MEASUREMENT OF THE ATMOSPHERIC NEUTRINO INDUCED MUON FLUX AT THE SUDBURY NEUTRINO OBSERVATORY

Christopher Conrad Maximillian Kyba

A DISSERTATION

in

Physics and Astronomy

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

2006

Eugene W. Beier, Supervisor of Dissertation

Randall D. Kamien, Graduate Group Chair

For my parents

Acknowledgements

This work would not have been possible without the collaboration, guidance, assistance, and friendship of many people. Gene Beier was my thesis advisor, and I am very greatful to him for taking me on, and for providing me with direction on occasions when I wasn't sure what to do next. Gene has the remarkable ability to say twice as much in half the time it takes most of us, and hearing him summarize the statements of others demonstrated the importance of clarity and precision in speaking. I was lucky to have substantial interactions with Josh Klein, who personifies my idea of a "scientist". Bill Heintzelman sat at the desk next to me three days a week, and was an amazing resource for problems related to statistics and programming. Rick Van Berg taught me about the trigger system on my first day underground, and over the next several years taught me everything I know about electronics.

I owe a lot to my fellow graduate students Monica Dunford and Vadim Rusu, who were both fantastic colleagues and friends. As far as I'm concerned they were the best part of grad school. I would also like to thank Greg Grason, Yehuda Sneer, Olivia Halt, and Reiko Nakajima for their friendship and physics advice over the years. As I began analysis work the Penn SNO postdocs were very helpful. I discussed muon event selection with Neil McCauley, and the CITR cut was his suggestion. I had several productive discussions with Huaizhang Deng, who suggested I use a linear discriminant to demonstrate the instrumental contamination was small. Jeff Secrest brought a collider perspective to our group, and taught me the word "cepstrum". I know I will miss having Paul Keener as a sys admin, as well as the technical help of Chuck Alexander, Godwin Meyers, and Mike Reilly. Oveida Norris and Puneeth Kalavase were Penn undergrads that I had the pleasure of teaching and working with.

I spent a great deal of time in Sudbury, and I would like to thank Jeanne Wilson, Jess Dunmore, and Charlotte Sims from Oxford for the terrific times we had. Teresa Spreitzer was fun to be with and a huge help when she worked for our group. Thanks to Noel Gagnon and Peter Skensved for your increadible dedication to SNO. I would also like to thank Fraser Duncan, Richard Van de Water, Kate Frame, and the SNO operations group for keeping the experiment going. A big thanks to the SNO muon working group, Tyron Tsui, Charles Currat, and especially to Joe Formaggio for all the productive discussions. Thanks also to my fellow run selectors, and to the MCProd group, particularly David Waller and Feng Zhang who helped with the muon analysis. James Cameron from developed the SNO phototube calibration, and I referred to his thesis on many occasions. Nick West did an amazing job of managing SNOMAN and squashing a few memorable bugs.

Philadelphia was a great place to live, especially thanks to Jake Munger, Solveig Argeseanu, Maya Ravindranath, and Christine Cox. My colleagues in GET-UP gave me hope for the future of America and the academy; Tina Collins, Clay Bennet, Elise Carpenter, and Kevin Jude were especially great friends. As an international student I would also like to thank the government and the people of the United States for funding this research and welcoming me into their country.

I have dedicated this work to my amazing parents, who have always inspired and believed in me. Thanks also to my extended families on both sides of the Atlantic for your love and support over all these years. Finally, my thanks to Tatjana Scheffler for her love, support, and LaTeX assistance. A complete description of my gratitude to Tanja is beyond the scope of this acknowledgement.

ABSTRACT

MEASUREMENT OF THE ATMOSPHERIC NEUTRINO INDUCED MUON FLUX AT THE SUDBURY NEUTRINO OBSERVATORY

Christopher Conrad Maximillian Kyba

Supervisor: Eugene W. Beier

This work presents a measurement of the neutrino induced through-going muon flux in a 337.25 day exposure at the Sudbury Neutrino Observatory. Due to the 6km w.e. depth and the flat overburden of the SNO detector, cosmic muons range out at a zenith angle near $\theta_z = 66°$. A total of 149 neutrino induced through-going muon candidate events were observed in the range $-1 < \cos\theta_z < 0.4$, corresponding to a muon flux of $2.54 \pm 0.21(\mathrm{stat}) \pm 0.08(\mathrm{syst}) \times 10^{-13} \mu/(\mathrm{cm}^2\,\mathrm{s}\,\mathrm{sr})$. A maximum likelihood parameter estimation is used to extract the best fit neutrino flux and two-flavor neutrino oscillation parameter $\Delta\mathrm{m}^2$. The data are found to be most consistent with a flux of 1.2 times that predicted by the Bartol group, with $\Delta\mathrm{m}^2 \approx 3 \times 10^{-3}\mathrm{eV}^2$. Due to the limited statistics of this data set it is not possible to rule out the no-oscillation hypothesis.

The depth to vertical intensity relationship for cosmic muons at the SNO site is studied, and is found to be consistent with the results of past experiments. The expected seasonal variation of the cosmic muon flux is observed over a period of almost four years. A sinusoidal fit to this annual variation is found to have an amplitude of $2.4 \pm 0.6\%$. A study of neutron events that occur within 200ms after the passage of muons through the SNO $\mathrm{D_2O}$ is performed in two phases of the SNO experiment. It is found that for muons with energy loss below 400MeV/m the production of neutrons increases linearly with muon radiative energy loss.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Ich [...] wende mich vertrauensvoll an Euch, liebe Radioaktive, mit der Frage, wie es um den experimentellen Nachweis eines solchen Neutrons stände.

Wolfgang Pauli, 1930

It has been three quarters of a century since Pauli urged his ”radioactive” colleagues to dream of ways to detect the neutrino, and in the last half century experimentalists have obliged. Experiments have detected neutrinos produced at nuclear reactors [1], in the core of our Sun [2], at particle accelerators [3], in the Earth’s atmosphere [4], and even in the explosion of a giant star in a companion galaxy [5, 6]. Experiments have shown us that there are three families of neutrinos that couple to the weak force [7]. Solar, reactor, and atmospheric neutrino experiments have shown us that neutrinos can mix, or oscillate, between these families as they propagate [8, 9, 10].

As time has passed, the goal of experiments has changed from observing neutrinos to studying their properties. The principle goal of this thesis is to study atmospheric neutrinos in the GeV-TeV energy range at a much shorter distance from their production point than has previously been possible. At this short distance it is expected that the neutrinos will not have had time to oscillate, and thus by observing these neutrinos the total flux of atmospheric neutrinos will be measured.

