \right)^2 \qquad (6.4)$$

The common form of the above equation has a factor of two in the denominator. The final modified log-likelihood value shown in equation 6.8 accounts for the factor of two. Since
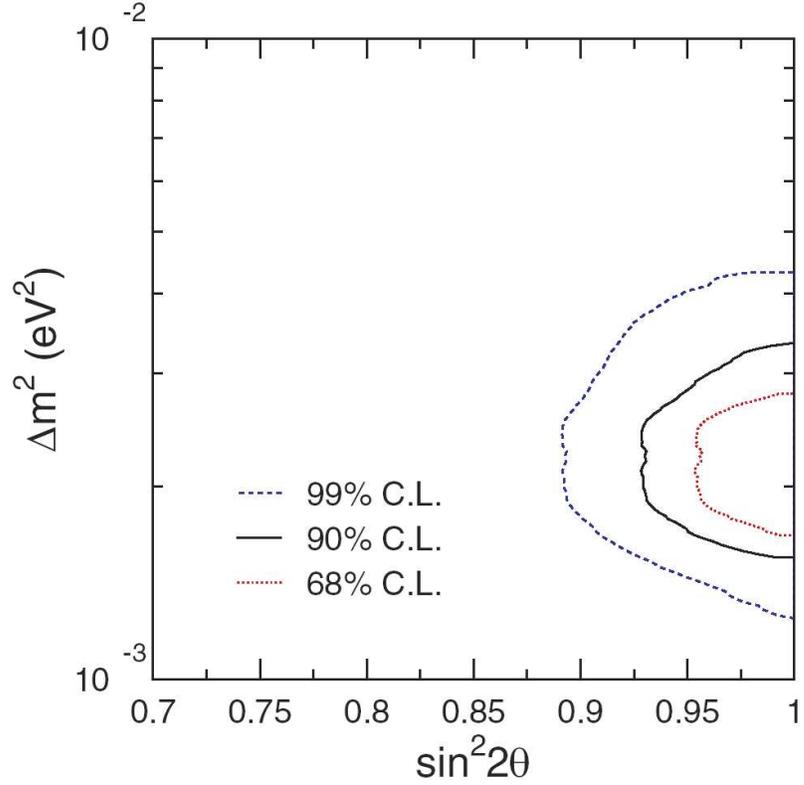
Figure 6.2: This shows the allowed oscillation parameters from the Super-K results [53]. The constraints used in this analysis were extracted from the 1-D versions of the individual oscillation parameter χ^2 distributions. The contours suggest the two parameters were independent.

the errors in the $\Delta m^2{}_{SK}$ value are asymmetric, a modified version of the ΔL correction is used and given in equations 6.5-6.7, where σ_+ and σ_- are the errors in either direction. This treatment of asymmetric errors is taken from [55].

$$\Delta L_{asym} = \left(\frac{\hat{x} - x}{\sigma + \sigma'(x - \hat{x})} \right)^2 \tag{6.5}$$

110

$$\sigma = \frac{2\sigma_+ \sigma_-}{\sigma_+ + \sigma_-} \tag{6.6}$$

$$\sigma' = \frac{\sigma_+ - \sigma_-}{\sigma_+ + \sigma_-} \tag{6.7}$$

The modified log-likelihood value used in the constrained fit is

$$-2\ln\mathcal{L} = \sum 2\left(N_{data}\ln\frac{N_{data}}{N_{mc}} - (N_{mc} - N_{data})\right)$$
$$+ \Delta L^{SK}_{\Delta m^2} + \Delta L^{SK}_{sin^2 2\theta} + \Delta L^{MINOS}_{\Delta m^2} \tag{6.8}$$

6.1.2 Systematic Errors

The treatment of systematic errors in this analysis were described fully in [56]. The systematic effects summarized in table 6.2 include those based on detector model, neutrino cross-section and muon propagation uncertainties.

Detector Systematics

Effects categorized as detector systematics include uncertainties in the reconstructed impact parameter bias and those arising from the cuts imposed on the data. The bias on the reconstructed impact parameter was estimated using the external muon tracking system installed during the NCD phase. The series of wire chambers was placed above the detector and were designed to track and tag muons that enter SNO since no other external muon calibration source existed. A Gaussian was fit to the distribution of the reconstructed impact parameter difference between FTI and the wire chamber system. The Gaussian had a mean of 7.3 cm and a width of 6.1 cm. Since the width was almost as large as the mean bias and the impact parameter was a radial measurement of the distance of closest

approach, the effect was incorporated into the analysis as a systematic uncertainty. The bias value of 7.3 cm was used as the size of systematic uncertainty. This corresponded to an error on the measured flux of $\pm 1.8\%$.

The simulated MC muons helped to understand the data. As described in the section 4.2 a number of cut distributions were shifted when comparing the MC muon and real candidate muon events. The source of these shifts was due to uncertainties and approximations taken in the simulation of charge and time in the detector. The DEDX shift of roughly 4.5% between data and MC muon events was translated into an error of $\pm 2.6\%$ in the observed flux for the model of energy loss that the fitter uses. The shift of roughly 1.5 GSU in the RAWQRMS cut value from data to MC resulted in a $\pm 1\%$ error in the PMT charge model used. Finally the shift observed in the linear discriminant affected the observed flux result by $\pm 2.1\%$. For each of these cuts, the differences between MC and data distributions were used to estimate the impact of an uncertainty on the cut, on the final calculated flux. The higher-level cuts being more effective at removing backgrounds and retaining signal resulted in a larger uncertainty than the lower-level RAWQRMS cut as expected. The combined uncertainty for the detector effects was 3.9% on the observed flux.

Neutrino Cross-Section Systematics

Three interactions were responsible for the neutrino induced muons observed at SNO. Each of these was modeled in `nuance` with associated uncertainties in the processes. The cross-section for the quasi-elastic charged current reaction,

$$\nu_\mu + n \to p + \mu^-, \tag{6.9}$$

has an uncertainty of 10% [41]. Factoring in the fraction of this type of interaction results in an uncertainty of 0.8% in the analysis.

At energies higher than a few GeV, the neutrino interactions probe scales smaller than the nucleus and scatters against individual quarks through

$$\nu_\mu(\bar{\nu}_\mu) + N \rightarrow \mu^-(\mu^+) + X. \tag{6.10}$$

The uncertainties associated with the cross-section of this deep-inelastic scattering are 3% [57]. Accounting for the fraction of these events at SNO a $\pm 2.1\%$ uncertainty was associated with this simulated process.

The first two process account for the majority of the muon production. A small fraction of upward-going muons was also produced in pion resonance production,

$$\bar{\nu}_\mu + p \rightarrow \mu^+ \Delta^0 \rightarrow \mu^+ \pi^0 n. \tag{6.11}$$

Although this was only a small contribution to muon production, the cross-sectional uncertainties are roughly 20% [44]. This corresponds to a 1.9% effect in this analysis. A 3.1% uncertainty was attributed to total the neutrino cross-section model.

Muon Propagation Systematics

This group of systematics include those associated with muon properties. The uncertainty in the muon propagation model was estimated from comparing PROPMU and MUSIC as described in section 2.5.1. MUSIC more completely models muon propagation, but because of the manner in which the code used here was designed PROPMU was used for muon propagation. The results from the two simulations agreed in the energy range we were most concerned with ($\sim$50 GeV). The difference between the two simulations allowed for a $\pm 2\%$ systematic error to be attributed to the propagation model.

The neutrino-induced muon flux varies with the solar cycle. The expanding magnetized

plasma from the Sun decelerates and partially excludes lower energy cosmic rays from the inner solar system [31]. All of the simulations used here assume the Sun was at solar maximum. This was the case for the D_2O and salt phases of the experiment, but through the NCD phase there was a transition to solar minimum. At solar minimum the flux of neutrino-induced muons with an energy less than $\sim$10 GeV increases. Above this energy there is no change. The effect was modeled in [48] and a $\pm$1% systematic error was assigned.

The muon fitter was designed for single through-going muon events. The response of the muon fitter was never fully characterized, but the systematic effect of multiple muons can still be estimated. From MACRO data [58] on multiple muons the estimated fraction of double to single direct cosmic ray muons at SNO is roughly 1%, with 225 double events out of approximately 25000 muon events. For the upward-going sample, the rate of muons is smaller by a factor of 200. Accounting for the lower chance of coincident muons, the systematic error for multiple muons in the atmospheric sample is much less than 1%. The combined uncertainties in the rock density, muon transport and the seasonal variations produced a 2.2% uncertainty on the observed flux.

Implementation

The systematic errors listed above were incorporated directly into the likelihood fitting routine. The general premise was that each systematic error would theoretically impact the number of observed events. Since the zenith angle PDFs used here were built from the 2-D $H(L/E, \cos\theta)$ distributions, the first step taken to understand how the errors would propagate through the analysis was to predict the variation of H with respect to a given systematic. H provides the expected number of muon events for a given $(L/E, \cos\theta)$ and gives the number N_{mc} in the likelihood equation 6.3 and 6.8. Given n systematic errors that each depend on a respective parameter α, the vector $\vec{\alpha} = \alpha_1, \alpha_2, \alpha_3, ..\alpha_n$ can be formed to

Systematic	Error
Detector	
Impact Parameter Bias	$\pm 1.4\%$
Energy Loss Model	$\pm 2.6\%$
PMT Charge Model	$\pm 1.0\%$
Linear Discriminant Cut	$\pm 2.1\%$
Total Detector Model	$\pm 3.8\%$
Neutrino Cross-Section Model	
Quasi-Elastic	$\pm 0.8\%$
Resonance	$\pm 1.9\%$
Deep Inelastic	$\pm 2.1\%$
Total Monte Carlo Model	$\pm 3.1\%$
Muon Propagation Model	
Transport Model	$\pm 2\%$
Seasonal Variation	$\pm 1\%$
Multiple Muons	$\ll 1\%$
Total Propagation Model	$\pm 2.2\%$
Total Systematic Error	$\pm 5.4\%$

Table 6.2: Summary of systematic uncertainties for upward-going muons divided into detector, neutrino model and muon propagation systematic uncertainties. Errors are assumed to be uncorrelated.

combine these systematic errors. The expression for H including the effects of systematics can be rewritten as a linear expansion with respect to $\vec{\alpha}$.

$$H^{sys}(\cos\theta_z, L/E_\nu) \simeq H^0(\cos\theta_z, L/E_\nu) \cdot (1 + \frac{1}{H}\frac{\delta H}{\delta\vec{\alpha}} \cdot \delta\vec{\alpha}) \qquad (6.12)$$

$$H^{sys}(\cos\theta_z, L/E_\nu) \simeq H^0(\cos\theta_z, L/E_\nu) \cdot (1 + \vec{\beta} \cdot \delta\vec{\alpha}) \qquad (6.13)$$

$$\vec{\beta} = \frac{1}{H}\frac{\delta H}{\delta\vec{\alpha}} \equiv \frac{\delta \ln H}{\delta\vec{\alpha}} \qquad (6.14)$$

In the above equations, H^{sys} are the $H(L/E, \cos\theta)$ (figure 6.1) distributions modified

by the list of systematics. In the above description of the systematic errors, an uncertainty on the observed flux was calculated. This was because the H distribution for a given zenith angle bin was filled with the expected number of events and changes in the observed flux affect the fit of that flux to the simulated distributions.

To obtain $\vec{\beta}$, the next step was to parameterize the $H(L/E, \cos\theta)$ distributions, with the following empirical functional form.

$$H_i(x) = A_i \begin{cases} e^{-(\frac{\mu_i - x}{\sigma_-^i})^4} & \text{if } x \leq \mu_i \\ e^{-(\frac{\mu_i - x}{\sigma_+^i})^2} & \text{if } x \geq \mu_i \end{cases} \tag{6.15}$$

The index i refers to zenith angle bins as shown in figure 6.1 and $x \equiv \log_{10} L/E_\nu$. The above form was not symmetric and was based on MC data. After parameterizing each zenith angle bin of H, calculating the effect of the systematic on the zenith angle distribution was a matter of varying the systematic values and extracting the new fit values for H to obtain $\vec{\beta}$. The $\vec{\beta}$ vector functionally scales each zenith angle bin by an amount which accounts for the systematic error. Now for any given expected number of events in the zenith angle distribution, one can factor in the effect of the systematics on the distribution and calculate the resulting final change to the likelihood value. The expanded version of the likelihood which takes into account systematic errors is given as

$$-2 \ln \mathcal{L}_{total} = 2 \sum N_{data} \ln \frac{N_{data}}{N_{mc}^0} - (N_{mc}^0 - N_{data}) - \delta\vec{\alpha}_{min}^T \mathcal{S}^2 \delta\vec{\alpha}_{min}. \tag{6.16}$$

The extra term at the end represents the effect of the systematics on the data in the case that the deviations on the true number are Gaussian distributed [59]. In the above equation, $\delta\vec{\alpha}_{min}$ and $\mathcal{S}^2$ are defined as

$$\delta\vec{\alpha}_{min} = (\sum (N_{data} - N_{mc}^0)\vec{\beta})\mathcal{S}^{-2} \tag{6.17}$$

116

$$\mathcal{S}^2 = \sigma^{-2} + \sum \vec{\beta} \times \vec{\beta}^T \tag{6.18}$$

The diagonal error matrix, σ^{-2}, has entries which represent the size of the systematic error constraints. A calculation of $\vec{\beta}$ and the estimated size of the systematic errors from the above equations show how the systematic uncertainties were incorporated into the likelihood function.

6.1.3 Bias Study

Fake datasets were used to test for biases in the signal extraction process. Given one set of input oscillation and flux values, we looked at the distribution of the extracted values over repeated trials. Each of the resulting distributions were expected to be roughly Gaussian with a mean of the input value and an error on the mean dependent on the number of trials.

We randomly sampled 300 datasets from the $H(L/E, \cos\theta)$ distribution generated by the MC simulations. The number of events for a given dataset therefore formed a Poisson distribution. The expected occurrence of the distribution was the average number of events for one simulated dataset. A known set of oscillation and flux normalization values was applied to the fake datasets, then the signal extraction was performed on these data. The input values used were the best fit MINOS values ($\Delta m^2 = 2.43 \times 10^{-3}~eV^2$, $\sin^2 2\theta = 0.88$) and $\Phi_0 = 1$.

A set of 300 trial runs were processed. The ranges of the fit parameters were extended to check for boundary effects as well as any errors imposed on the analysis because of the fixed parameter ranges that were not previously considered. The extended ranges were $0 < \Delta m^2 < 0.05~\text{eV}^2$, $0 < \sin^2 2\theta < 2.0$, and $0. < \Phi_0 < 2..$

Figures 6.3 and 6.4 are the trial distributions for each parameter. The $\sin^2 2\theta$ and Δm^2 distributions are relatively flat and not at all Gaussian. The Δm^2 distribution had a
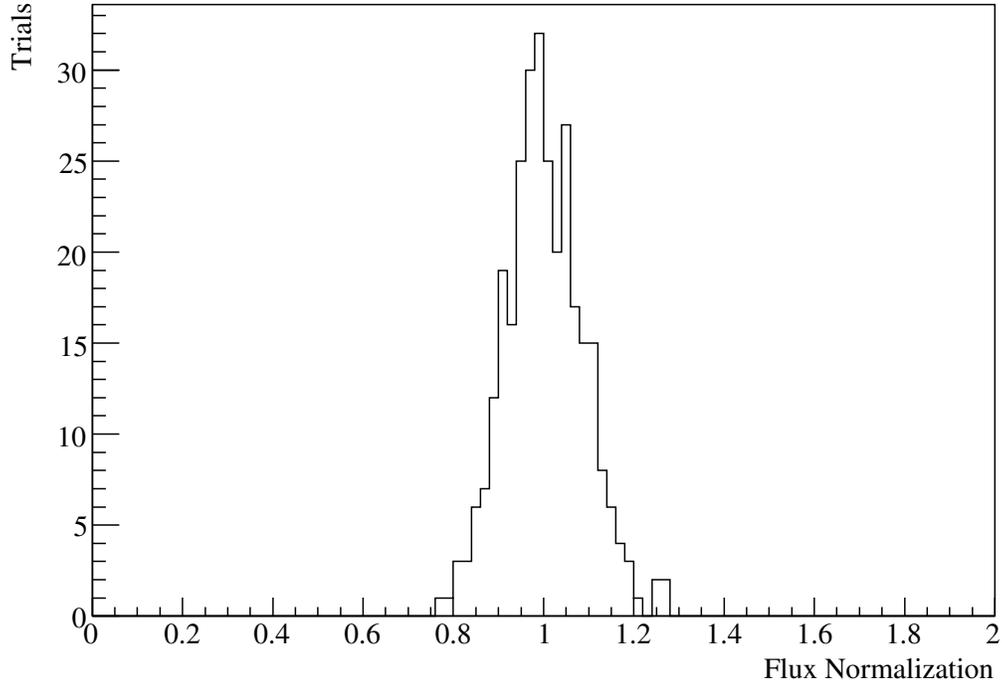
Figure 6.3: Distribution of best fit flux values over 300 trials given three-phase statistics. A Gaussian fit gives a mean and sigma of 1.00±0.08.

small peak around the input value, but many of the trials were spread through the parameter range. The extracted flux distribution is roughly Gaussian and agrees with the input value of $\Phi_0 = 1$.