1.1 Neutrino Masses and Mixing

Neutrinos are the only uncharged fermions, and for reasons that are as yet unknown their mass is many orders of magnitude smaller than that of the quarks and charged leptons. The recent solar, reactor, and atmospheric neutrino experiments have taught us that neutrinos do have mass, and that lepton number is not a conserved quantity. This section will discuss the theory behind neutrino masses and neutrino oscillation, and the experimental evidence for neutrino mass.

1.1.1 Neutrinos in the Old Standard Model

From the very beginning of the neutrino's status as a particle it was known that its mass had to be very small. In his famous letter to the conference at Tübingen, Pauli suggested the hypothetical particle's mass would be at most on the order of the electron mass. Within four years Fermi's theory of nuclear beta decay indicated that the neutrino had either no mass, or at least orders of magnitude less mass than the electron. With the discoveries that parity is violated in weak decays and that all observed neutrinos are left handed [11], it became natural to model the neutrino as a massless two component field. The first generation of quarks and leptons in the Standard Model was represented by a set of left handed doublets, with corresponding right handed singlets for massive particles.

$$\text{quarks:} \begin{pmatrix} u \\ d \end{pmatrix}_L \; u_R \; d_R \qquad \text{leptons:} \begin{pmatrix} \nu_e \\ e^- \end{pmatrix}_L \; e^-_R \tag{1.1}$$

Mass is generated in the Standard Model through a coupling of left and right handed fields through the Higgs mechanism. The mass term for an electron is given by

$$\mathcal{L} = m\overline{e_L}e_R + m\overline{e_R}e_L \tag{1.2}$$

where the e_L and e_R are left and right chiral electron fields. The right chiral neutrino field was left out by design [12] in order to produce the result that there are no neutrino mass terms, making all neutrinos left handed[1].

There are three generations of quarks and leptons in the Standard model, and all twelve of these particles have been observed experimentally. In the lepton sector the order of the charged particles in terms of increasing mass is e, μ, τ, so although the neutrinos were massless in the old Standard Model it was customary to give the neutrinos the same ordering: ν_e, ν_μ, ν_τ.

1.1.2 Neutrinos in the Current Standard Model

With the discovery that neutrinos are massive it was no longer necessary to exclude the right chiral neutrino field in Equation 1.1. With a right chiral neutrino field the standard Dirac mass terms can easily be used to generate neutrino mass:

$$\mathcal{L} = m_D\overline{\nu_{e_L}}\nu_{e_R} + m_D\overline{\nu_{e_R}}\nu_{e_L} \tag{1.3}$$

Neutrinos are the only neutral fermions in the standard model, so an open question in neutrino physics is whether or not the neutrino is actually distinct from the antineutrino. If the two are actually the same particle (as is the case for the photon and Z boson) then

[1]For massless particles helicity is a good quantum number, and is identical to chirality. Since the neutrino was by construction massless, all neutrinos in the old Standard Model had left handed helicity.

Majorana mass terms are also allowed:

$$\mathcal{L}_M = m_M \overline{\nu_{e_L}^c} \nu_{e_L} + m_M \overline{\nu_{e_L}} \nu_{e_L}^c \tag{1.4}$$

In this situation the right chiral anti neutrino field is the charge conjugate of the left chiral neutrino field, $\nu_R = \nu_L^c$. The nature of neutrino masses has no bearing on the analyses in this thesis, but two points are worth mentioning briefly. The first is that if neutrinos are Majorana particles then the smallness of the neutrino mass relative to the charged fermion masses could be explained through the so-called "see-saw mechanism" [13, 14]. If neutrinos are Dirac particles there is no known reason for their tiny mass. The see-saw mechanism is viewed as attractive because it explains this phenomenon, albiet at the expense of introducing another new (and unobserved) field. The second is that if neutrinos are Majorana particles and the lightest neutrino mass is large enough, it could be possible to observe "neutrinoless double beta decay". If neutrinoless double beta decay is observed it would demonstrate that neutrinos are Majorana particles, and the absolute neutrino mass scale could also be measured.

1.1.3 Neutrino Oscillation

The first indication that there was a problem with the Standard Model description of massless neutrinos came from Ray Davis's chlorine solar neutrino experiment in the Homestake gold mine. Over a period of decades the Homestake experiment observed [15] far fewer neutrinos than were predicted [16, 17] based upon the standard solar models. The Kamiokande-II water Cherenkov experiment confirmed the deficit, and furthermore noticed a deficit of upward going atmospheric neutrinos. These two problems came to be known as the solar neutrino problem and the atmospheric neutrino anomaly. A host of models were proposed to attempt to explain the discrepant results; after several decades of research the massive neutrino oscillation hypothesis is the only one which can explain the data. The massive neutrino hypothesis asserts that the neutrino weak flavor eigenstates are not the same as the neutrino propagation (mass) eigenstates.

Neutrinos are produced and absorbed by the weak charged current interaction. The measurement of the invisible width of the neutral Z boson at the LEP experiments indicates that there are three families, also known as flavors, of neutrinos with masses below 45GeV, ν_e, ν_μ, and ν_τ. The massive neutrino oscillation hypothesis asserts that there exist three propagation eigenstates, ν_1, ν_2, and ν_3, which are not aligned with the flavor eigenstates. Consequently, neutrinos produced in one of the three flavor eigenstates can be expressed as a mixed mass state. The conversion between flavor and propagation eigenstates is expressed in terms of a unitary mixing matrix, U. The pure neutrino flavor state $|\nu_\alpha>$ is then expressed in the propagation eigenstates as

$$|\nu_\alpha>=\sum_{k=1}^{3} U^*_{\alpha k}|\nu_k> \tag{1.5}$$

If ν_1, ν_2, and ν_3 are not degenerate then a neutrino produced in a pure flavor state will oscillate between the three flavor states as it propagates. As soon as the neutrino begins to travel away from the production point it will no longer be in a pure flavor eigenstate. This means that experiments such as the Homestake experiment, which are only sensitive to the original flavor eigenstate, may therefore observe a deficit of neutrinos. The algebra involved in two neutrino mixing is considerably simpler than that for three neutrinos, and in most cases two neutrino mixing does an adequate job of describing oscillations for a single experiment, for reasons we will see in a moment. Two neutrino mixing is briefly sketched below.

For two flavors the unitary mixing matrix which converts between the flavor and propagation eigenstates can be described simply as

$$\mathrm{U} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \tag{1.6}$$

Where θ is a "mixing angle" that parameterizes how misaligned the flavor and propagation eigenstates are. A neutrino produced in a pure ν_μ state will evolve with distance[2] according to

$$|\nu(L)>= \sum_{k=1}^{2} U^*_{\mu k} e^{-i\frac{m_k^2}{2E}L} |\nu_k> \tag{1.7}$$

Thus, the transition probability depends on the difference of the squares of the two masses, $\Delta m^2_{23} = \mathrm{m}_3^2 - \mathrm{m}_2^2$. Specifically, for the mixing of ν_μ and ν_τ neutrinos:

$$P_{\nu_\mu \to \nu_\tau} = \sin^2(2\theta_{23})\ \sin^2\left(1.27\Delta m^2_{23}\frac{L}{E}\right) \tag{1.8}$$

where Δm^2_{23} is expressed in eV^2, L in km, and E in GeV. Rather than measuring the transition probability, atmospheric neutrino and solar neutrino experiments usually measure the survival probability. For $\nu_\mu \to \nu_\tau$ this is:

$$P_{\nu_\mu \to \nu_\mu} = 1 - P_{\nu_\mu \to \nu_\tau} \tag{1.9}$$

$$P_{\nu_\mu \to \nu_\mu} = 1 - \sin^2(2\theta_{23})\ \sin^2\left(1.27\Delta m^2_{23}\frac{L}{E}\right) \tag{1.10}$$

Equation 1.10 has three parameters, two of which are set by nature (and under investigation) and one of which is experimental, the oscillation parameter $\frac{L}{E}$. The mixing angle tells us how large the oscillation signal will be. If the mixing angle is zero then the flavor and propagation eigenstates are identical, and there will be no mixing. If the angle is maximal ($\theta = 45°$) then as L increases the survival probability drops from one to zero, and

[2] The neutrino mass is so small that they are always relativistic, so time and distance are interchangeable.

then rises back to one. Detector effects such as energy resolution tend to smear out the oscillation signal, so at large enough values of $\frac{L}{E}$ the survival probability is given by:

$$P_{\nu_\mu \to \nu_\mu} = \frac{1}{2} \sin^2 2\theta \tag{1.11}$$

In section 1.3 we will consider what effect the survival probability equation has in different atmospheric neutrino energy windows.

Returning to Equation 1.10 we can also see why the two neutrino approximation is valid in the atmospheric neutrino sector. The measured value of Δm^2 from solar neutrino signals is a factor of 30 smaller than for atmospheric neutrinos, so for the $>$GeV neutrinos which are the subject of this thesis atmospheric muon neutrinos have not propagated far enough to have a chance to oscillate. At sub GeV energies the situation is not trivial. Peres and Smirnov have shown [18] that at these low energies the oscillations are masked by the fact that both muon and electron neutrinos are present in the atmospheric flux. They show that effect of the oscillations on the electron neutrino flux is proportional to the value $(r \cos^2 \theta_{23} - 1)$, where r is the ratio of the muon neutrino to electron neutrino fluxes. In the following sections we will see that θ_{23} is close to $45°$ and that at sub GeV energies r is very close to 2. Thus, the electron neutrino oscillations are suppressed, and the oscillations of atmospheric neutrinos are well described with two neutrino oscillation at all energies.

1.1.4 Evidence for Non-Zero Neutrino Mass

Several experiments have used the phenomenon of neutrino oscillation to demonstrate conclusively that at least two of the neutrino propagation eigenstates have non-zero mass. The SNO experiment was the first to demonstrate conclusively (at 5.3σ) that neutrinos change their weak interaction "flavor" between their production in the sun and their arrival at earth [8]. The combined analysis of the $\sim$MeV solar neutrino signal at SNO and the

~MeV reactor neutrino signal at KamLAND[3] [9] have observed neutrino oscillation with $\Delta m^2 = (8.0^{+0.6}_{-0.4}) \times 10^{-5}$ eV^2 and $\theta = 33.9^{\circ +2.4}_{-2.2}$ [19]. This indicates that one of the neutrino propagation states has a mass of at least $\sim 8.4 \times 10^{-3}$eV. By studying the higher energy atmospheric neutrinos, the Super-Kamiokande experiment has also observed very strong evidence for neutrino weak eigenstate flavor change. Through its observation of atmospheric neutrinos over an energy range of six orders of magnitude (30MeV-~10TeV), the Super-K experiment has observed neutrino oscillation with a different set of parameters, with $\theta \approx 45^\circ$ and $\Delta m^2 \approx 2.2 \times 10^{-3}$ eV^2. This indicates that one of the other neutrino propagation eigenstates must have a mass of at least 0.032eV [10]. While neutrino oscillation experiments have set lower limits on the minimum size of two of the three neutrino propagation eigenstate masses, there has not yet been a direct measurement of the lightest neutrino mass.

1.2 Cosmic Rays, Cosmic Ray Muons, and Atmospheric Neutrinos

While the space between planets, stars, and galaxies is a vacuum, it is not empty. In inter stellar space charged particles, mainly protons, stream in all directions. These particles are known as cosmic rays, and their energy spectrum extends from non-relativistic to ultra-relativistic (10^{20}eV) energies. The magnetic field of the Earth is known to accelerate particles up to energies in the keV to MeV range. The intense fields produced in solar flares have been known to accelerate particles up to the GeV scale [20]. Supernovas, binary star systems, and black holes are among the candidates for the generation of higher energy cosmic ray primaries. The maximum energy and the origin of the highest energy cosmic rays

[3]Because SNO observes neutrinos and KamLAND observes anti-neutrinos, the combined analysis requires the assumption that CPT is conserved. If CPT is conserved then the propagation eigenstates of neutrinos and anti-neutrinos must have the same mass.

are open questions in particle astrophysics.

When cosmic rays with TeV energies or higher strike the Earth's atmosphere they interact to produce secondary particles, and these secondary particles further interact to develop a "cosmic ray shower". Muons and pions are among the most common particles produced, and both of these particles were first discovered in cosmic ray shower studies. Neutrinos are produced through the decay of particles in cosmic ray showers, and these "atmospheric neutrinos" have been used to study neutrino oscillation. The following subsections will discuss the cosmic rays themselves, cosmic ray showers, cosmic ray muons in the air and underground, and finally the production of atmospheric neutrinos. The next section will discuss the detection of atmospheric neutrinos.

1.2.1 Cosmic Ray Spectrum

The cosmic ray energy spectrum is shown in Figure 1.1. The spectrum can be well described over many orders of magnitude by an inverse power law, with the differential flux given by

$$\frac{dN}{dE} \propto E^{-(\gamma+1)} \tag{1.12}$$

At the lowest energies $\gamma \sim 1.7$. Near $10^{6.5}$ GeV the spectrum steepens to $\gamma \sim 2.0$ at a feature that has become known as the "knee". Above $10^{17.7}$ eV the spectrum steepens further to $\gamma \sim 2.3$, and finally flattens back to $\gamma \sim 1.7$ at $10^{18.5}$ eV, the "ankle" [21]. The changes in the power law are suspected to be due to changes in the mechanism responsible for the acceleration.