The $\sin^2 2\theta$ distribution shows an excess in the upper boundary bin. The tendency to fit for unphysical values of $\sin^2 2\theta$ was seen by Super-K [53]. The strongest constraint on $\sin^2 2\theta$ came from the two bins with the smallest $\cos\theta$ values. In these two bins the path length through the earth was greatest and the statistics we had were the smallest. The fitter preferred the largest possible value for $\sin^2 2\theta$ to compensate for the small statistics. To check that this was a statistical effect, we took one trial and increased the number of events

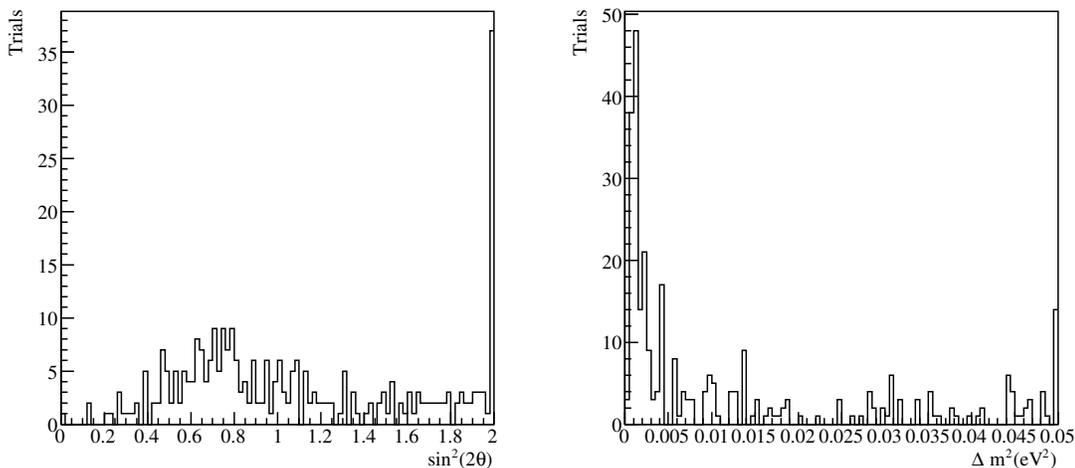Figure 6.4: Distribution of best fit $\sin^2 2\theta$ and Δm^2 values over 300 trials given three phase statistics. The $\sin^2 2\theta$ plot shows the large number of trials that prefer the maximum value. The Δm^2 distribution has a small peak around the input value and a large spread of trials over the rest of the parameter range. A Gaussian fit to only peaks in each plot gives a mean and sigma of 0.81±0.4 for $\sin^2 2\theta$ and $(1.56 \pm 0.79) \times 10^{-3}$ eV^2 for Δm^2.

by a factor of 100. Figure 6.5 shows the 2-D and log-likelihood distribution for increased statistics run. They both show that with higher statistics $\sin^2 2\theta$ was bound on both sides. The extracted flux in our analysis appears to not have a significant bias while the fit for $\sin^2 2\theta$ generally prefers the highest allowable value because of our limited statistics.

6.1.4   Results

Over the three phases at SNO, 514 muon events were observed over 1229.26 days of live-time with $\cos\theta < 0.4$. The resulting observed muon flux from those events was $(2.48 \pm 0.25) \times 10^{-13}$ cm^{-2}s^{-1}sr^{-1}. In the region above horizontal $(0. < \cos\theta < 0.4)$ we observed 201 events corresponding to an observed muon flux of $(3.48 \pm 0.25) \times 10^{-13}$ cm^{-2}s^{-1}sr^{-1}. Given current measurement of oscillation parameters, those 201 events were considered to
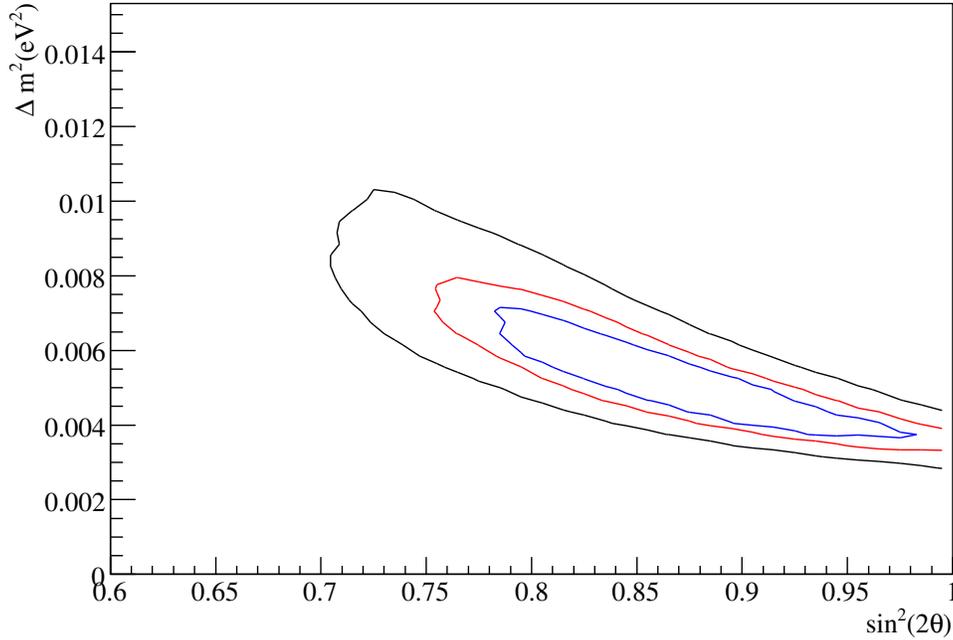
Figure 6.5: A contour plot of the log-likelihood values for a $100 \times$ MC sample dataset. The contour closes on the high $\sin^2 2\theta$ value side suggesting that the tendency for the $\sin^2 2\theta$ to fit the highest value was an effect of our limited statistics.

be from neutrinos that have not traveled a long enough distance to oscillate. The zenith angle distribution of events with $-1.0 < \cos\theta < 0.4$ was used to constrain neutrino oscillation parameters and the flux of neutrinos at SNO depth. The best fit values to our data were $\Delta m^2 = (2.6 \pm 2.0) \times 10^{-3}$ eV^2, $\sin^2 2\theta = 1.0 \pm 0.1$ and $\Phi_0 = 1.22 \pm 0.10$. The PDF with these parameter values is plotted with the data muon events and the PDF for no neutrino oscillations in figure 6.6. Using Super-K and MINOS constraints on the oscillation parameters we find the overall flux normalization to be 1.22 ± 0.09 times that expected by the Bartol group's 3D calculations. Figures 6.7 - 6.8 show the Δ log-likelihood values as a function of each parameter. The profile for a given parameter was built by scanning for the best
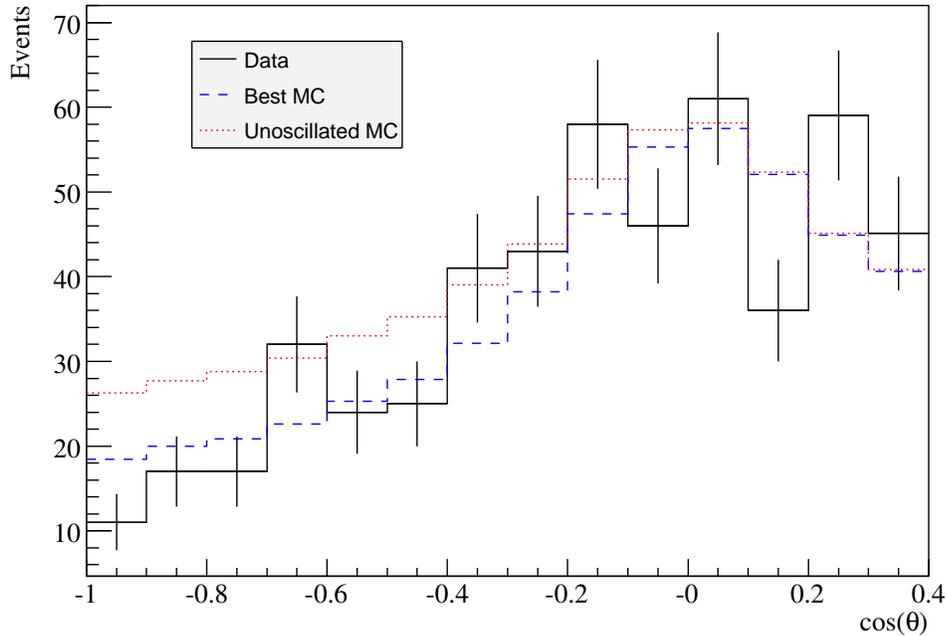
Figure 6.6: Zenith angle histogram of data muon events with statistical error bars (black solid), the best fitting MC (blue dashed) and no oscillations MC (red dotted).

log-likelihood value over the other two. The uncertainties associated with each of the best fit values were obtained from the 1σ or 68% confidence level of the profile.

Figure 6.7 shows a simple monotonic decrease in the log-likelihood value as $\sin^2 2\theta$ goes to 1. The plot of Δm^2 in the same figure contains a minimum and a slightly ragged curve. These types of features were present in other analyses done by Super-K and were attributed to the statistics in our sample. As mentioned above, a large part of oscillation parameters constraints came from the bins where the muons were coming up through the earth and where we had the fewest events.

The profile of the overall flux normalization (figure 6.8) has a single minimum and is asymmetric about that minimum. It is more difficult for the fitter to constrain higher

PSfrag replacements

$\hat{q}$

$\hat{t}$

Figure 6.7: Δ log-likelihood values plotted against $\sin^2 2\theta$ and Δm^2 with SNO data. While the distribution for Δm^2 has a minimum, the $\sin^2 2\theta$ plot shows the tendency of the fitter to chose the highest possible value.
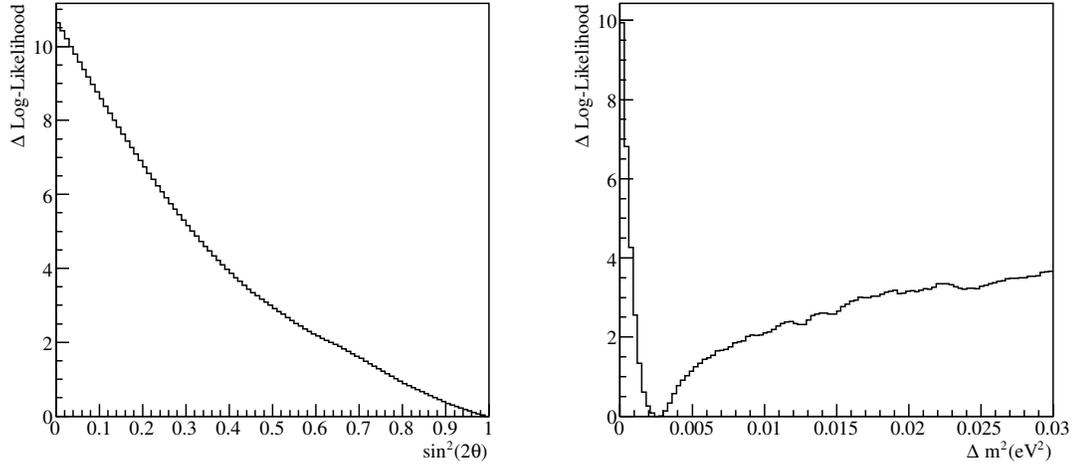
fluxes because those higher fluxes still fit the data with an appropriate oscillation pair. The oscillation parameters reduce the number of observed events and accommodate for the higher fluxes. The lower fluxes do not fit the data as well and cannot rely on the oscillations to bring them in line with the data since there are no oscillation scenarios in which the events would increase. Our statistically limited analysis and loose constraints on the mixing parameters result in the weaker limits on higher fluxes. Using Super-K and MINOS mixing values and uncertainties, strengthens our upper constraint on the flux normalization as shown in figure 6.11.

The features described in the 1-D plots of the parameters are also seen in the 2-D contour plots. Figures 6.9-6.10 show contour plots of the Δ log-likelihood values in each of 3 parameter planes ($\Delta m^2 - \sin^2 2\theta$, $\Phi_0 - \sin^2 2\theta$, $\Phi_0 - \Delta m^2$). The slightly ragged Δm^2 profile translated into rough contours in both of the 2-D Δm^2 plots and reflect our difficultly
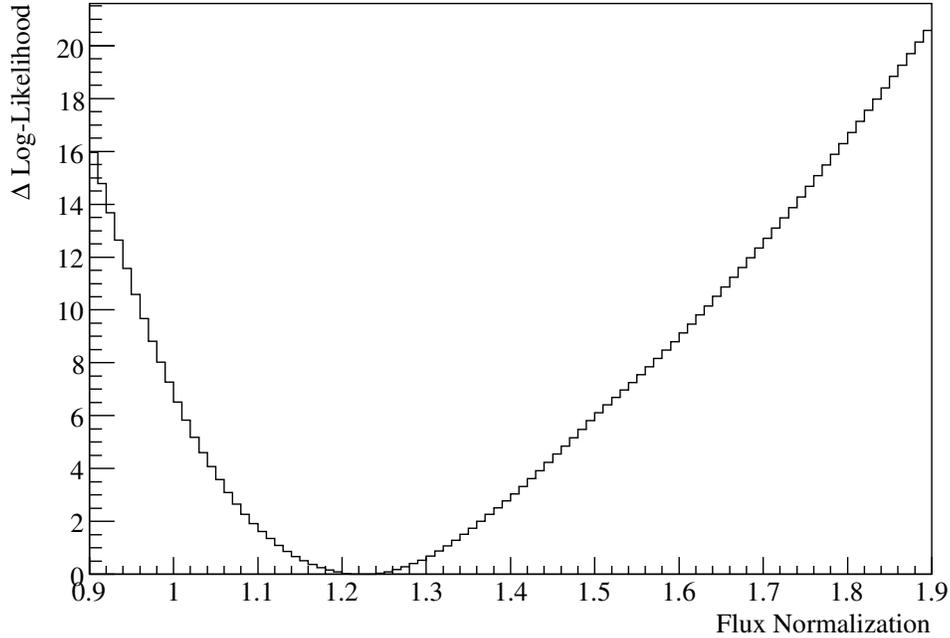
Figure 6.8: Δ log-likelihood values plotted against the overall flux normalization with SNO data. The best fit flux point and 1 σ uncertainties come from these data.

in constraining it. To eliminate discrete sampling as a cause of the rough contours, the code was reapplied with a 50% increase in sample points in Δm^2. No improvement to the smoothness of the contour was seen.

The likelihood value for our best fitting parameter values includes the impact of systematic effects as described above. Without the correction for systematic errors the best fitting flux normalization is 1.24 $\pm$ 0.08. Figure 6.11 shows the Δ log-likelihood values plotted against the overall flux normalization with and without the pulls correction. The pull correction shifted the distributions by roughly 2% towards higher fluxes and increased the uncertainty in the flux by 2%.
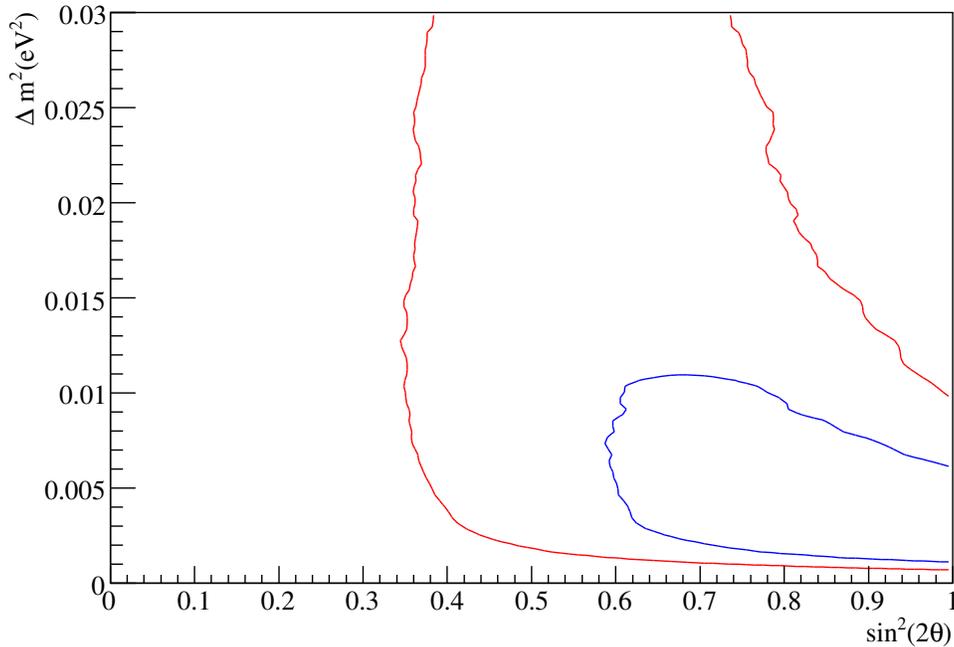
Figure 6.9: Contour plot of Δ log-likelihood values in $\Delta m^2 - \sin^2 2\theta$ space with SNO data. The innermost contour is the at the 68.3% level and the surrounding contours follow the values given in table 6.1. With SNO statistics it was not expected for the contours to close on the high $\sin^2 2\theta$ side.