1.2.2 Cosmic Ray Showers

When cosmic rays reach the upper atmosphere of the Earth they undergo hadronic interactions with the nuclei of atmospheric gases. Depending on the energy of the cosmic primary,

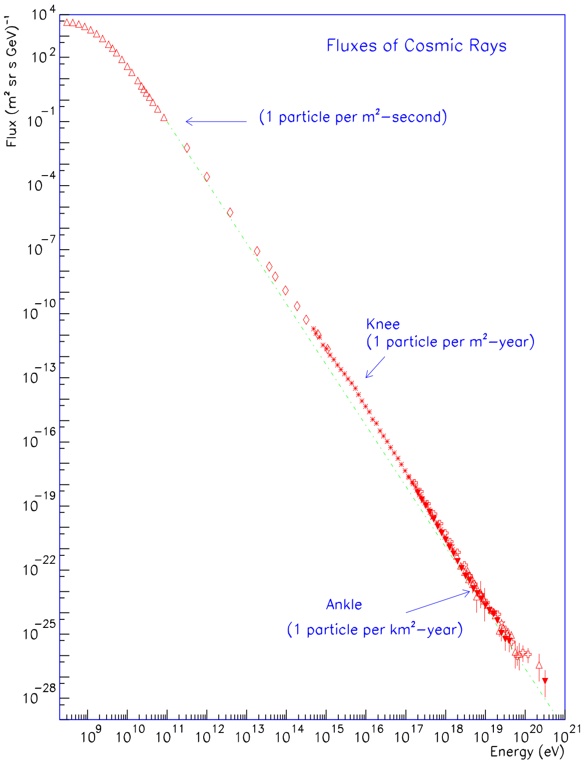

Figure 1.1: Observed flux of cosmic rays over 14 orders of magnitude. The dotted line shows E^{-3} for comparison. Plot is taken from [21].

these interactions may take place at energies orders of magnitude larger than those accessible in modern particle accelerators. The result of these collisions is the production of secondary hadrons, most importantly pions and kaons. These secondary pions and kaons may decay or interact, depending upon their energy and their depth in the atmosphere. At low altitudes the atmospheric density becomes large enough that secondary hadrons are attenuated; this is why the pion was discovered by cosmic ray experiments situated on mountain-tops rather than in the earlier sea-level experiments.

As the primary cosmic ray and its daughter particles continue to interact with the atmosphere a “cosmic ray shower” develops. The size, geometry, extent, and particle content of cosmic ray showers depend upon the energy of the primary, the height in the atmosphere at which the primary first interacts, and the angle (relative to surface vertical) at which the primary is traveling. Higher energy cosmic rays result in larger, more extensive showers. The higher in the atmosphere and the more “horizontal” the shower is, the more time the secondary hadrons have to decay rather than interact. Due to interaction kinematics, the highest energy particles are bunched in the center of the shower; secondary particles at the edges have more isotropic directions and far lower energies.

1.2.3 Cosmic Ray Muons in the Atmosphere

At sea-level the hadronic component of the cosmic rays has been almost entirely attenuated. What remains is a “hard component” of cosmic ray muons, and a “soft component” of photons and electrons. It is possible to attenuate the soft component with a few meters of dirt, so for our purposes we need only consider the muon production in cosmic ray showers. At low energies, muons are produced in cosmic ray showers primarily through the following decay:

$$\pi^+ \longrightarrow \mu^+ + \nu_\mu \tag{1.13}$$

The charge conjugate process contributes at a lower level due to the fact that the hadronic cosmic ray primaries have positive charges. If the pions are energetic enough they will interact before they have time to decay, so the decay length to interaction length ratio for pions in the atmosphere determines whether the decay is likely to occur. As was mentioned earlier, deep in the atmosphere the interaction length is far shorter than the decay length, so muon production from pions decaying in flight shuts down. The atmospheric density gradient is less steep for muons traveling parallel to the Earth's surface, and for this reason the cosmic ray muon flux from the horizon is greater than that from the vertical. At low energies, and above $\theta_z \sim 60°$ the muon flux goes as $\sec\theta_z$, where θ_z is the zenith angle. At energies above $\sim$100GeV the decay length to interaction length ratio for pions becomes small, so the following kaon decay becomes important:

$$K^+ \longrightarrow \mu^+ + \nu_\mu \qquad\qquad (\sim 63\%\ \mathrm{BR}) \tag{1.14}$$

1.2.4 Cosmic Ray Muons Underground

Muons are not only able to penetrate the several kilometers of atmosphere needed to reach sea level, the most energetic ones are also able to penetrate to great depths underwater or underground. As muons propagate through material they lose energy through interactions with that material. At low energies the energy losses are dominated by ionization of the surrounding medium. At higher energies, losses through discrete radiative interactions begin to dominate. The "critical energy" is defined as the energy at which muons lose an equal amount of energy through ionization and radiative processes; the critical energy for muons in rock is about 500GeV. Above the critical energy muons lose energy primarily through bremsstrahlung, direct e^+e^- pair production, and electromagnetic interactions with nuclei. Losses through pair production are just slightly larger than losses through bremsstrahlung, while losses through interactions with nuclei are about an order of magnitude smaller. In

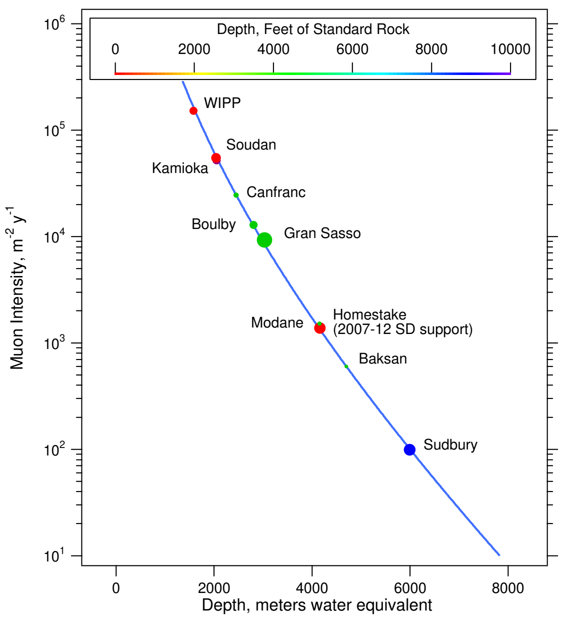

Figure 1.2: The integrated muon flux is shown as a function of depth, with the location of several underground laboratories overlaid. The relative sizes of the laboratories are indicated by the circle size. Figure courtesy Hamish Robertson. An equivalent figure that indicates whether the sites are in mountains or under a flat overburden can be found in [23].

contrast to ionization loss, which in general may be treated as continuous, radiative loss is discrete. Thus simulations of muon propagation must be done stochastically. The uncertainty on each of the muon energy loss mechanisms is in the range 1-2%, with the exception of photonuclear interactions, which have an uncertainty of 10-20% [22].

Each energy loss mechanism described above is proportional to the density of the material that the muon is traveling through. In order to compare the relative depths of experimental sites underwater, under ice, or underground, the unit “meters of water equivalent” (m w.e.) is often used to describe the detector’s depth. By this measure SNO is an extremely deep experiment at 6km w.e., as can be seen in Figure 1.2.

Any cosmic ray muon which manages to propagate deep underground has to have been produced in a shower caused by a very energetic primary. For this reason it is possible for an underground detector to observe more than one muon at once. The MACRO experiment in the Gran Sasso Laboratory has observed that the muon multiplicity distribution in their 12m×76.5m detector extends to greater than 35 muons in a single event [24]! About 6% of the muon events in the MACRO detector were reported to have a multiplicity greater than one.

1.2.5 Production of Atmospheric Neutrinos

Atmospheric neutrinos are predominantly produced by pion, muon, and kaon decay in the atmosphere. At low energies pion and muon decays dominate, giving rise to the following decay chain:

$$\begin{aligned} \pi^+ &\longrightarrow \mu^+ + \nu_\mu \\ &\mu^+ \longrightarrow e^+ + \nu_e + \overline{\nu}_\mu \end{aligned} \tag{1.15}$$

with the charge conjugate process contributing as well. One can easily see that if this chain is dominant then the atmospheric neutrino flavor ratio,

$$\frac{\nu_e + \overline{\nu}_e}{\nu_\mu + \overline{\nu}_\mu} \tag{1.16}$$

should have a value of about 1/2. At an energy of about 2.5GeV muons in the atmosphere become energetic enough to reach the Earth before decaying. Due to the inability of the muons to decay the atmospheric neutrino flavor ratio quickly decreases. At very high energies electron neutrinos are only produced in the decay of the relatively rare neutral kaon,

$$K_L^0 \longrightarrow \pi^+ + e^- + \overline{\nu}_e \qquad\qquad (\sim 39\%\ \mathrm{BR}) \tag{1.17}$$

so at high energies the electron neutrino flux is very small compared to the muon neutrino flux.

At low energies (below the critical energy of pions in the upper atmosphere) the energy spectrum of the neutrinos has the same power as the cosmic rays ($\gamma = 1.7$). As the energy increases the spectrum steepens, until at energies much larger than the pion critical energy in air the spectrum has a power of $\gamma \approx 2.7$. The reason that the spectrum changes at the pion critical energy is that above that energy the decay length to interaction length ratio becomes small, so the neutrino flux from the first line of Equation 1.15 is reduced. Above the pion critical energy, kaons quickly become more important for neutrino production. This is because the critical energy for charged kaons in the atmosphere is near 850GeV, more than seven times greater than the critical energy of pions (115GeV). The kinematics of meson decay also favors kaons, as the available energy is split more evenly between the produced μ and ν_{μ}.

As was the case for muons, the flux of atmospheric neutrinos is increased for cosmic rays in which the zenith angle of the primary is close to horizontal. This is once again due to the fact that the decay length to interaction length ratio is larger at low atmospheric densities, and the shallower the cosmic ray's path the less quickly the density increases. This also means that the energy at which kaon decay becomes more important than pion decay for muon production is actually zenith angle dependent.

The atmospheric neutrinos which will be studied in this thesis are produced by cosmic rays in with energies in the range of 1 GeV-100TeV Experimental measurements of the cosmic ray flux in this range shown in Figure 1.3. The uncertainty due to the cosmic primary flux is the dominant uncertainty in the prediction of the neutrino flux, at 20-30% [25].

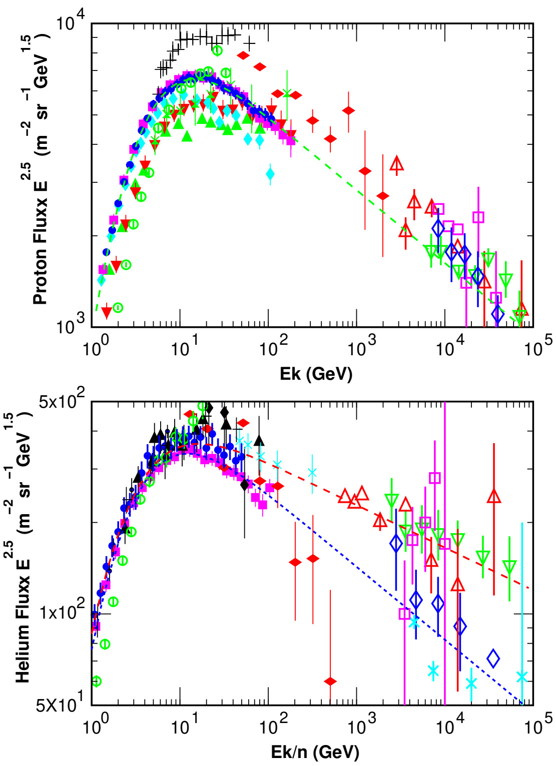

Figure 1.3: Observed flux of cosmic-ray protons and helium as a function of kinetic energy per nucleon. The disagreement of the data contributes to the large overall normalization uncertainty in the atmospheric neutrino flux. Plot is taken from [25].

The Effect of Earth's Magnetic Field on the Neutrino Flux

The Earth's magnetic field has two effects upon the observed cosmic ray spectrum. At low enough energies the field deflects cosmic ray particles before they ever enter the atmosphere. At higher energies the deflection causes an asymmetry in the azimuthal distribution of cosmic rays relative to the detector z. Both of these effects are dependent upon the latitude of the detector relative to Earth's magnetic pole, and are considerably less important for experiments like Soudan and SNO, which are located near the magnetic North Pole, than they are for experiments with larger geomagnetic latitudes, such as the Gran Sasso site in Italy or the Kamioka site in Japan. In Sudbury geomagnetic effects may be ignored for cosmic energies greater than about 10GeV.

Barr *et al.* have kindly provided SNO with their "Bartol04" three dimensional flux predictions [26] for the Sudbury site. Plots of the ν_μ flux at SNO are shown for three energies in Figure[4] 1.4. Note that the "peaking" of the flux at the horizon ($\cos\theta_z = 0$) increases at higher energies, as was discussed earlier in the chapter.

Seasonal Flux Variation

The flux of cosmic ray muons is known to vary with the season, being largest in summer and smallest in winter. This effect is due to the "breathing" of Earth's atmosphere due to changes in its temperature. During the higher-temperature summers the atmosphere expands. This makes it possible for cosmic ray primaries to interact at higher altitudes, and increases the pion decay length to interaction length ratio. Due to this change more pions are able to decay in the summer than in the winter, and the cosmic muon flux is therefore proportionally larger. In the final chapter of this thesis it will be shown that the size of this effect is $\sim$2% for cosmic muons which reach the SNO detector.