6.2   Direct Cosmic Ray Muons

The direct cosmic ray muons produced in the atmosphere travel down through the overburden at SNO. The method of presenting the observed intensity is through a vertical muon intensity plot. These muons were seen with an incoming zenith angle of up to $66°$. The distance the muon traveled through rock is an increasing function of zenith angle. The idea is to plot the vertical intensity as a function of distance traveled through standard rock. This allows for a simpler comparison of data from different experiments at varying depths and rock type.
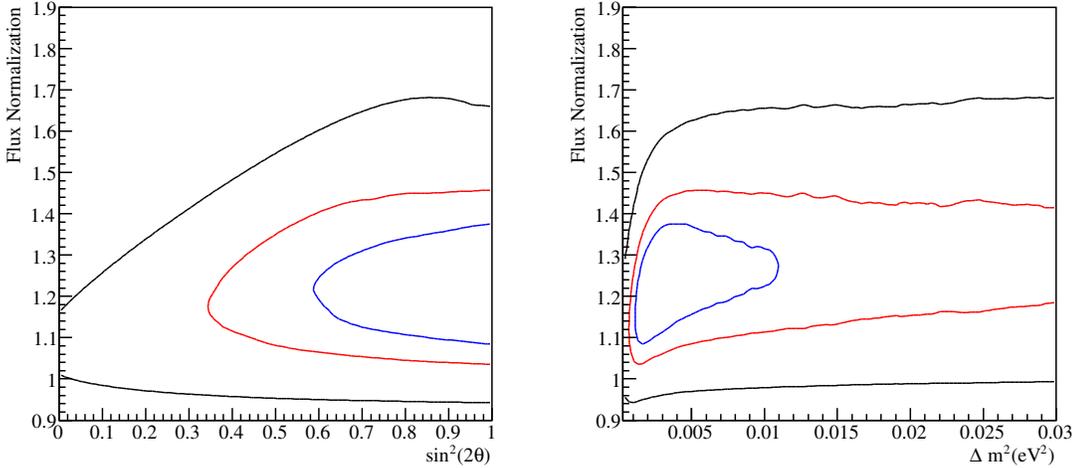
Figure 6.10: Contour plots of Δ log-likelihood values in $\Phi_0-\sin^2 2\theta$ space and $\Phi_0-\Delta m^2$ space with only SNO data. The inner-most contour is the at the 68.3% level and the surrounding contours follow the values given in table 6.1. While the flux was well constrained, these plots show that with SNO statistics the constraints on Δm^2 and $\sin^2 2\theta$ were not as strong.

For each zenith angle bin, we have an observed number of muon events. In order to convert that to a vertical intensity or a number of events per unit area, time and solid angle, the livetime (L), area of the detector (A) and the solid angle covered (Ω) must be considered. The number of muons observed in each $\cos\theta$ bin is converted to a vertical intensity through

$$I_\mu^v(\cos\theta) = \frac{1}{LA\Omega\eta} \sum_{i=1}^{N} \cos\theta_i. \tag{6.19}$$

Here, η is the detection efficiency. The sum is over the number of events in the zenith angle bin and the $\cos\theta$ accounts for the intrinsic variation in the intensity with zenith angle. The incoming muon angle was converted to a depth (x_{SNO}) given the flat overburden at SNO. While the physical depths were converted into metres water equivalent in order to

125

PSfrag replacements

$\hat{q}$

$\hat{t}$
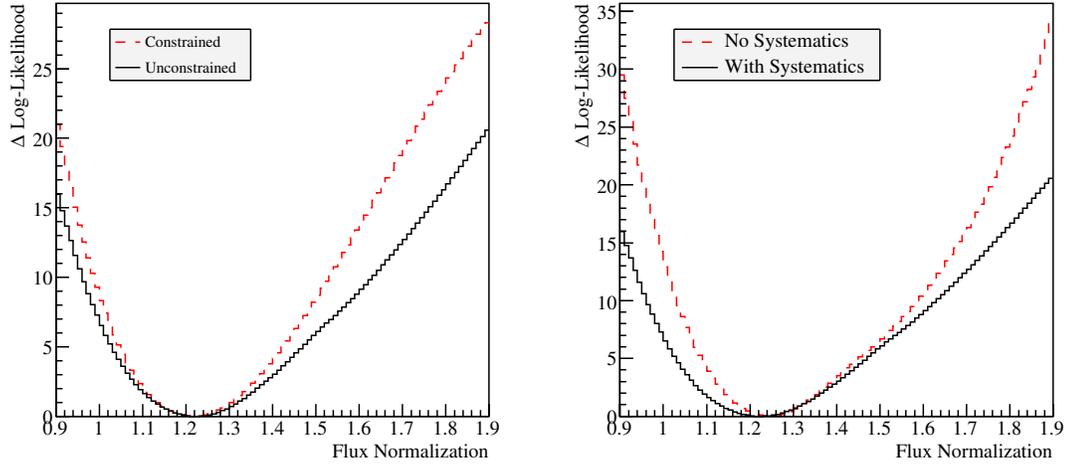
Figure 6.11: Δ log-likelihood values plotted against the overall flux normalization. The right plot shows the effect of the systematic pulls function. The general trend was to decrease the Δ log-likelihood values making our final estimate on the flux less constrained. The left plot shows the results from constraining the oscillation parameters with Super-K and MINOS data. This improves our constraint on the overall flux normalization.

make comparisons with other experiments in different densities and types of rock, the depth values needed a further, second order correction to account for the effects of varying rock composition. The depth values were converted to depth in standard rock ($CaCO_3$ and $\rho = 2.65$ g/cm^3) in units of m.w.e. through

$$x_{standard} = 1.015 x_{SNO} + \frac{x^2_{SNO}}{4 \times 10^5 m.w.e.}. \tag{6.20}$$

The above conversion was derived from simulating muon propagation through both standard rock and norite rock with MUSIC and comparing the vertical intensities at a range of depths.

The data can now be plotted as a vertical muon intensity as a function of standard rock

126

Parameter	SNO MC	SNO	SNO (Fixed)	SNO (Corr.)	LVD
$I_0(\times 10^{-6})$	1.15 ± 0.20	0.93 ± 0.05	2.32 ± 0.10	3.72 ± 0.53	2.15 ± 0.08
x_0 (km.w.e.)	1.34 ± 0.31	2.44 ± 0.31	1.09 ± 0.10	1.15 ± 0.04	1.16 ± 0.06
α	2.24 ± 0.42	5.62 ± 0.40	1.77	2.26 ± 0.30	1.77 ± 0.05
χ^2	61.32	38.9	126.8	149.1	$-$

Table 6.3: Results from the vertical muon intensity depth parameterization from both this analysis and existing results. SNO data reaches depths where the upward-going muons are a significant contribution. This shifted the overall α parameter as shown by the SNO corrected results. The correction accounted for the flux of atmospheric-neutrino induced muons in the sample. Fixing it to the value found by LVD brings the other two parameters to agree with existing measurements.

depth as in figure 6.13. The distribution has traditionally ([15],[58]) been parameterized by

$$I_\mu^v(x) = I_0 \left(\frac{x_0}{x}\right)^\alpha e^{-\frac{x}{x_0}}. \tag{6.21}$$

The standard depth we calculated is x and we fit for I_0, the overall normalization constant, x_0 and α. The spectral index α is highly correlated with that of the primary cosmic ray spectrum (γ in equation 1.1) since each depth bin we have effectively sets a minimum muon energy required to reach that particular depth.

Table C.1 in appendix C shows the range of depth bins, the observed events in each depth bin and the vertical intensity for the MC and real data. The bin sizes increased in two steps with depth to help with the extreme attenuation of the muon signal at larger depths. Up to 7600 m.w.e. the bin widths were 50 m.w.e., from 7600 m.w.e to 10 000 m.w.e. they were 200 m.w.e. and above 10 000 m.w.e. they were 500 m.w.e..

Figure 6.12 shows the vertical muon depth intensities for SNO, LVD, and MACRO data. The data are consistent where they overlap in depth. The line shows the function (equation 6.21) with best-fit value for I_0 and x_0 with $\alpha = 1.77$. At depths up to roughly 12000 m.w.e., SNO data appears to agree with the fit values from LVD. Figure 6.13 shows the vertical

muon depth intensity for MC, data with the best-fit function and data with a correction for the neutrino-induced muon flux. The best-fit function shown in figure 6.13 was from the full fit to SNO data where all three parameters were unconstrained. The data from 7000 - 12500 m.w.e. are consistently lower than expected values from the MC. All of the errors provided are statistical. A more detailed study of the systematic uncertainties in the simulation code which were mentioned in chapter 2 could help resolve the difference.

Table 6.3 summarizes the fit results and compares them to existing results. The results from the full 3-parameter fit to SNO data differed significantly from those of LVD [15]. This was due to new SNO data at larger depths where the parameterization was not the best representation of the vertical intensity. The parameterization and α are derived from the primary cosmic ray spectrum and characterize those direct cosmic ray muons produced in the atmosphere. SNO data above the line at depths greater than 12000 m.w.e. can be due to the atmospheric neutrino-induced muons produced in the rock surrounding SNO which have comparable numbers to the downward events at those large depths (larger zenith angles). This flattens the spectrum at greater depths, with the data deviating from the parameterization.

Two steps were taken to check that the data at larger depths were the cause of the disagreement. First, fixing $\alpha = 1.77$ showed that SNO data at depths comparable to LVD provided consistent fit results (figure 6.12). With the direct measurement of the atmospheric neutrino-induced muons described in the previous section, it was also possible to subtract the expected number of neutrino-induced muon events from the direct cosmic sample. This is shown as the green data in figure 6.13. The expected number of events from the atmospheric neutrino-induced MC and scaling those numbers by our best fit value to the overall normalization provided estimates of the correction in each depth bin. With this correction the best fit values to the SNO data were more consistent with the existing
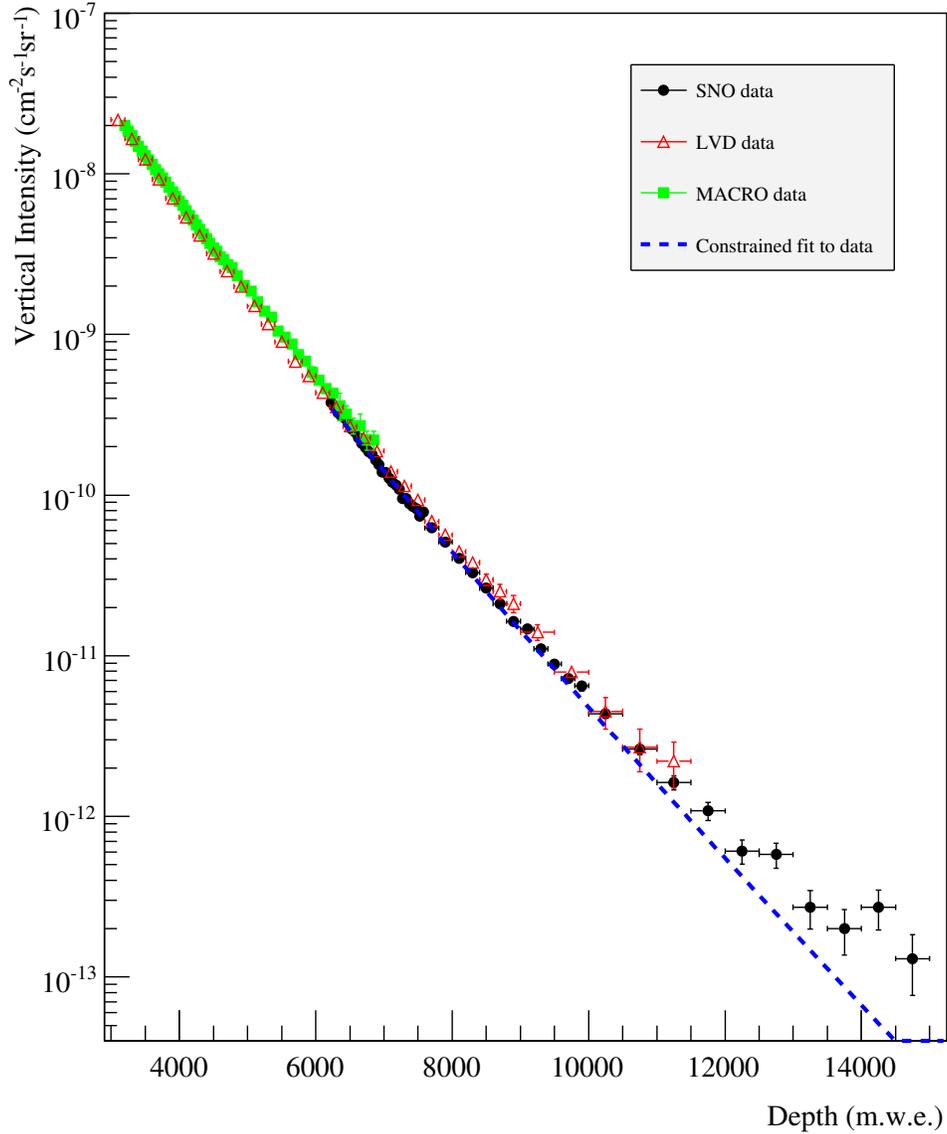
Figure 6.12: Vertical muon depth intensities for SNO, LVD, and MACRO data. The data are consistent where they overlap in depth. The line shows the function (equation 6.21) with best-fit value for I_0 and x_0 with $\alpha = 1.77$. With α derived from the primary cosmic ray spectrum, the SNO data above the line at depths greater than 12000 m.w.e. can be due to the upward-going muons which have comparable numbers to the downward events at those large depths (larger zenith angles).
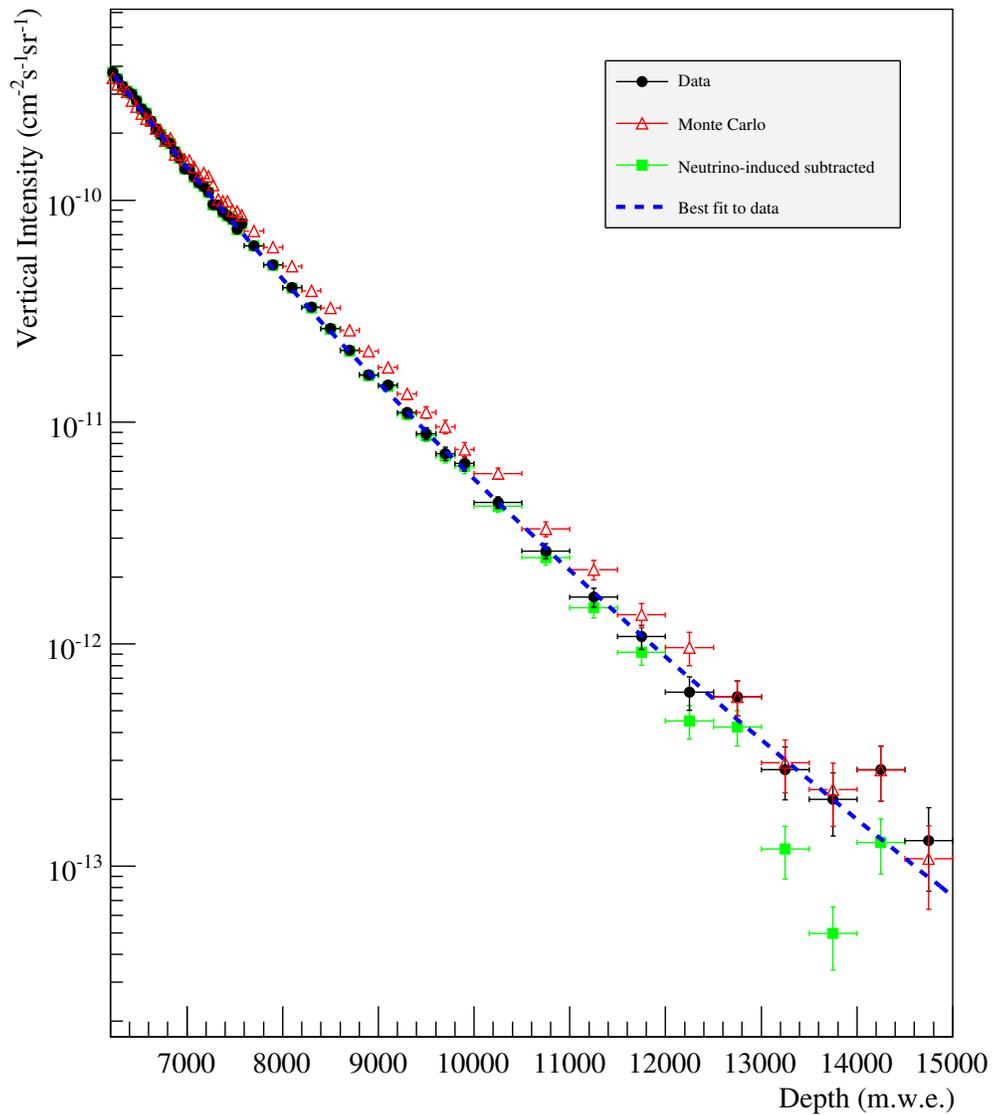
Figure 6.13: Vertical muon depth intensities for SNO data and MC. The data from 7000-12500 m.w.e. are consistently lower than expected values from the MC. The green data show SNO data that were corrected for the neutrino-induced muons. The error bars are represent only the statistical errors. The blue dashed line shows the function (equation 6.21) with best-fit values to the data.

values.

The fit parameters themselves are strongly correlated with each other allowing for significant uncertainties in the fit parameters. The largest source of systematic error were multiple muons. With no multiple muon fitter at SNO, these events were not characterized in this analysis. Previous work in [33], showed that the multiple to single muon fraction at SNO is roughly 3%.

6.3   Summary

We observed 514 muon events with $\cos\theta <$ 0.4. The resulting observed muon flux from those events was $2.48 \pm 0.25 \times 10^{-13}$ cm^{-2}s^{-1}sr^{-1}. In the region above horizontal ($0 < \cos\theta <$ 0.4) we observe 201 events corresponding to an unoscillated observed muon flux of $3.48 \pm 0.25 \times 10^{-13}$ cm^{-2}s^{-1}sr^{-1}. This data with $0 < \cos\theta <$ 0.4 from the three phases of SNO provide a measurement of the underground flux of atmospheric neutrinos which were unaffected by neutrino oscillations given current best estimates of oscillation parameters. Previous measurements of the underground neutrino flux were only estimated by simultaneously fitting for both the overall flux and the oscillation parameters. The best fit oscillation and flux values to our data were $\Delta m^2 = (2.6 \pm 2.0) \times 10^{-3}$ eV^2, $\sin^2 2\theta =$ 1.0 ± 0.1 and $\Phi_0 = 1.22 \pm 0.10$. The no oscillation hypothesis was ruled out at the 99.8% confidence level. The oscillation values are consistent with Super-K and MINOS values ($\Delta m^2{}_{SK} = 2.1^{+0.6}_{-0.4} \times 10^{-3}$ eV^2, $\sin^2 2\theta_{SK} = 1.000 \pm 0.032$ [53], and $\Delta m^2{}_{MINOS} = (2.43 \pm 0.13) \times 10^{-3}$ eV^2 [54]). Using those Super-K and MINOS values to help constrain our measurement of the overall flux normalization gives 1.22 ± 0.09 times that expected by the Bartol group simulations [42].

The vertical muon depth intensity was fit with equation 6.21 and the resulting best fit parameters were $I_0 = (0.93 \pm 0.05) \times 10^{-6}$ cm^{-2}s^{-1}sr^{-1}, $x = 2.32 \pm 0.10$ km.w.e. and

$\alpha = 5.62 \pm 0.02$. We observed 77376 downward going muon events with $\cos\theta > 0.4$. This gave a vertical depth range of 6.2 km.w.e to 15.5 km.w.e.. These data extended the depth of available world data on direct cosmic ray muons by roughly 4 km.w.e. [15]. The results from the full 3-parameter fit to SNO data differed significantly from those of LVD. In two separate checks, fixing $\alpha = 1.77$ and correcting for the neutrino-induced muons in the direct cosmic ray sample showed that SNO data at depths comparable to LVD provided consistent fit results. The fit parameters themselves are strongly correlated with each other allowing for significant uncertainties in the fit parameters.

Chapter 7

Discussion

This thesis presented the analysis tools and results from the study of through-going muons in all the collected data at the SNO experiment. The muon events at SNO were the subject of two analyses. The first was performed on the downward-going muons and the second on upward-going muons.

7.1   This Thesis

Both through-going muon analyses were dependent on development of the tools described in this thesis. The definition of event cuts described in this thesis helped to isolate the sample of signal muon events from both physical and non-physical backgrounds. The bifurcated analysis provided a means to estimate the non-physical backgrounds contaminating the sample. With no explicit model of non-physical backgrounds at SNO, reported values for the contamination of background events were previously only estimates. The bifurcated analysis relied only on the cuts and data itself to provide robust calculations of the contamination. Finally, the algorithms produced to extract physical quantities from the data in both analyses were presented.

The resulting observed muon flux from the 514 events with $\cos\theta < 0.4$ was $2.48 \pm 0.25 \times 10^{-13}$ cm^{-2}s^{-1}sr^{-1}. The best fit oscillation and overall flux normalization values to our data were $\Delta m^2 = (2.6 \pm 2.0) \times 10^{-3}$ eV^2, $\sin^2 2\theta = 1.0 \pm 0.1$ and $\Phi_0 = 1.22 \pm 0.10$. The no-oscillation

hypothesis was ruled out at the 99.8% confidence level. The oscillation values are consistent with Super-K and MINOS values ($\Delta m^2_{SK} = 2.1^{+0.6}_{-0.4} \times 10^{-3}$ eV^2, $\sin^2 2\theta_{SK} = 1.000 \pm 0.032$ [53], and $\Delta m^2_{MINOS} = (2.43 \pm 0.13) \times 10^{-3}$ eV^2 [54]). Using those Super-K and MINOS values to help constrain our measurement of the overall flux normalization gives 1.22 ± 0.09 times that expected by the Bartol group simulations [42].

All atmospheric neutrino oscillation experiments and simulations require an estimate of the actual flux of atmospheric neutrinos. This result from the three phases of SNO provides a measurement of the underground flux of neutrinos which were unaffected by neutrino oscillations given current best estimates of oscillation parameters. Previous measurements of the underground neutrino flux were only estimated by simultaneously fitting for both the overall flux and the oscillation parameters. Given that neutrino oscillations act to reduce the number of detected events, the uncertainties that arise in those analyses were significant. The direct measurement of the unoscillated neutrino flux provided here gives an experimental means to constrain this flux for future analyses.

Not only do the data presented in this thesis include the first sample of unoscillated atmospheric neutrino-induced muon events, they also probe a part of the primary cosmic ray spectrum which is not well sampled. As mentioned in the introduction this intermediate energy range ($10^{11} - 10^{14}$ GeV) of the cosmic ray spectrum is difficult for the two standard types of experiments to observe. These particles are too energetic to be detected by the single detectors and too low energy to produce large showers. That energy range is best probed underground by atmospheric neutrino-induced muon detectors. Converting the overall flux normalization reported here back into a quantity that will help to fill in this intermediate energy range will require a separate analysis involving all aspects of the muon simulations and propagation calculations.

We observed 77376 downward going muon events with $\cos \theta > 0.4$. This gave a vertical

depth range of 6.2 km.w.e to 15.5 km.w.e.. The vertical muon intensities calculated here can help to improve uncertainties in both the cosmic ray spectrum and theoretical muon propagation algorithms. These data extended the depth of available world data on direct cosmic ray muons by roughly 4 km.w.e.. The data at these deeper levels can be used to help constrain the primary cosmic ray spectrum. Full 3D muon propagation simulations like MUSIC have not had experimental data to constrain the simulation through the depths that are accessible to SNO. Being able to compare the simulated results to those present here, will help to refine assumptions and cross-sections that are in the simulation.

The resulting best fit parameters were $I_0 = (0.93 \pm 0.05) \times 10^{-6}$ cm^{-2}s^{-1}sr^{-1}, $x = 2.32 \pm 0.10$ km.w.e. and $\alpha = 5.62 \pm 0.02$. These results differed significantly from those of LVD. This was due to new SNO data at larger depths where the parameterization was not the best representation of the vertical intensity. In two separate checks, fixing $\alpha = 1.77$ and correcting for the neutrino-induced muons in the direct cosmic ray sample showed that SNO data at depths comparable to LVD provided consistent fit results. The fit parameters themselves are strongly correlated with each other allowing for significant uncertainties in the fit parameters.

The work in this thesis will be published as part of the final overall muon analysis at SNO. In particular, the signal extraction code and results from the upward-going neutrino-induced muon analysis will be those presented in the SNO through-going muon paper.

7.2 Future SNO Analyses

There is still information in the stopping and multiple muon analyses. The current muon fitter was optimized to reconstruct single through-going muon tracks. Multiple muon events and stopping muons at SNO can provide additional information on direct cosmic rays and neutrino oscillations. Multiple muons at large depths can be used to help understand pri-

mary cosmic ray compositions. Both MACRO and LVD results include single and multiple muons. While the depth of the SNO experiment was greater than those two detectors, it would be difficult for SNO to resolve multiple muons. MACRO and LVD have multiple segmented detectors capable of distinguishing multiple muons events. With only 1 target volume at SNO, any fitter that is developed for multiple muons will have an efficiency of tracking multiple muons that increases with the distance between the muon tracks.

Muons that stop in the detector originate from neutrinos with lower energy ($\sim$10 GeV) than those muons that travel all the way through ($\sim$100 GeV). With the energy dependence of neutrino oscillations as seen in equation 6.1, being able to reconstruct both stopping muons and through-going muons can allow for two simultaneous oscillation studies. Taking a ratio of the two fluxes can also cancel out many of the theoretical uncertainties in the flux normalization and interaction cross sections. Super-K has analyzed their stopping and through-going muon and report results in [60]. The group measured a stopping to through-going ratio of 0.22$\pm$0.02. This value was significantly lower than the theoretical ratio assuming no neutrino oscillations (0.37$\pm$0.05).

While the oscillation parameters obtained here are independent of the theoretical flux, a similar stopping-muon analysis at SNO could still help to constrain theoretical uncertainties. The largest effort would be in developing and characterized a new stopping muon fitter. An overall stopping to through-going ratio could be extracted and that ratio could be investigated as a function of zenith angle. Similar to the analysis done here, each zenith angle bin would provide information to an overall oscillation parameter fit.

7.3   Beyond SNO

A number of experiments look to probe both cosmic ray and neutrino physics in the near future. Understanding the origins of the highest energy cosmic rays is one of the main goals

of the Auger Project [13]. With the first results published from their original detector array in Argentina mentioned in the introduction, the group is currently constructing a matching array in southeastern Colorado. This will provide nearly uniform coverage of the entire sky allowing them to look for both point sources of cosmic rays and to investigate large-scale variations or tendencies of cosmic rays in the universe.

With respect to neutrino physics, the current and near future efforts are going into constraining the third mixing angle (θ_{13}) so that next generation neutrino experiments hoping to investigate the CP violation phase in neutrino oscillations can begin. Experiments like T2K [27] and Double Chooz [28] aim to measure θ_{13} and help complete the understanding of neutrino oscillations. If θ_{13} is found to be roughly equal to or greater than 0.01, intensified conventional neutrino beams should be able to probe CP violation. The CP violating phase in neutrino oscillations allows for a difference in oscillation probability between neutrinos and antineutrinos. A measurement of this value could help to resolve the matter-antimatter asymmetry in the Universe. If θ_{13} is found to be significantly less than 0.01, experiments may need to rely on "neutrino factories" [31]. These factories will be designed to produce a very pure and intense source of neutrinos.

Bibliography

[1] M. W. Friedlander. *A Thin Cosmic Rain*. Harvard University Press, Cambridge, 2000.

[2] R. A. Millikan and I. S. Bowen. *Phys. Rev.*, 27:353–361, 1926.

[3] A. H. Compton. *Phys. Rev.*, 50:1119–1130, 1936.

[4] H. E. Tatel and J. A. Van Allen. *Phys. Rev.*, 73:87–88, 1948.

[5] S.H. Neddermeyer and C.D. Anderson. *Phys. Rev.*, 51:884–886, 1937.

[6] H. J. Bhabha and W. Heitler. *Proc. Roy. Soc.*, 159:432, 1937.

[7] P. Auger and T. Grivet. *Rev. Mod. Phys.*, 11:232–234, 1939.

[8] G. P. S. Occhialini C. M. G. Lattes, H. Muirhead and C. F. Powell. *Nature*, 159:694–697, 1947.

[9] T. K. Gaisser J. Cronin and S. P. Swordy. *Sci. Amer.*, 276:44, 1997.

[10] T. Stanev. *High Energy Cosmic Rays*. Springer-Praxis, Chichester, UK, 2004.

[11] E. S. Seo *et al*. *Ap. J.*, 378:763, 1991.

[12] D. J. Bird *et al*. *Ap. J.*, 424:491, 1994.

[13] J. Abraham *et al*. (Pierre Auger Collaboration). *Nature*, 318:938–943, 2007.

[14] T. K. Gaisser. *Cosmic Rays and Particle Physics*. Cambridge University Press, Cambridge, 1990.

[15] M. Aglietta *et al.* (LVD Collaboration). *Phys. Rev.*, D58:092005, 1998.

[16] K. S. Hirata *et al.* *Phys. Lett.*, B205:416, 1988.

[17] D. Casper *et al.* *Phys. Rev. Lett.*, 66:2561, 1991.

[18] R. Davis Jr. *Phys. Rev. Lett.*, 12:303, 1964.

[19] Q. R. Ahmad *et al.* (SNO Collaboration). *Phys. Rev. Lett.*, 89:011302, 2002.

[20] J. N. Bahcall. *Neutrino Astrophysics*. Cambridge University Press, Cambridge, 1989.

[21] N. Nakagawa Z. Maki and S. Sakata. *Prog. Theor. Phys.*, 28:870, 1962.

[22] V. N. Gribov and B. M. Pontecorvo. *Phys. Lett.*, B28:493–496, 1969.

[23] D. H. Perkins. *Introduction to High Energy Physics*. Cambridge University Press, Cambridge, 2000.

[24] B. Aharmim *et al.* (SNO Collaboration). *Phys. Rev. Lett.*, 101:111301, 2008.

[25] P. Adamson *et al.* (MINOS). *Phys. Rev.*, D77:072005, 2008.

[26] M. Apollonio *et al.* (Chooz Collaboration). *Eur. Phys. J.*, C27:331–374, 2003.

[27] Y.Itow *et al.* The jhf-kamioka neutrino project (`arXiv:hep-ex/0106019`). 2001.

[28] F. Ardellier *et al.* (Double Chooz). Double chooz, a search for the neutrino mixing angle θ_{13}(`arXiv:hep-ex/0606025v4`). 2006.

[29] V. Fanti *et al.* *Phys. Lett.*, B465:335–348, 1999.

[30] B. Aubert *et al.* (BABAR). *Phys. Rev. Lett.*, 93:131801, 2004.

[31] C. Amsler *et al.* *Phys. Lett.*, B667:1, 2008.

[32] A. Osipowicz *et al.* (KATRIN). Katrin: A next generation tritium beta decay experiment with sub-ev sensitivity for the electron neutrino mass(`arXiv:hep-ex/0109033`). 2001.

[33] N. Tagg. *The ^8Li calibration Source and Through-going Muon Analysis in the Sudbury Neutrino Observatory.* PhD thesis, The University of Guelph, 2001.

[34] The SNO Collaboration. *Nuclear Instruments and Methods in Physics Research*, A449:171–207, 2000.

[35] M. C. Chen. Sno and sno+. In *AIP Conf. Proc.*, volume 944, pages 25–30, 2006.

[36] R. W. Le Maitre (Editor). *Igneous Rocks: A Classification and Glossary of Terms, Recommendations of the International Union of Geological Sciences, Subcommission of the Systematics of Igneous Rocks.* Cambridge University Press, Cambridge, 2002.

[37] C. Stormer. *Astrophysics*, 1:237, 1930.

[38] C. Tunnel and S. Seibert. Measuring the trigger efficiency in the ncd-phase and the all-phase implications. SNO technical note, 2006.

[39] J. R. N. Cameron. *The Photomultiplier Tube Calibration of the Sudbury Neutrino Observatory.* PhD thesis, Wolfson College, Oxford University, 2001.

[40] MCNP4A. A monte carlo n-particle transport code system. Radiation Shielding Information Center, Los Alamos National Laboratory, 1993.

[41] D. Casper. *Nucl. Phys. Proc. Suppl.*, 112:161, 2002.

[42] G. D. Barr, T. K. Gaisser, P. Lipari, S. Robbins, and T. Stanev. *Phys. Rev.*, D70:023006, 2004.

[43] R. A. Smith and E. J. Moniz. *Nucl. Phys.*, B43:605, 1972.

[44] A. Bodek and U. K. Yang. *Nucl. Phys. Proc. Suppl.*, B112:70, 2002.

[45] D. Rein and L. Sehgal. *Ann. Phys.*, 133:79, 1981.

[46] P. Lipari and T. Stanev. *Phys. Rev.*, D44:3543–3554, 1991.