[4] The data in the plot are for the period of maximum solar magnetic field activity, the effect of the solar field on the neutrino flux is discussed in the final chapter.

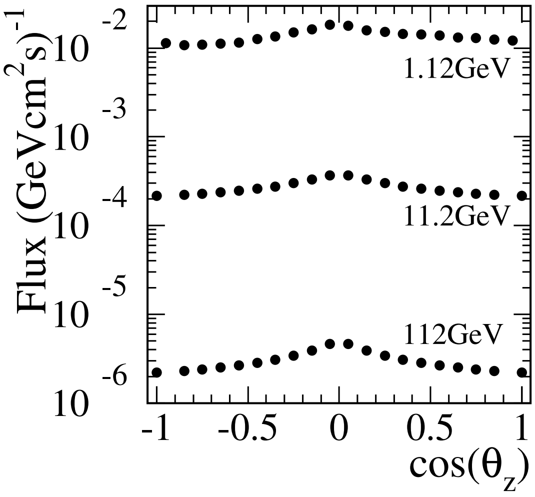

Figure 1.4: Bartol04 [26] flux of atmospheric muon neutrinos at SNO during the solar maximum. The variations at 1GeV are due to the Earth's magnetic field. As energy increases charged pions begin to reach the Earth's surface before decaying. This effect is larger for muons traveling straight down than for those from the horizon. This is the reason for the shape change with increasing energy.

The IceCube collaboration has estimated the amplitude of seasonal variations on the integrated neutrino flux above 50GeV to be $\sim 0.25\%$ at latitudes below 40°, rising to $\sim 2.2\%$ at the pole [27]. Given the latitudes and statistics available in the SNO and Super-Kamiokande experiments it is unlikely that either could observe a statistically significant variation in the atmospheric neutrino flux. The effect of atmospheric breathing on the atmospheric neutrino flux will probably be observable at the future km^3 detectors, especially in the IceCube detector at the South Pole. Due to IceCube's location there is an exact relation between detector zenith angle and latitude at which the atmospheric neutrino was produced, and this should increase the effect of the variation on the flux.

1.3 Detection of Atmospheric Neutrinos

Because of their neutral charge and tiny cross sections with matter, neutrinos cannot be directly detected in the same manner that hadrons and charged leptons are. Rather, experiments look for evidence of neutrinos by studying charged particles produced by interactions of neutrinos with matter. In Reines and Cowan's original neutrino detection, for example, anti-neutrino interactions with protons resulted in the conversion of the proton to a neutron with the production of an anti-electron. The neutrino signal in that case was the timing coincidence of the electron-positron annihilation with the neutron capture on cadmium. The Kamiokande-II experiment used the fact that in the neutrino-electron elastic scattering interaction the electrons are forward scattered in the neutrino propagation direction. This property allowed them to demonstrate that the sun is a neutrino source, and to measure the solar neutrino flux using a statistical extraction. The signature of atmospheric neutrinos can be observed in a detector in one of two ways. The neutrino can interact inside the detector, resulting in the production of particles the experiment is able to detect. Alternatively, a high energy muon neutrino can interact with the rock outside the detector to produce a muon, and that muon can then propagate through the rock to the detector.

1.3.1 Atmospheric Neutrino Event Classes

Atmospheric neutrino experiments of the past have typically distinguished between four types of neutrino interaction events. In order of increasing neutrino energy these are fully contained (FC), partially contained (PC), stopping upward traveling muons and through-going upward traveling muons.

Fully contained events are events in which the neutrino interacts with target material inside the detector and all charged particles remain inside the detector. The maximum energy at which a muon can be fully contained thus depends on the volume (length) of the detector. At low energies muons lose energy through ionization at a rate of near 2MeV/(g cm^{-2}), so a water Cherenkov detector with a length of 15m can only contain muons up to an energy of about 3GeV. Partially contained events are those in which the neutrino interacts inside the detector but one or more of the charged particles produced in the interaction leaves the detector. In water Cherenkov detectors it is possible to estimate the energy that the neutrino transferred in its interaction in fully contained events. In charge-current interactions in particular this provides a fairly accurate measurement of the neutrino energy. Water Cherenkov detectors can only put a lower limit on the neutrino energy for partially contained events. The rate of FC and PC events in a detector is proportional to its volume.

Through-going and stopping muon events are those in which a muon neutrino interacts with the target (rock, water, or ice) outside the detector and produces a muon with enough energy to reach the detector. Because of the great range of energetic muons this greatly increases the effective detector volume at large muon neutrino energies. Stopping muon events are those in which the externally produced muon stops in the detector (often producing a decay electron), through-going muons are those that pass right through the detector. Both the ratio of stopping to through-going muons and the average energy of the neutrinos which produced stopping muons increase with increasing detector size. Because the muon traveled for an unknown distance before entering the detector, it is not possible to estimate the neutrino energy for an individual through-going or stopping muon event. The total rate of through-going and stopping neutrino induced muons is proportional to the cross sectional area of the detector.

1.3.2 Neutrino Cross Section With Nuclei

The total cross section of neutrinos with nuclei is approximately linear for neutrino energies below about 3.5TeV [20], with a measured world average value given by [28] as

$$\sigma_\nu \sim 0.667 \pm 0.014 \times 10^{-38} cm^2 \times E_\nu \tag{1.18}$$

$$\sigma_{\overline{\nu}} \sim 0.334 \pm 0.008 \times 10^{-38} cm^2 \times E_\nu \tag{1.19}$$

Where E_ν is in GeV. At higher energies the denominator in the W propagator term of the weak charged current that produces muons begins to grow, and the total cross section increases more slowly. The uncertainties in the normalization of the neutrino flux and in neutrino cross sections with rock are very large ($\sim$20%) at contained event energies, and the uncertainties grow with energy. The predicted neutrino rate for three different flux models at Super-K, for example, differ by $\sim$5% to $\sim$20% depending upon the neutrino energy and signal type. In the past, atmospheric neutrino experiments have used "ratio of ratio" measurements, in order to reduce the importance of the systematic uncertainty on the neutrino flux. Two examples are the muon flux ratio of ratios

$$R_R = \frac{(\Phi_{stopping}/\Phi_{through-going})_{observed}}{(\Phi_{stopping}/\Phi_{through-going})_{theory}} \tag{1.20}$$

and the muon-electron ratio of ratios,

$$R_R = \frac{(\Phi_e/\Phi_\mu)_{observed}}{(\Phi_e/\Phi_\mu)_{theory}} \tag{1.21}$$

While ratio of ratio measurements make single experiment analyses easier, they make it difficult to compare the measurements of different experiments, as each experiment floats their estimate of the neutrino flux to a different value.

1.3.3 Atmospheric Neutrino Oscillations

Several experiments have observed atmospheric neutrinos, but the great volume ($\sim$50 kiloton) and running time ($\sim$1500 days) of the Super-Kamiokande experiment dwarfs any of the other measurements. The Super-K experiment has observed tens of thousands of neutrinos, and their most recent oscillation results are shown in Figure 1.5. The figure shows eleven different event samples, with the fully contained events split into three energy windows for both “electron-like” and “muon-like” events, two energy samples of events in which more than one Cherenkov ring was observed, partially contained events in which a muon left the detector, and neutrino induced through-going and stopping muons.

The Super-Kamiokande experiment defines a 39 parameter χ^2 statistic to compare predictions for each of the 180 bins of Figure 1.5 with the data. Their “best fit” oscillation parameters as a result of minimizing this χ^2 are maximal $\nu_\mu \rightarrow \nu_\tau$ mixing ($\sin^2 2\theta = 1$) with $\Delta_m^2 = 2.1 \times 10^{-3}\mathrm{eV}^2$. In Figure 1.5 the Super-K data are shown as points, the Monte Carlo prediction for one flux prediction with no oscillations is shown as bands, and the Monte Carlo prediction with Super-K’s best fit oscillation parameters and floated flux are shown as a histogram. Note that flux is not floated in MC bands; this is why in some bins the oscillated prediction is larger than the non-oscillated prediction.

With the Super-Kamiokande data as a backdrop let us now return to Equation 1.10:

$$P_{\nu_\mu \rightarrow \nu_\mu} = 1 - \sin^2(2\theta_{23})\ \sin^2\left(1.27\Delta m_{23}^2 \frac{L}{E}\right) \tag{1.10}$$

At maximal mixing the $\sin^2(2\theta_{23})$ term may be dropped, as it is equal to one. With $\Delta m^2 = 2.1 \times 10^{-3}\mathrm{eV}^2$ the first minimum of the survival probability equation takes place at an L/E value of $\sim$590km/GeV. The oscillation probability for three different energies is shown in the top frame of Figure 1.6. The distance that neutrinos travel from the top of the atmosphere to the detector as a function of $\cos\theta_z$ is shown in the bottom frame. Note that in the range $0 < \cos\theta_z < 1$ the distance to the atmosphere changes very little with $\cos\theta_z$, but that “below the horizon” it begins to change very rapidly.

The atmospheric neutrinos detected at Super-K have an energy range from 30MeV for the lowest detected atmospheric electron neutrinos to $\sim$few TeV in neutrino induced muon events, and their distance from production ranges from $\sim$15km to $\sim$12500km. While this means that atmospheric neutrino experiments can sample a wide band of L/E values it

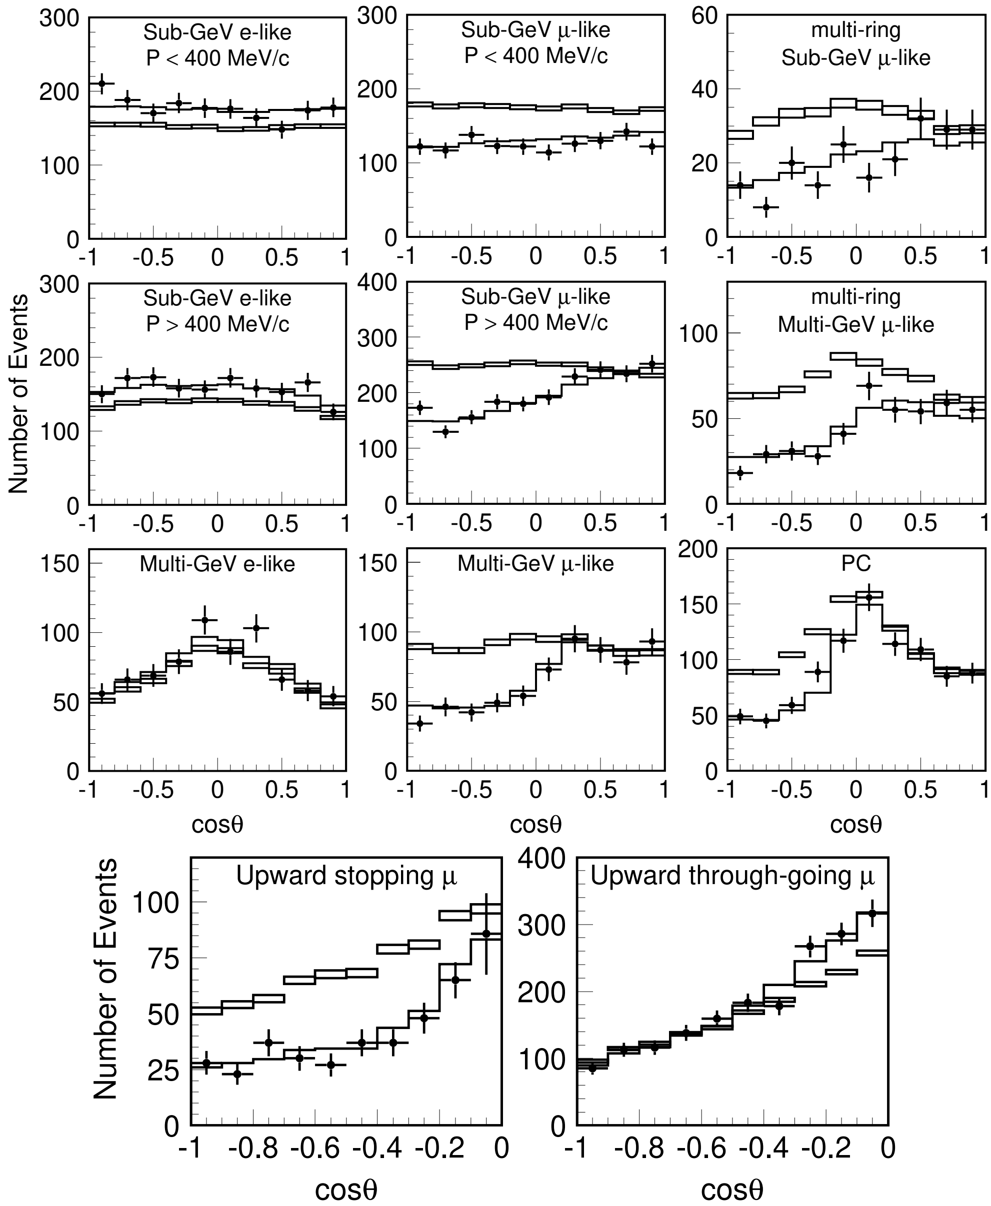

Figure 1.5: Zenith angle distribution of Super-Kamiokande data. The plots include 1489 live days of fully contained and partially contained events, and 1645.9 live days of upward going muon events. The data sample in each frame is individually labeled. Crosses show the Super-K data, the histogram shows the best-fit oscillation simulation, and the bands show the non-oscillation prediction. Plots are taken from [10].