[47] P. Antonioli *et al. Astroparticle Physics*, 7:357–368, 1997.

[48] C. C. M. Kyba. *Measurement of the Atmospheric Neutrino Induced Muon Flux at SNO.* PhD thesis, University of Pennsylvania, 2006.

[49] C. Kyba. An improved snoman pmt transit time distribution with realistic pre-pulse. SNO technical note, 2006.

[50] Tom Walker. *Muon Flux at SNO (In progress).* PhD thesis, Massachusetts Institute of Technology, 2009.

[51] J.R. Wilson. *A Measurement of the ^8B Solar Neutrino Energy Spectrum at the Sudbury Neutrino Observatory.* PhD thesis, Jesus College, Oxford University, 2004.

[52] C. Currat. Personal communication. 2006.

[53] Y. Ashie *et al.* (Super-Kamiokande). *Phys. Rev.*, D71:112005, 2005.

[54] D. G. Michael *et al.* (MINOS). *Phys. Rev. Lett*, 97:191801, 2006.

[55] R. Barlow. Asymmetric statistical errors (`arXiv:physics/0406120v1`). 2004.

[56] J. A. Formaggio. Applying systematic pulls to a maximum likelihood analysis. SNO analysis note, 2007.

[57] P. S. Auchincloss *et al. Z. Phys.*, C48:411, 1990.

[58] M. Ambrosio *et al.* (MACRO Collaboration). *Phys. Rev.*, D52:3793, 1995.

[59] M. W. E. Smith. *An Investigation of Matter Enhanced Neutrino Oscillation with the Sudbury Neutrino Observatory.* PhD thesis, University of Washington, 2002.

[60] Y. Fukuda *et al.* (Super-Kamiokande). *Phys. Lett.*, B467:185–193, 1999.

[61] D. Ballard. *Computer Vision.* Prentice-Hall, New Jersey, 1982.

[62] I. Pitas. *Digital Image Processing Algorithms.* Prentice-Hall, New York, 1993.

[63] M. Desbrun *et al.* Anisotropic feature-preserving denoising of height fields and images. In *Graphics Interface 2000 Conference Proceedings*, pages 145–152, 2000.

[64] M. Desbrun *et al.* Implicit fairing of arbitrary meshes using diffusion and curvature flow. In *SIGGRAPH 1999 Conference Proceedings*, pages 317–324, 1999.

Appendix A

Individual Effort

A group as large and diverse as the SNO collaboration provided an incredible pool of effort, knowledge and expertise. The entire muon analysis at SNO was completed largely by the muon working group led by Joe Formaggio at MIT and included Chris Kyba at the University of Pennsylvania, Charles Currat at Lawrence Berkeley National Lab, Tom Walker at MIT and myself.

The general division of work had Joe Formaggio compiling the final paper and developing the systematic error analysis, Chris Kyba producing the muon event fitter, Tom Walker focusing on the external muon detection system and me developing the instrumental background study and the final signal extraction code for the neutrino-induced muon analysis. This code extracted the neutrino oscillation parameters, overall normalization of the neutrino flux and associated uncertainties for the atmospheric neutrino analysis at SNO. These physics results will be those presented in the SNO through-going muon paper. The analysis described in this thesis was my responsibility. The only exceptions are the description of the muon event fitter and the development of the systematic uncertainties corrections. Both are included here because they were such a significant part of the entire analysis.

A lot of the preliminary work I was involved with included testing and developing simulation and analysis programs. Before MUSIC and `nuance`, both external particle MC simulators, were integrated into the standard SNO analysis software, I helped to test and characterize the code. I was also tasked with maintaining and updating the previous muon

event fitter developed by Nathaniel Tagg [33], as it was used to compare and check newer results from the muon fitter described in here.

In addition to the work detailed in this thesis, I was given a number of other responsibilities within the collaboration. Spending time at the SNO site while the detector was still operational was important in helping with the day-to-day operation of the detector and in learning the processes and tasks that occur there. In addition to the 4 months I spent at SNO doing maintenance and shift work, there were 3 or 4 trips a year to site to fulfill the operator shift requirement at UBC. This resulted in over 60 detector operator shifts over the course of my graduate work. Detector operator shifts involved ensuring the detector was running optimally, performing calibrations and checks. This helped a tremendous amount with understanding the functioning of the detector.

Graduate students were also given the opportunity to help with some of the lower level analysis tasks. I was given the responsibility of maintaining the database of technical reports at SNO. These reports were not directly related to analysis, but for example described technical calibrations and the construction of SNO.

Appendix B

Hough Transform Event Fitter

The current muon fitter was designed to fit for the track of a through-going muon fitter. The performance of that fitter on muons that stop in the detector was not well characterized and was a function of the amount of track left in the detector. The probability distribution functions of the current fitter were developed and tuned by through-going muon events and all the information available in the event (hit pattern, tube timing and tube charge). Work put into developing a muon fitter that would fit for the track of stopping muons and through-going muons similarly well is described here. While the hit patterns of stopping muons and through-going muons differ, the light coming from stopping muons and through-going muons share a common source; it is produced by Cherenkov light emitted in a cone. Using the geometrical fact that the Cherenkov cone intersects the spherical detector, a Hough transform based fitter was developed.

The fully functional fitter was not completed, but described here are the step taken in its development. Presented first are the principles behind Hough transforms. Next, the anisotropic smoothing method which is used to reduce noise in the PMT charge distribution is explained. Finally, the application of the Hough transforms to full SNO muon events is described.

B.1 Hough Transforms

Hough transforms are used for detecting parameterized curves (e.g. circles) in digitized images [61][62]. The transform takes the problem from straightforward search for the parameterized curve in image space to a search for the curve in parameter space. In the case of SNO, the roughly 9600 PMTs act as sampling points or pixels. The charge registered by a PMT can be compared with the amount of light a pixel receives in a digital image. The intent here was to extend this idea to search for the cone vertex and direction produced by the muon event. The sampling information available is restricted to the surface of the sphere that defines the edge of the detector.

A simple example helped to demonstrate the Hough transform method and characterize its performance. Simulated images were produced 100×100 pixels in size with each pixel either on or off. Circles and noise were artificially placed into the simulated image and a Hough transform was used to search for the circles.

A circle can be parameterized as

$$(x - a)^2 + (y - b)^2 = r^2 \tag{B.1}$$

where the parameters of interest are the circle centre (a, b) and the circle radius, r. Given that each parameter in the description of circle has a certain range and resolution defined by the image boundaries and the user, a hit pixel is an element of a finite number of circles. Each of these circles are represented by a point in parameter space (a, b, r) and will contribute a vote to an accumulator array, $P(a, b, r)$. A second hit pixel will have another finite set of circles which it satisfies, but only a subset of circles will be contained in both sets. These circles show up as points which overlap in parameter space. This continues for all hit pixels and the best fitting circle for the image will be the point in parameter space

146

that has the greatest number of votes. For each point in the image, the method iterates over all possible (a, b) combinations so that the number of operations goes as (number of hit pixels $\times N_a \times N_b$), where N_a and N_b are the number of steps in a and b. With a (x, y) pair and a (a, b) pair, from equation B.1, the value of r is determined so no iteration over r is necessary.

The simulated images consisted of a list of (x, y) pairs that defined the hit pixels. The images were generated by specifying the circle centre and radius as well as a background noise level and scatter in the circle itself. This list of (x, y) pairs were passed to the Hough transform which perform the operations outlined above. The resolution in parameter space is 1 unit in x, y and r. Figures B.1-B.5 show the image in the upper left and three projections of parameter space in the other areas. Since there are 3 parameters, $P(a, b, r)$ is shown by fixing one of the parameters and plotting a 2D histogram of the other two, i.e. $P(a, b, r_{max})$, $P(a_{max}, b, r)$ and $P(a, b_{max}, r)$. The maximum values correspond to the most voted for point in parameter space found by the Hough transform. These plots are intended to show the noise surrounding the peaks and not to check whether the Hough transform found the correct circle. In all but one of the examples shown, the Hough transform was successful to the accuracy of the accumulator array in finding the correct circle.

Figure B.1 is an image with the circle $(a = 26, b = 73, r = 18)$, consisting of 100 points with roughly 500 extra points placed randomly in the image. The peaks in parameter space stand clearly above any noise by roughly 6 times. In the bottom two plots, the noise is greater for lower values of r. This is most likely due to the resolution in parameter space being 1 unit. With rounding errors it is easy for a number of hit pixels to belong to a small group of small circles ($r < 3$). This effect will be the main cause of the transform finding the wrong circle in a later example. Figure B.2 shows the same circle with the same number of points, but with a scatter introduced to the ring. The noise level in parameter space
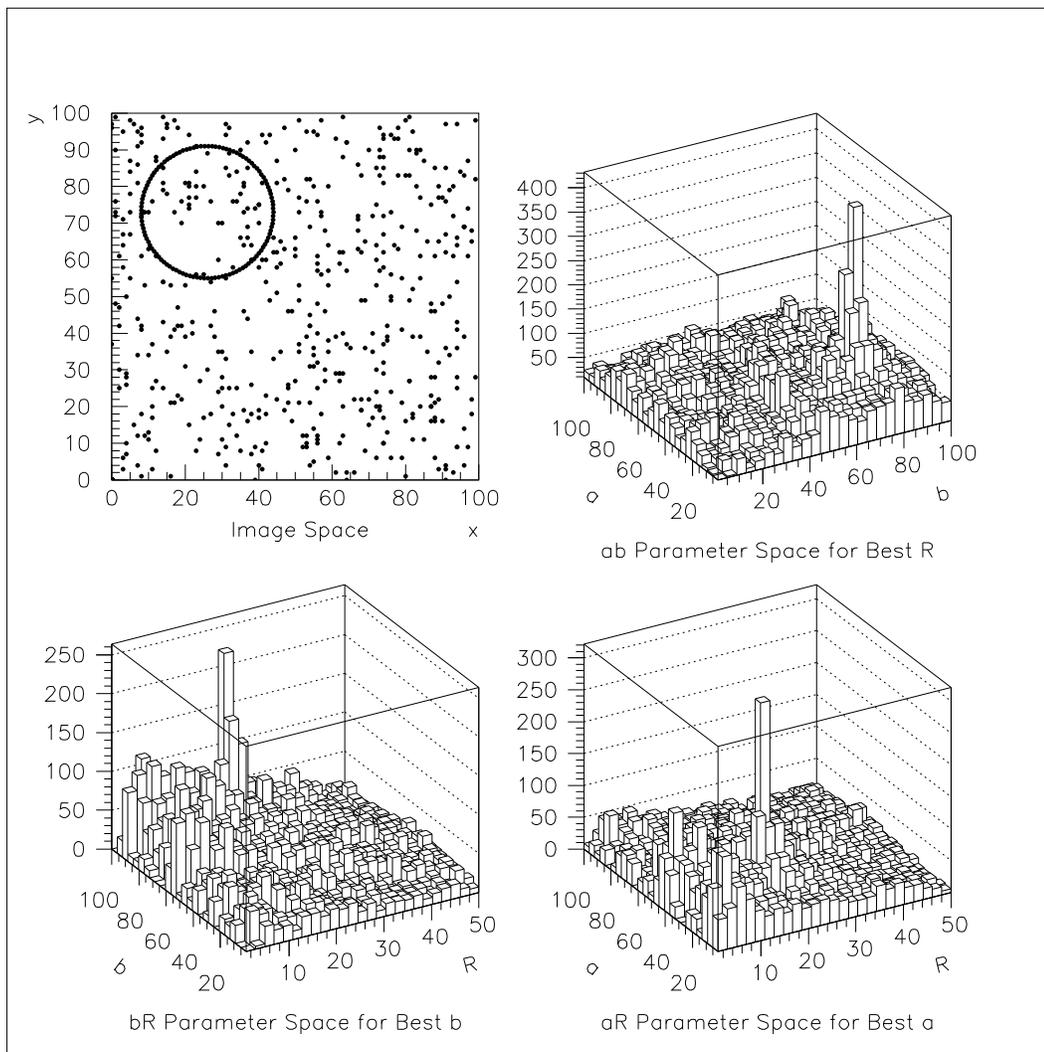
147

Figure B.1: A circle with background noise in a simulated image and parameter space distributions. The top left shows a circle with 100 points and 500 background noise points. See text for a description of the parameter spaces. The peak that corresponds to the best fit circle is at roughly the 6σ level.
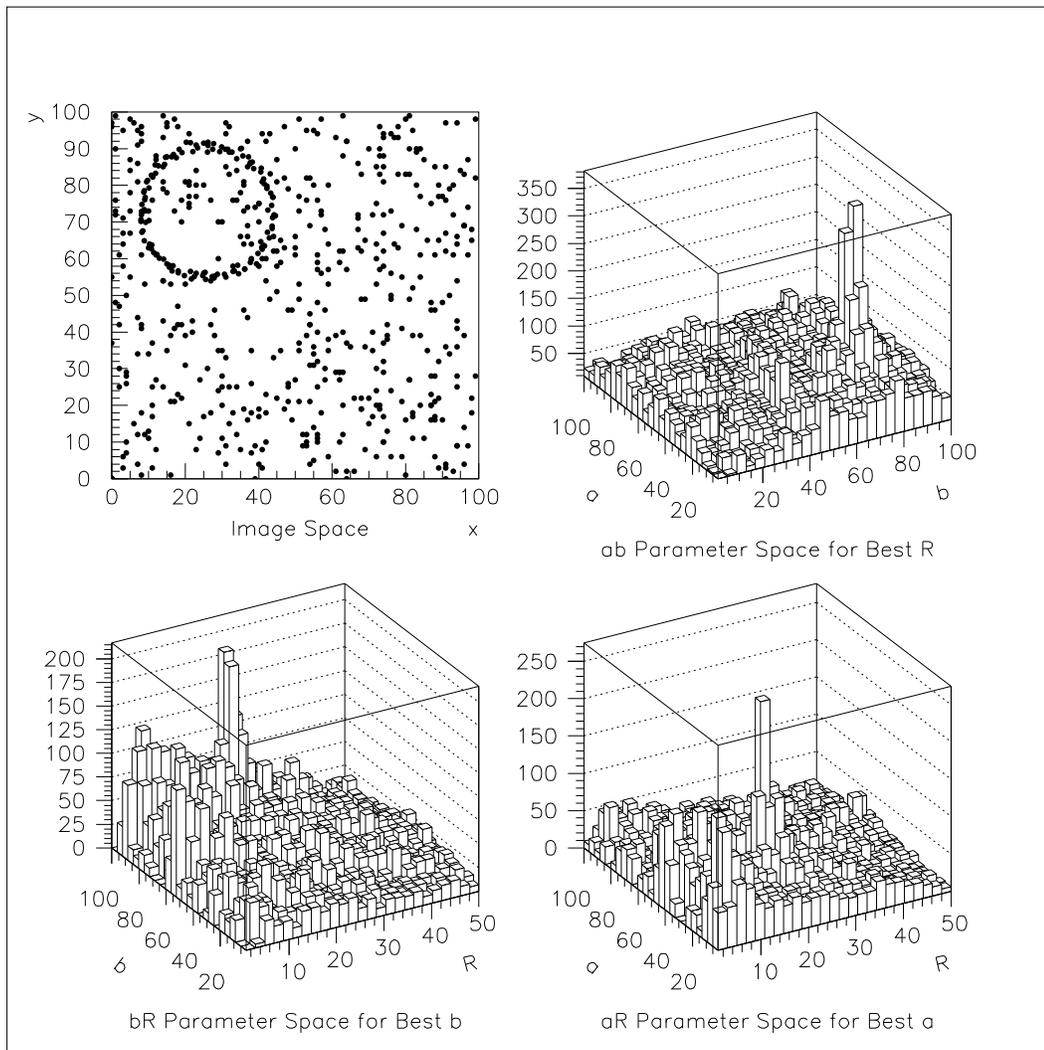
Figure B.2: A ragged circle with background noise in a simulated image and parameter space distributions. The top left shows a scattered circle with 100 points and 500 background noise points. The peak that corresponds to the best fit circle is at roughly the 5σ level.
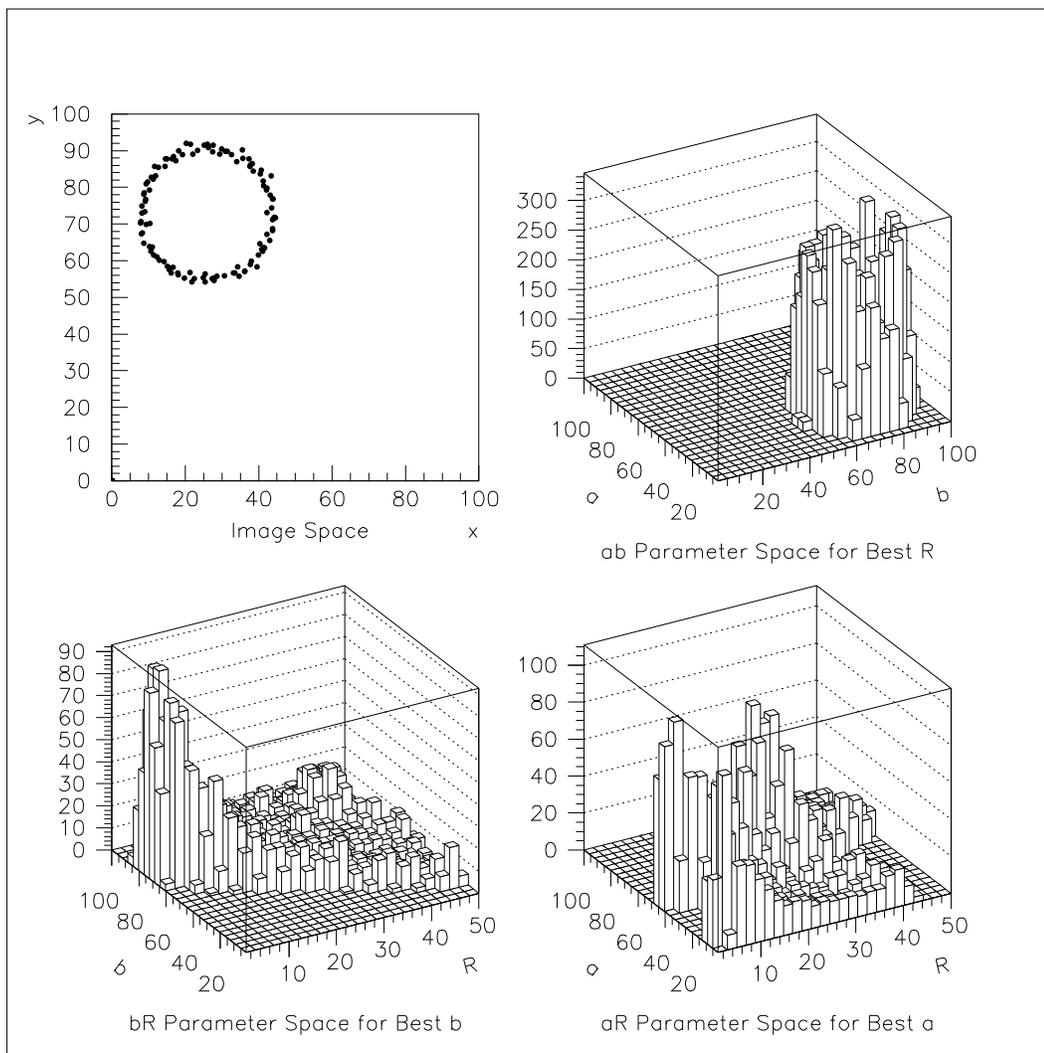
Figure B.3: A ragged circle with no background noise in a simulated image and parameter space distributions. The top left shows a more scattered circle with 100 points and 500 background noise points. The Hough transform failed in this case because scattering the point in the circle took votes away from the peak as well as adding noise to the surrounding parameter space.
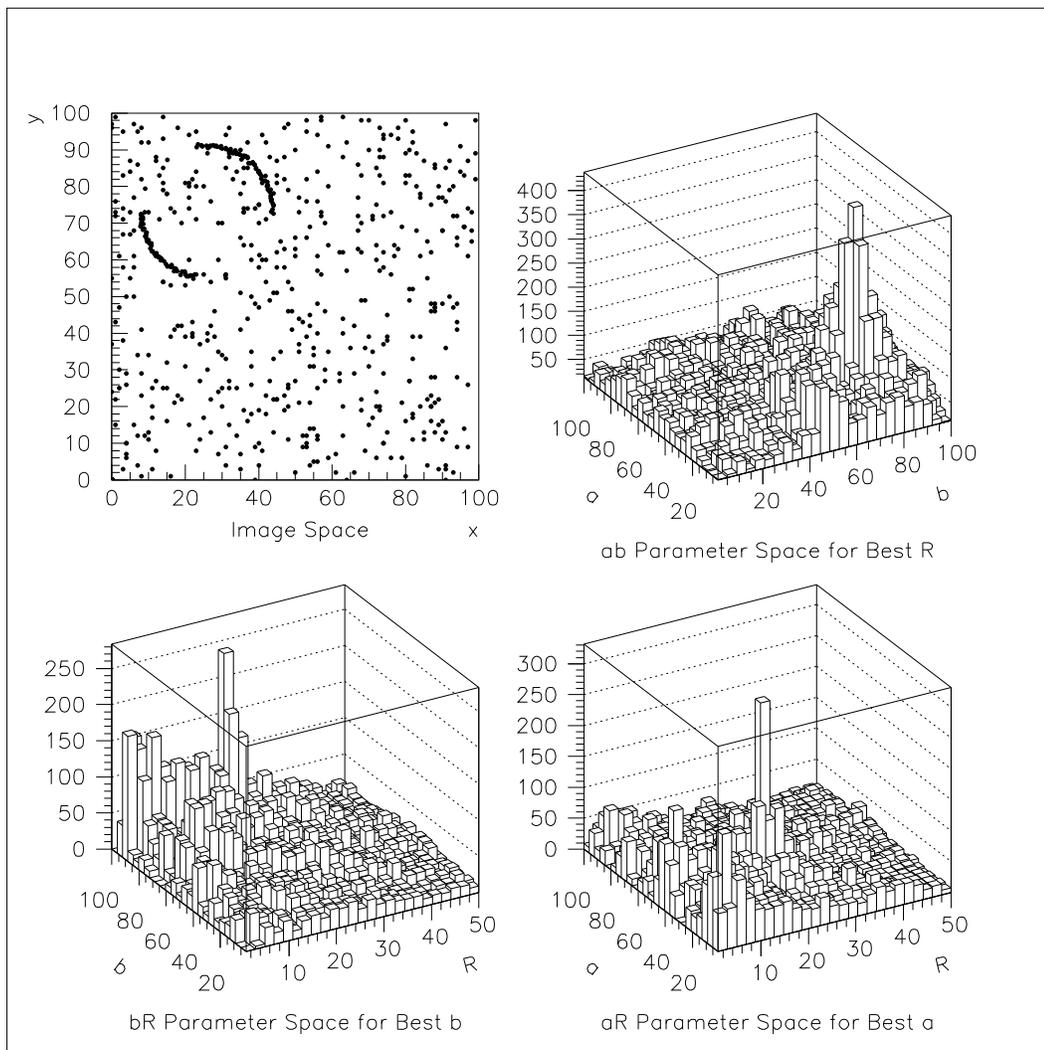
Figure B.4: A segmented circle with background noise in a simulated image and parameter space distributions. The top left shows a scattered and segmented circle with 100 points and 500 background noise points. The peak that corresponds to the best fit circle is at roughly the 5σ level.
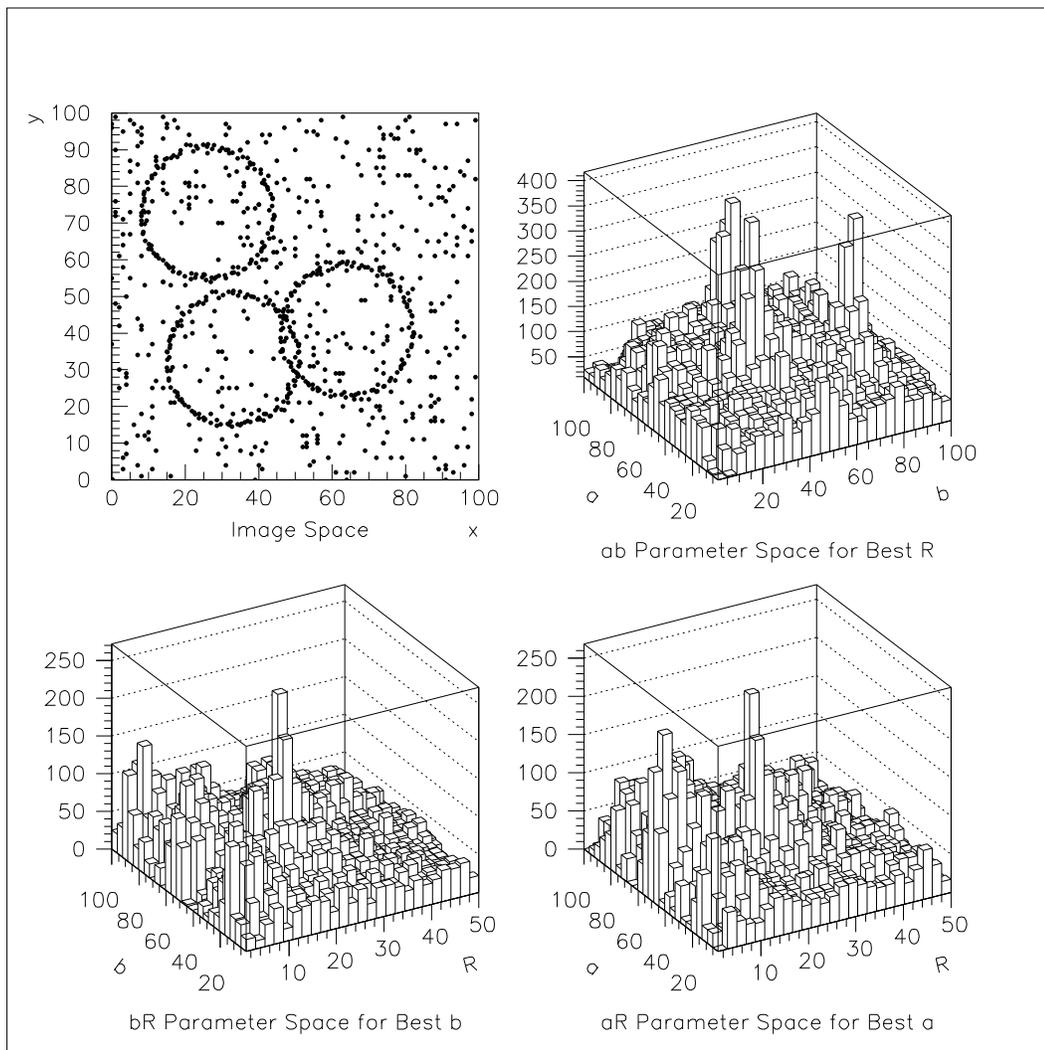
Figure B.5: Multiple ragged circle with background noise in a simulated image and parameter space distributions. The top left shows 3 scattered circles with 100 points each and 500 background noise points. Since the 3 circles all have the same radius, 3 peaks show up in the top right projection of parameter space. The best R histogram shows all 2 circles because all of the circles have the same r. The bottom two do no show 3 peaks because the 3 circles do not share the same centre coordinates.

is about the same but the peak value has dropped by roughly 30%. Figure B.3 shows the same circle with the a slightly larger spread in the circle. It has chosen an incorrect circle with radius 2 as the best fit. The plots in br and ar space represent projections along the wrong a and b, so this shows that the peaks that dropped by 30% before have dropped below the noise at low r. Figure B.4 shows the same circle with the same number of points, but with the points confined to 2 arcs creating a segmented circle. There is a spread in the points as well. This shows that the Hough transform is not seriously affected by gaps in the circle. Figure B.5 shows how the Hough transform can find multiple rings. In order to visualize the peaks in parameter space, all the circles were chosen to have the same radius so that the maximum r value will be the same for all circles and therefore will show up in the projections as three distinct peaks. The bottom 2 plots do not provide much information since only one of the circles has the peak a and b.

To qualitatively summarize, the Hough transform is relatively unaffected by gaps in the hit pattern or background noise. Noise in the actual circle has the largest effect on the ability of the Hough transform to find circles because it both removes votes from the peak and adds to the surrounding noise level in parameter space. Essentially, the signal is reduced and the noise is increased in one step. Multiple circles will show up as multiple peaks in parameter space and a simple thresholding scheme or more involved peak-finding routine should be able to pick the circles out. These properties of the Hough transform make it suited for searching for muon Cherenkov cones since the PMT hit pattern of a muon suffers from gaps and a large amount of noise. Only a few hundred PMTs should be hit by direct light from a muon Cerenkov cone, while in any muon event almost all PMTs can be hit by indirect light. The main difficultly, as described in the next section, is extracting set of PMTs which best defines the Cerenkov cone.

B.2 Applying the Hough Transform to SNO

Extending the Hough transform to search for cones based on hit PMTs requires two main steps. The first is to select the PMTs which best represent the cone and the second is to chose a parameterization for the cone.

B.2.1 Tube Selection

An average muon event produces light in the detector which fires thousands of PMTs. Some of the light is direct light from Cherenkov cone and some is reflected light. Using the PMTs as pixels in an image, and the measured charge as the counts registered in each pixel, only the PMTs hit by direct light are useful in reconstructing the Cherenkov cone. In addition, since the muon produces Cherenkov light along its entire path, only the PMTs hit by the first Cherenkov light as the muon just enters the detector will be useful in determining that cone vertex.

The simplest way to reduce the number of PMTs is to apply a charge cut to the PMTs. Those hit by indirect light will have a lower charge than PMTs hit by direct light from the Cherenkov cone. As shown in Figure B.6 the scatter in the charge distribution makes cutting on charge difficult. This figure shows a simulated muon traveling from the left to the right. Low charge PMTs that were hit were removed from the image, blue represents low charge and red represents high charge. Pink is over-scaled high charge. The PMTs that add to the background noise will not significantly affect the method, but reducing scatter in the main curve of tubes will be important in ensuring the Hough transform will provide accurate results.

Figure B.7 shows the same event with an Gaussian smoothing mask applied. This simple mask reduces the amount of noise both in the background and in the main curve, but still leaves the curve as a thick band. Even after the Gaussian smoothing applied, a
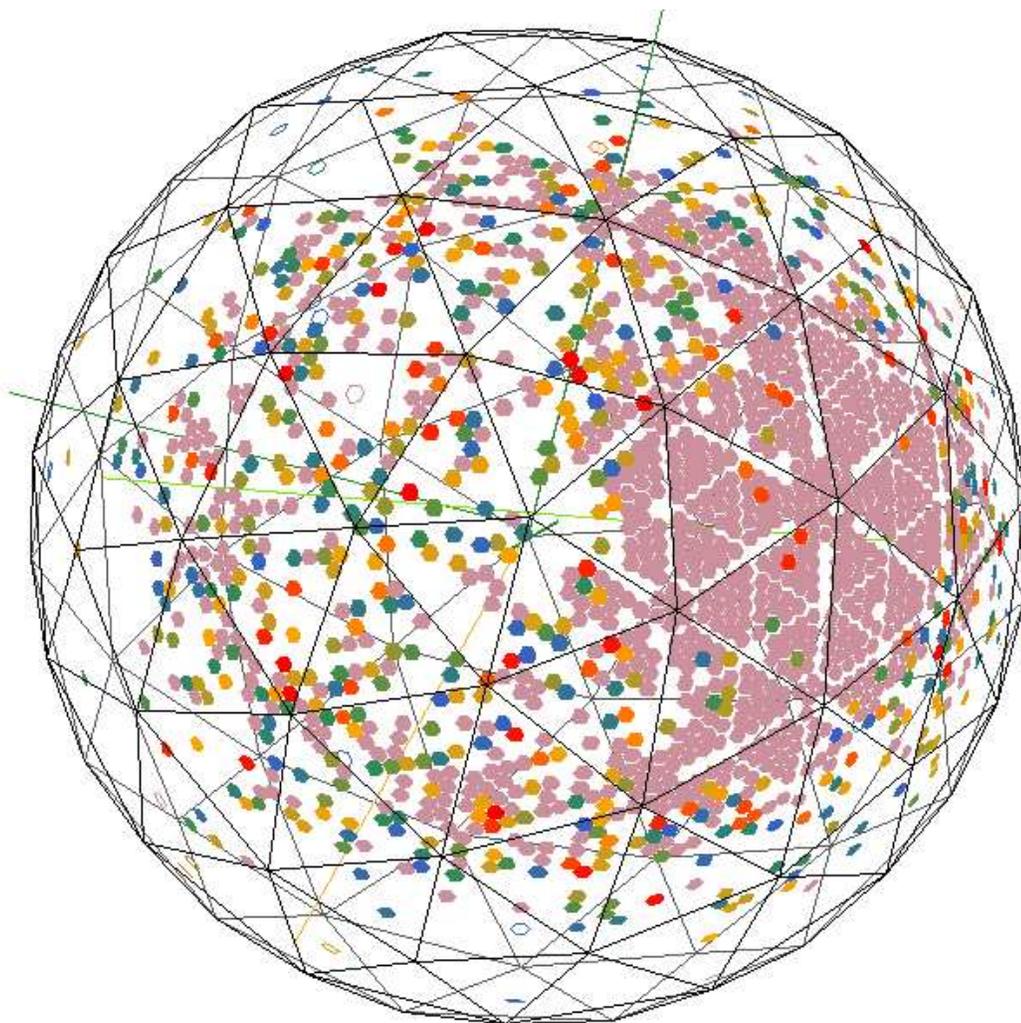
154

PSfrag replacements

$\hat{q}$

$\hat{t}$

Figure B.6: A low impact parameter MC muon event where all the low charge tubes have been cut. The colour scheme shows charge where blue is low and red is high.

simple charge cut cannot be placed to select the edge tubes. Isotropic smoothing techniques reduce noise uniformly and a consequence of this is that sharp edges and features are blurred in the process. This is why a thick band of PMTs are left still after using the Gaussian smoothing mask.

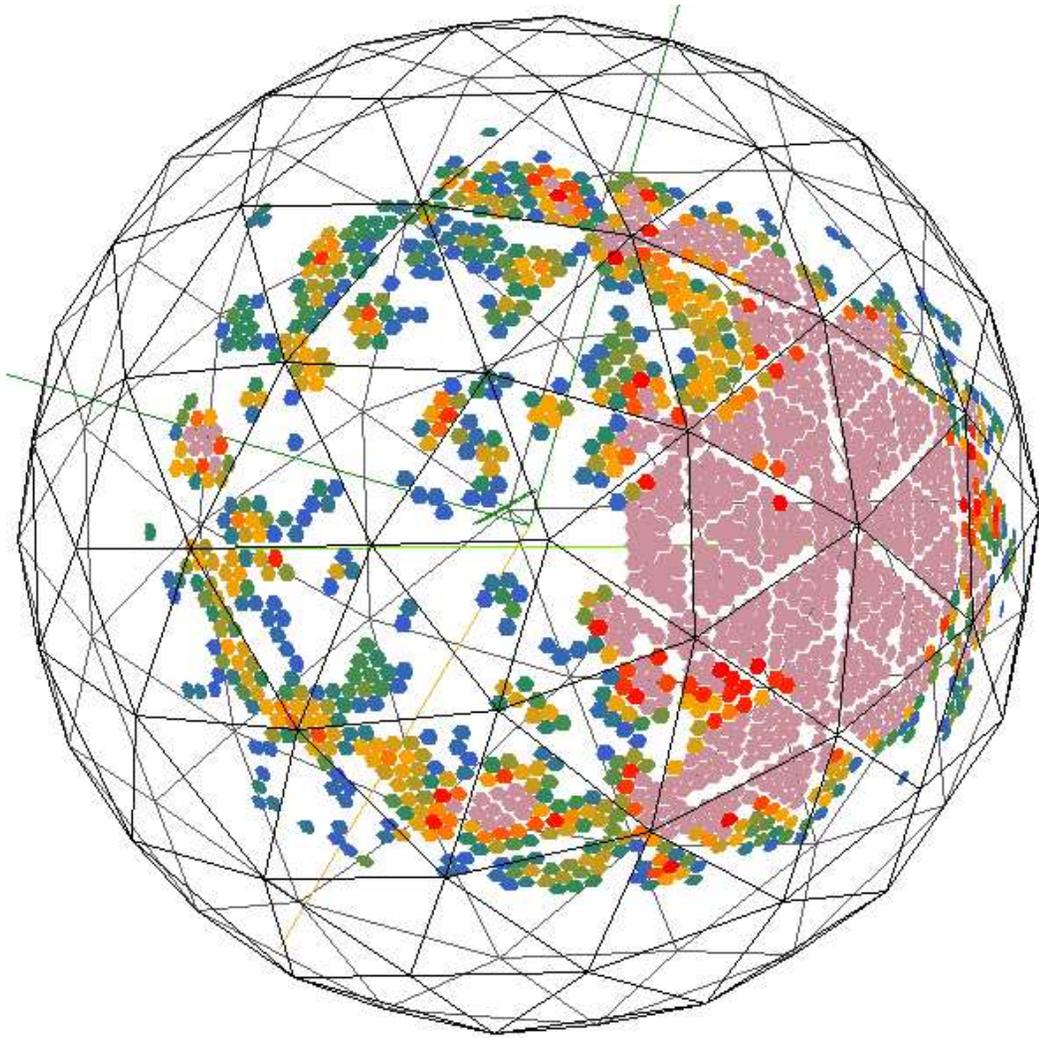PSfrag replacements

$\hat{q}$

$\hat{t}$

Figure B.7: A low impact parameter MC muon event where the PMT charges have been Gaussian smoothed and the low charge tubes have been cut. The colour scheme shows charge where blue is low and red is high.

Anisotropic smoothing methods are more commonly used when sharper features (like edges) wish to preserved as best as possible. The method described in [63] is applicable to data on an arbitrary grid or mesh. In the case of SNO each PMT has usually 6 nearest neighbours. The 7 PMTs can be taken as locally flat and are analogous to a centre pixel in an standard 2D digitized image and its surrounding 8 neighbours. The basic idea is that the smoothing depends on the local mean curvature. Uniformly noisy regions are smoothed the strongest while discontinuities like edges are preserved. The curvature flow technique is used here to smooth surfaces by moving along the surface normal at a speed equal to the mean curvature.

$$\frac{\delta x}{\delta t} = -\bar{\kappa}\mathbf{n} \tag{B.2}$$

In equation B.2, $\bar{\kappa}$ is the mean curvature and $\mathbf{n}$ is the normal to the surface. The mean curvature normal at a point, $\mathbf{P}$ is the limit of the surface area variation with respect to $\mathbf{P}$ as smaller and smaller pieces of surface are taken. This gives a method for calculating the mean curvature normal as in the following equation.

$$2\bar{\kappa}\mathbf{n} = \frac{\nabla \mathcal{A}}{\mathcal{A}} \tag{B.3}$$

$\mathcal{A}$ is a small area around $\mathbf{P}$ and ∇ is the derivative with respect to $\mathbf{P}$. Since information available here are in discrete points, a discretized approximation the surface area gradient is needed and given by

$$\nabla \mathcal{A} = \frac{1}{2} \sum (\cot \alpha_{ij} + \cot \beta_{ij})(\mathbf{X}_i - \mathbf{X}_j). \tag{B.4}$$

The derivation from the continuous case to the above equation is given in [64] and the elements of the equations are shown in figure B.8. With the mesh of sample points, triangles
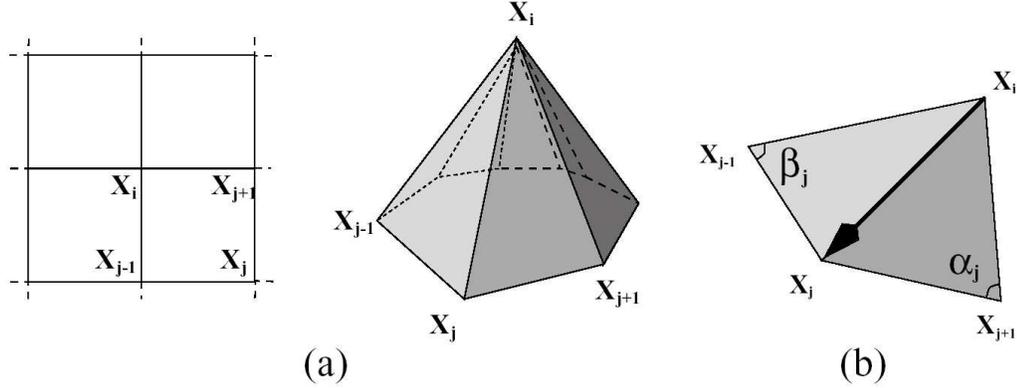
Figure B.8: The left figure shows a 2D grid with points representing pixel positions. The central figure shows a similar grid where the height at a point **X** is the image intensity or in this case PMT charge. The right figure shows the central point and three points connected to it. This forms two triangles with then distances and angles labeled in reference to equation B.4. The figure is from [63].

were formed to calculate the distances, angles and areas shown in equation B.4. For any given flat surface, the above gradient is zero.

The final implementation of the smoothing is given as

$$\frac{dI_i}{dt} = \frac{1}{2\mathcal{A}} \sum (\cot \alpha_{ij} + \cot \beta_{ij})(I_i - I_j). \tag{B.5}$$

Here, I_i is the image intensity at point i or in the case of the SNO PMT data, the charge at PMT_i.

A toy model was used here to help understand and characterize the anisotropic smooth scheme adopted. Figures B.9-B.10 shows how this method is applied to images. The image intensity is translated into a height field and anisotropically smoothed.

The left plot in figure B.9 shows the height field of a image with a thick bar, while the right plot is the same image with a constant level of background noise added. Figure B.9

shows two smoothing methods. The left is a result of an isotropic Gaussian smoothing mask and the right is the result of the anisotropic method described above. While the isotropic method smoothed the noise well, it also removed the edges of the central bar. In the right plot, the noise was reduced while the edges remain sharp.

Figure B.11 shows this anisotropic smoothing method applied to a simulated muon event at SNO. The muon traveled straight down through the centre of the detector. The colour represents the amount of charge a particular PMT detects with blue being a low charge and red representing high charge. The light blue and pink points are under and over-scaled charge respectively.

B.2.2 Cone Parameterization

With the PMTs defining the initial Cherenkov cone selected, a cone parameterization needed to be defined to apply the hough transform. A cone with an axis parallel to the z-axis can be defined as the following

$$\frac{(x - a)^2 + (y - b)^2}{n^2} = (z - c)^2. \tag{B.6}$$

Taking n to be the tangent of Cherenkov cone opening angle in water ($41°$), and accounting also for the cones rotated in both zenith (η) and azimuthal (ψ) angle, this parameterization will have 5 parameters of interest (a, b, c, η, ψ). Given that this cone will have a vertex on a sphere of known radius ($R = 840$ cm), the number of free parameters can be reduced to 4 by representing the cone vertex in spherical coordinates.

$$a = R \sin\theta \cos\phi$$
$$b = R \sin\theta \sin\phi \tag{B.7}$$
$$c = R \cos\theta$$

159

PSfrag replacements

$\hat{q}$

$\hat{t}$

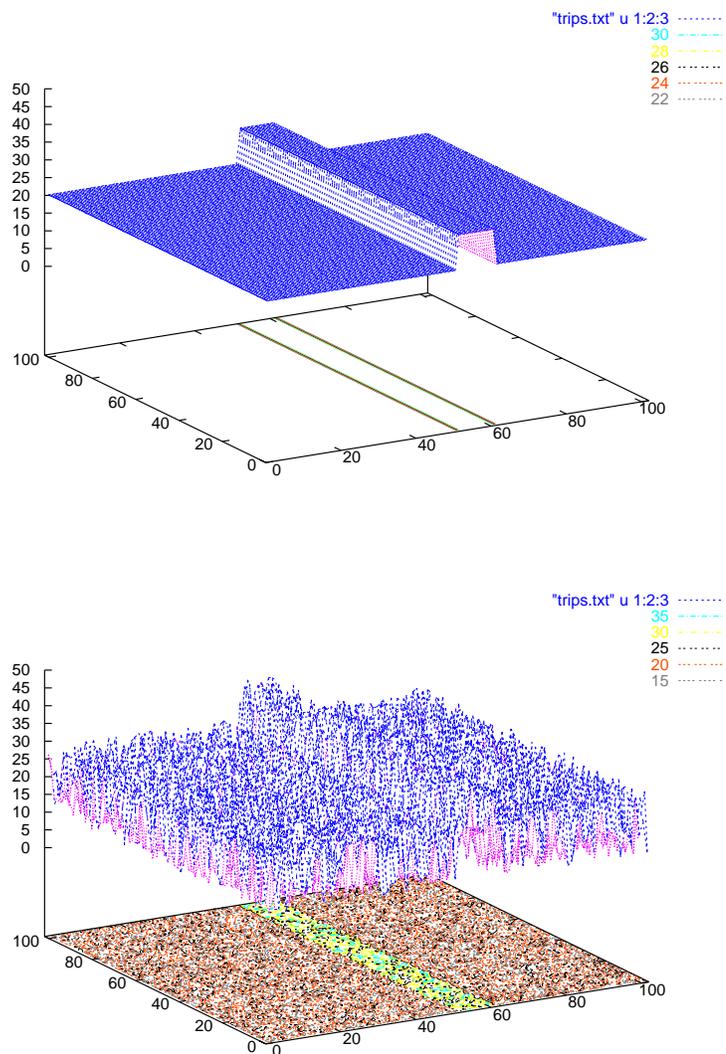PSfrag replacements

$\hat{q}$

$\hat{t}$

Figure B.9: These are height fields of simulated images. The main feature is the bright bar which is shown as a ridge. The top image is simply the ridge, while the bottom image is the ridge with a constant level of background noise added. By eye, it is somewhat hard to distinguish. The vertical axis represents the brightness of each image pixel in arbitrary units.
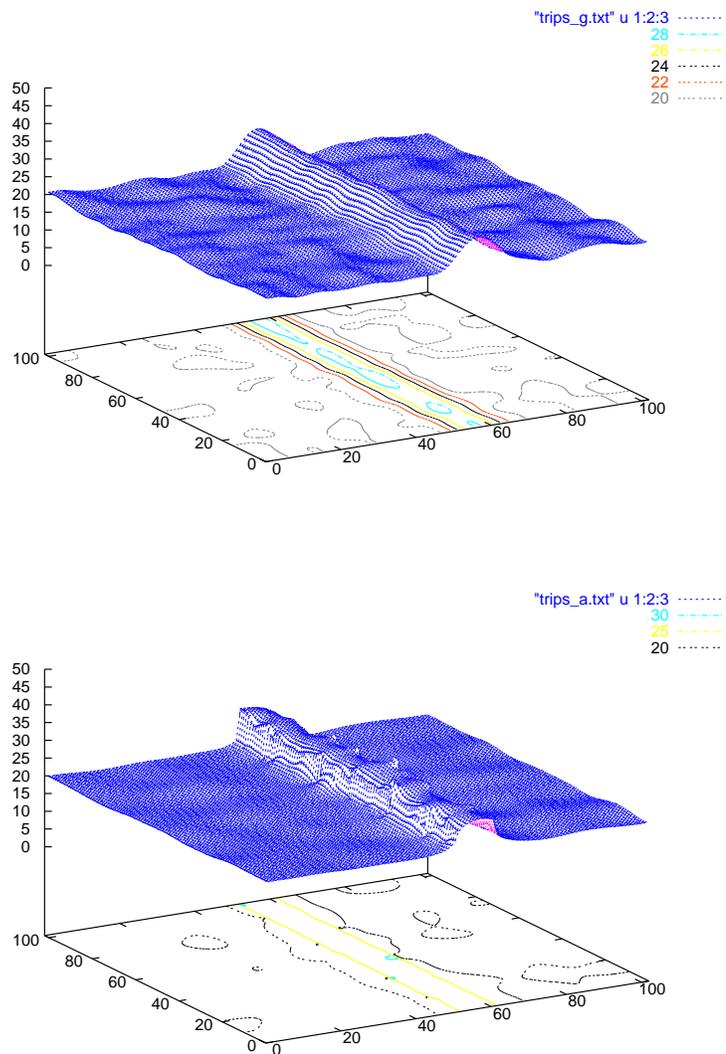
Figure B.10: These images are the previous noisy image smoothed with an isotropic Gaussian method (top) and the anisotropic method (bottom). The overall noise was easily reduced by both methods. The edges are distinct in the top image, while basically non-existent in the bottom. The vertical axis represents the brightness of each image pixel in arbitrary units.
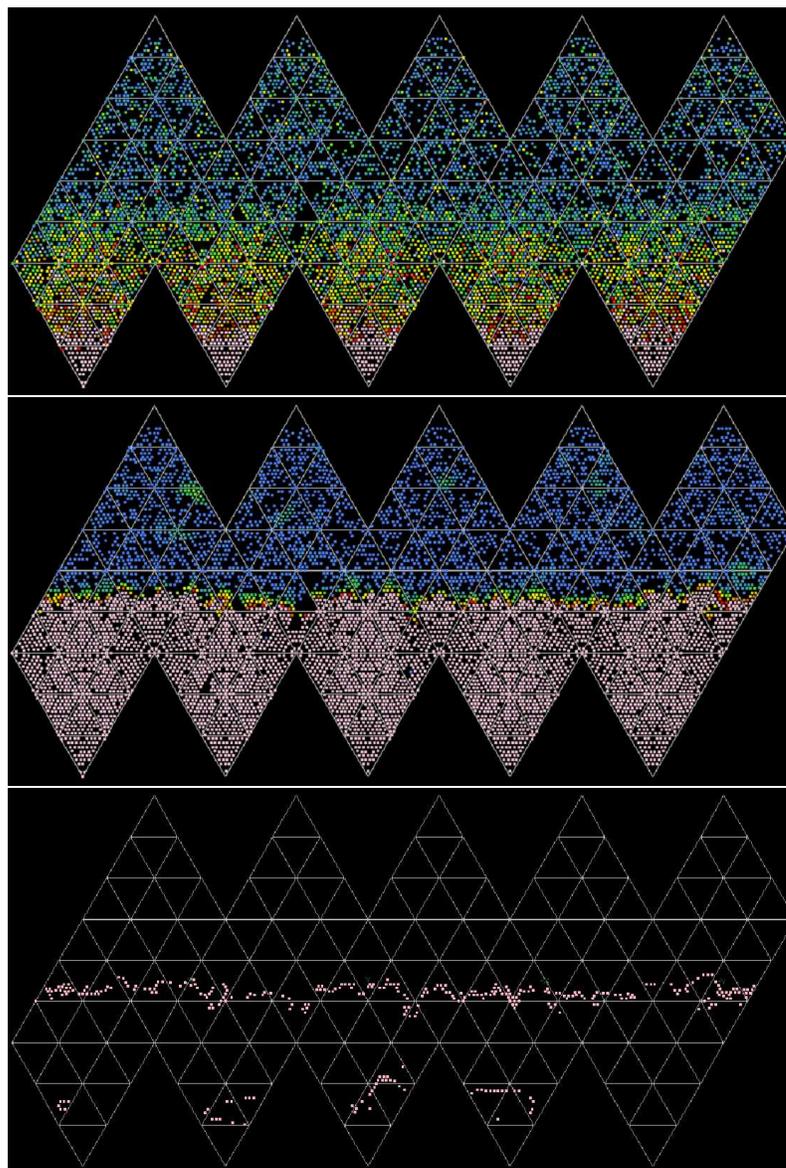
Figure B.11: These are flat charge maps of a simulated muon that travels down the middle of the detector. The colour scheme is the same as all previous charge maps. The top figure shows the raw charge distribution. The second image shows the charge smoothed with the anisotropic algorithm. The bottom image is a simple charge trim showing the PMTs selected by the algorithm to represent the first tubes hit by direct light as the muon enters the detector.

After replacing a, b, and c with the formulae in equation B.7, the accumulator array and loops over parameter space could be established. For each step in the loop, we have a defined value for θ, η and ψ, and the tube position (x, y, z). There were 32 θ and η bins and 63 ψ bins. The 4th parameter, ϕ, is calculated by inverting the above equations using *Mathematica*. There were 4 resulting solutions for ϕ, each of which have fairly complex forms. The four results were checked by reinserting the values of ϕ back into the equation of the cone. The value which resulted in the opening angle closest to $41°$ was selected.

A number of simulated cones were generated with the actual opening angle and PMT positions. PMTs were considered hit by the cone if they were at the intersection of the cone and sphere. Figure B.12 shows the results from the hough transform using the above parameterization of a cone. The residual impact parameters are plotted against impact parameter squared. The impact parameter was found from the cone vertex and direction obtained from the hough transform. The hough transform performs well at higher impact parameters where the PMT hit patterns are more distinct. Lower impact parameter events, given the resolution in parameter space step size were harder to fit. The pattern is mostly likely a result of the discrete resolution steps used in the rotations.

B.3   Simulated Trials

The next steps were to apply the smoothing algorithm and hough transform to actual MC simulated muons and extract the cone direction and vertex from the accumulator arrays. The difficulty in the last step turned out to limit the full development of the entire algorithm.

Figures B.13 and B.14 show two examples of the resulting PMTs after the smoothing algorithm is applied. The red points are the PMTs selected from the smoothing algorithm and the green points are those from a perfect cone with the same vertex and direction as the muon intersecting the PMTs. In both cases, the smoothing algorithm is performing
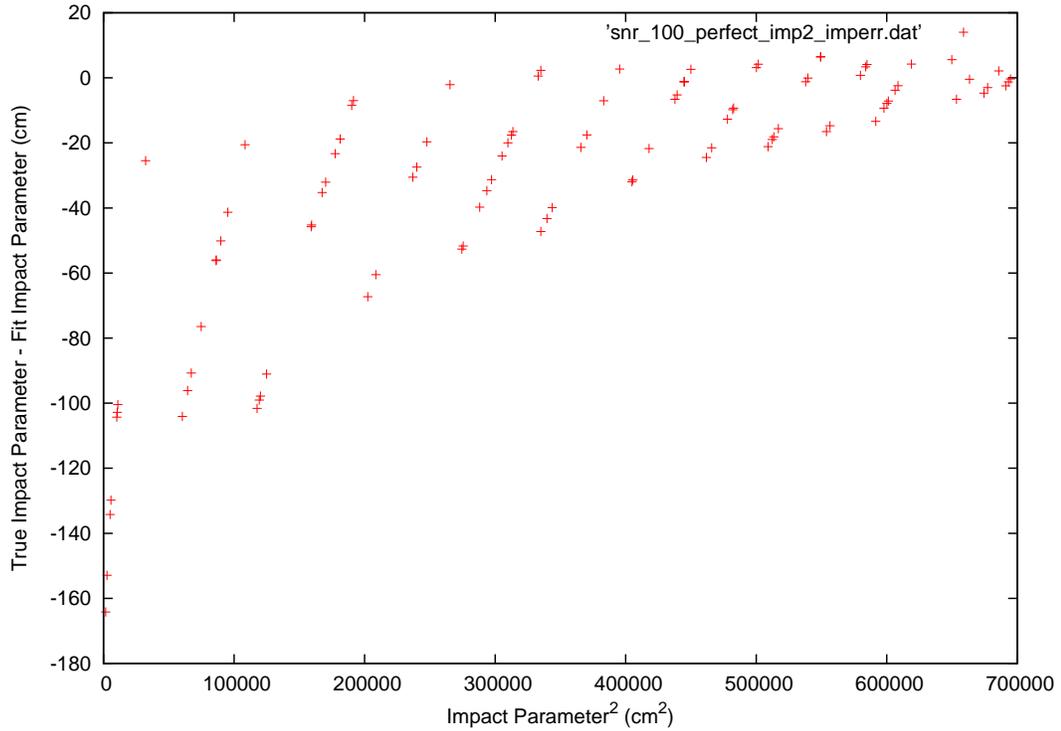
Figure B.12: This shows the results from the hough transform on simulated perfect cone-sphere intersections. The residual impact parameters are plotted against impact parameter squared. The hough transform performs well at higher impact parameters where the PMT hit patterns are more distinct. The pattern is mostly likely a result of the discrete resolution steps used in the rotations.

well, but in figure B.14 the effects of a high impact parameter muon are seen. The PMT light concentrators block light coming in at large incident angles. High impact parameter muons produce the double-edged pattern seen because of this incident angle cut off and the smoothing algorithm retains it. A complete solution to this problem has not yet been developed.

The final accumulator array has all the information from the event, but extracting specific elements was difficult. A method for characterizing the performance of the transform needs to be developed. Events that fit well only require a search for the peak in
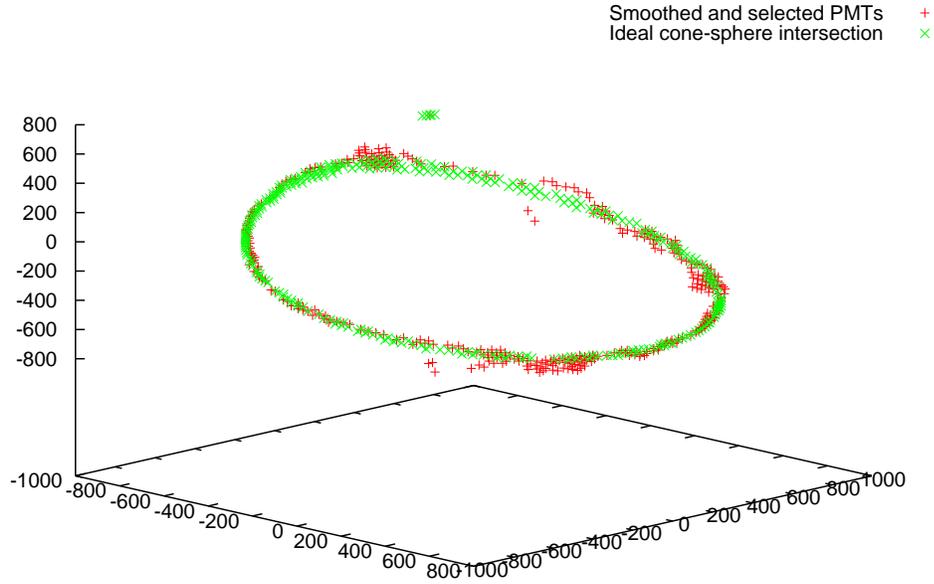
Figure B.13: Plotted here are the PMT positions of a perfect cone-sphere intersection in green and the result of smoothing a MC simulated muon event at SNO. This is a lower impact parameter event and two simulations match well. The axes are x, y, z PMT positions.

5-dimensional parameter space. Events that fail to recover the input muon vertex are harder to investigate and visualize in parameter space.

B.4   Summary and Improvements

Algorithms to search for the muon vertex and direction were developed. The method of anisotropically smoothing and selecting PMTs to pass to the hough transform performed well both in simulated flat images an in preserving the edge of PMTs in muon events. The hit pattern properties resulting from PMT light concentrators prevented the method
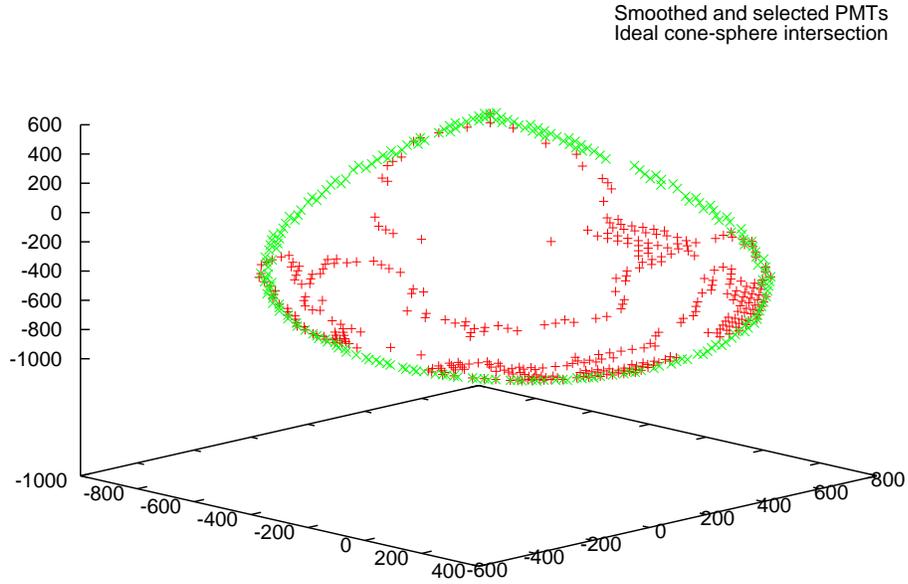
Figure B.14: Plotted here are the PMT positions of a perfect cone-sphere intersection in green and the result of smoothing a MC simulated muon event at SNO. This is a higher impact parameter event and two simulations match quite poorly. The edge finding routine is working as intended, but the PMT concentrators cause the inner edge. The axes are x, y, z PMT positions.

from selecting the proper edge in higher impact parameter muon events. Some method of removing the tubes along the inner edge needs to be developed to address this.

Extracting the most voted for position and direction from the accumulator array only corresponded to the actual position direction for roughly half the events that were simulated. The peak or signal for the other events were either lost in noise or had fewer votes than random noise spikes similar to what was shown in the first example with simulated circles. A method of obtaining the information from the huge array is what is needed at this point.

A good start would be a peak finding algorithm in 5 dimensions and something to quantify the amount of "power" in a peak or the amount of neighboring points in parameter space that also received a large number of votes. This would help to remove single spikes and noise.

Another element of the algorithm that could be improved upon was mentioned at the start. The resolution steps in parameter space are limited by processing power and storage space. The parameter spacing chosen here resulting in an average muon event requiring 20 to 30 minutes of processing to complete. This is roughly the actual muon detection rate at SNO. Increasing the resolution of the accumulator array by a factor of two in each direction results in 2^3 times more processing steps and 2^4 time more space for storing the accumulator array.

Overall, the two methods seem to provide promising results. While the smoothing algorithm performs well, handling the large accumulator array from the hough transform appears to be biggest challenge at this point. An increase in computing power and resolution elements will certainly help in refining the overall method.

Appendix C

Muon Vertical Depth Intensities

Depth (m.w.e.)	MC Events	Data Events	MC Vertical Intensity $\times 10^{-11}(\mathrm{cm}^{-2}\mathrm{s}^{-1}\mathrm{sr}^{-1})$	Data Vertical Intensity $\times 10^{-11}(\mathrm{cm}^{-2}\mathrm{s}^{-1}\mathrm{sr}^{-1})$
6225	3979	4207	35.746 ± 0.567	37.790 ± 0.583
6275	3681	3906	33.336 ± 0.549	35.375 ± 0.566
6325	3519	3584	32.129 ± 0.542	32.721 ± 0.547
6375	3352	3369	30.850 ± 0.533	31.006 ± 0.534
6425	3037	3250	28.172 ± 0.511	30.146 ± 0.529
6475	2827	2997	26.431 ± 0.497	28.020 ± 0.512
6525	2617	2737	24.661 ± 0.482	25.790 ± 0.493
6575	2459	2612	23.348 ± 0.471	24.805 ± 0.485
6625	2398	2372	22.947 ± 0.469	22.700 ± 0.466
6675	2192	2185	21.135 ± 0.451	21.069 ± 0.451
6725	2121	2043	20.607 ± 0.447	19.848 ± 0.439
6775	1902	1909	18.621 ± 0.427	18.687 ± 0.428
6825	1934	1831	19.074 ± 0.434	18.058 ± 0.422
6875	1623	1678	16.127 ± 0.400	16.673 ± 0.407
6925	1603	1544	16.047 ± 0.401	15.454 ± 0.393
6975	1508	1375	15.204 ± 0.392	13.862 ± 0.374

Table C.1 – Continued from previous page

Depth (m.w.e.)	MC Events	Data Events	MC Vertical Intensity $\times 10^{-11}(\mathrm{cm}^{-2}\mathrm{s}^{-1}\mathrm{sr}^{-1})$	Data Vertical Intensity $\times 10^{-11}(\mathrm{cm}^{-2}\mathrm{s}^{-1}\mathrm{sr}^{-1})$
7025	1494	1362	15.175 ± 0.393	13.832 ± 0.375
7075	1386	1249	14.179 ± 0.381	12.777 ± 0.362
7125	1254	1167	12.922 ± 0.365	12.023 ± 0.352
7175	1279	1123	13.272 ± 0.371	11.653 ± 0.348
7225	1230	1038	12.854 ± 0.367	10.847 ± 0.337
7275	1114	906	11.724 ± 0.351	9.534 ± 0.317
7325	954	899	10.109 ± 0.327	9.527 ± 0.318
7375	927	831	9.892 ± 0.325	8.869 ± 0.308
7425	923	791	9.917 ± 0.326	8.499 ± 0.302
7475	832	764	9.001 ± 0.312	8.264 ± 0.299
7525	817	679	8.898 ± 0.311	7.396 ± 0.284
7575	783	717	8.586 ± 0.307	7.863 ± 0.294
7700	2612	2244	7.284 ± 0.143	6.258 ± 0.132
7900	2149	1792	6.152 ± 0.133	5.129 ± 0.121
8100	1717	1380	5.042 ± 0.122	4.052 ± 0.109
8300	1303	1097	3.921 ± 0.109	3.302 ± 0.100
8500	1065	858	3.284 ± 0.101	2.646 ± 0.090
8700	824	673	2.602 ± 0.091	2.126 ± 0.082
8900	647	507	2.091 ± 0.082	1.638 ± 0.073
9100	536	443	1.772 ± 0.077	1.465 ± 0.070
9300	398	330	1.345 ± 0.067	1.115 ± 0.061

Table C.1 – Continued from previous page

Depth (m.w.e.)	MC Events	Data Events	MC Vertical Intensity $\times 10^{-11}(\mathrm{cm}^{-2}\mathrm{s}^{-1}\mathrm{sr}^{-1})$	Data Vertical Intensity $\times 10^{-11}(\mathrm{cm}^{-2}\mathrm{s}^{-1}\mathrm{sr}^{-1})$
9500	321	257	1.109 ± 0.062	0.888 ± 0.055
9700	271	204	0.957 ± 0.058	0.720 ± 0.050
9900	210	181	0.757 ± 0.052	0.652 ± 0.048
10250	393	291	0.589 ± 0.030	0.436 ± 0.026
10750	211	169	0.331 ± 0.023	0.265 ± 0.020
11250	132	98	0.217 ± 0.019	0.161 ± 0.016
11750	80	63	0.136 ± 0.015	0.108 ± 0.014
12250	54	34	0.097 ± 0.013	0.061 ± 0.010
12750	31	31	0.058 ± 0.010	0.058 ± 0.010
13250	15	14	0.029 ± 0.008	0.027 ± 0.007
13750	11	11	0.022 ± 0.007	0.022 ± 0.007
14250	13	12	0.027 ± 0.008	0.025 ± 0.007
14750	5	6	0.011 ± 0.005	0.013 ± 0.005

Table C.1: Table of muon vertical depth intensities from MC data and real data. The width of the depth bins change at 7600 m.w.e. and 10 000 m.w.e..