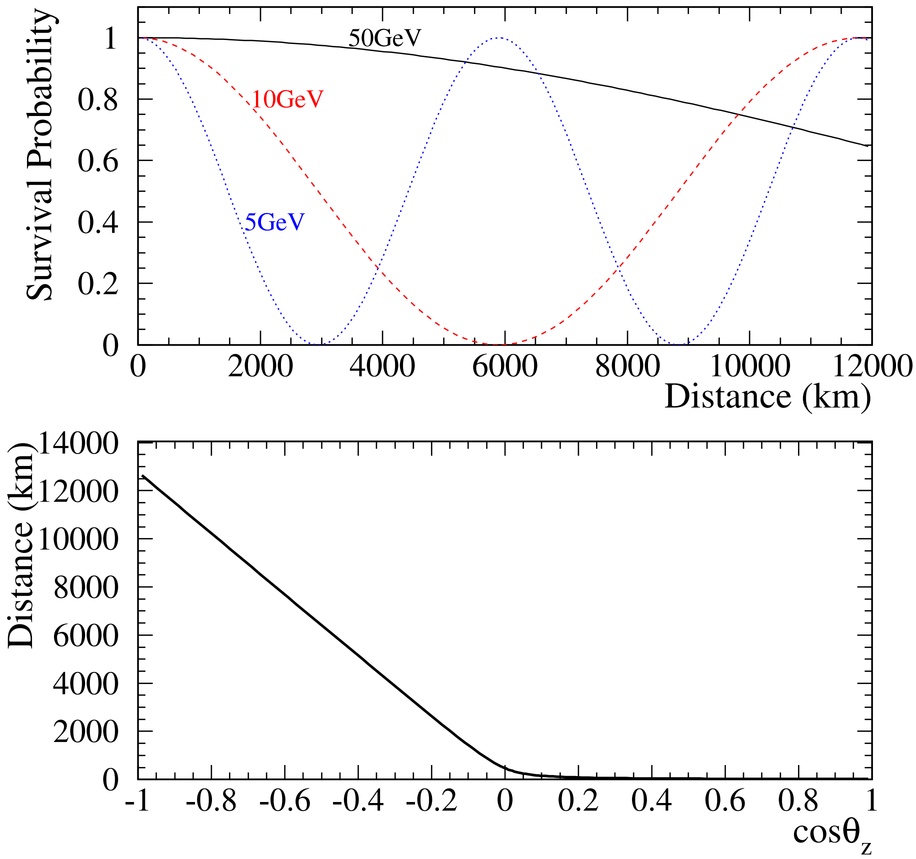

Figure 1.6: The top plot shows the survival probability as a function of distance from production for muon neutrinos at three exact neutrino energies. Super-Kamiokande's "best fit" oscillation parameters are used to calculate the survival probabilities. The bottom plot shows the distance from the detector to the top of the atmosphere as a function of $\cos\theta_z$.

also makes precision measurements of the oscillation parameters difficult. In order to be able to tell maximal mixing from $\sin^2 2\theta = 0.9$ it is necessary for a detector to have very good resolution on L/E, but this resolution is severely limited[5] in atmospheric neutrino experiments.

As can be seen in the bottom frame of Figure 1.6, in over half of the phase space the value of L is very dependent upon a precise measurement of the neutrino's $\cos\theta_z$. The angle between the incoming neutrino and outgoing lepton is large at low energies, making measurement of L completely impossible. This is why there is very little slope in the oscillated prediction in the middle top frame of Figure 1.5. On the other hand, if the atmospheric neutrino flux were well known then the these data would be very useful in constraining $\sin^2 2\theta$.

The effect of the smearing of L/E can also be seen well in the "Multi-GeV μ-like" and "PC" frames of Figure 1.5. In these frames the direction of the neutrino and its outgoing muon are correlated enough to tell "up" from "down", so the effect of oscillations is not very large for vertically downgoing events. Below the horizon ($\cos\theta_z < 0$) the predicted number of events under the oscillation hypothesis quickly becomes 1/2 of that predicted for no oscillations.

The final observation the reader should draw from Figure 1.5 is that because of the uncertainty in L/E there is an ambiguity between the neutrino flux and the oscillation parameters. This can be seen especially in the top center frame and the bottom right frame. In the top center frame the predicted flux would not change if the flux was increased while the mixing angle was moved away from maximal appropriately. In the bottom right frame the shape can be kept relatively constant by increasing Δm^2 while moving the mixing away

[5]The Super-Kamiokande collaboration has published an L/E analysis in which they see a dip and rise in the suvival probability at the expected location [29]. However, due to the resolution constraints discussed the minimum occurs at ~0.3 times the no-oscillation hypothesis rather than zero. This inability to resove the true depth reduces the sensitivity to θ_{23} in this analysis compared to the regular Super-K analysis.

from maximal. Because of these sorts of ambiguities the neutrino induced muons, which are at the highest energies and have the best resolution in L, play a key role in constraining Δm^2. The best constraint on the neutrino flux comes from the neutrino induced muons, and knowledge of the flux reduces the ambiguities between flux, $\sin^2 2\theta$ and Δm^2 in the other signals. This means that if SNO data can help constrain the overall flux, the parameter space currently allowed by Super-K could be reduced.

1.4 What Makes SNO a Unique Atmospheric Neutrino Detector

The SNO experiment was designed and optimized to search for evidence of solar neutrino flavor change. The Super-Kamiokande experiment, in contrast, was designed to study both solar neutrino and atmospheric neutrino oscillation. With regards to an atmospheric neutrino measurement, the greatest resultant design difference is the detector volumes. The Super-K detector has roughly 20 times the volume of SNO, and since the rate of fully contained and partially contained events scales with volume the statistics of such events at SNO are far to small to be competitive.

Despite the difference in size between SNO and Super-K, the extreme depth and flat overburden of SNO allow it to make a unique and new contribution to the atmospheric neutrino induced muon analysis. At six kilometers water equivalent (km w.e.) all but the highest energy cosmic ray muons are attenuated as they pass through the rock surrounding the detector; a cosmic muon traveling vertically downward must have a surface energy of at least 3TeV to reach SNO. In addition to being at shallower depths, most large atmospheric neutrino detectors have been located within mountains. Consequently their “slant depth”, the distance to surface along a line at angle θ_z from the vertical is not necessarily azimuthally symmetric, nor does it increase rapidly until θ_z is very close to 90^o. The depth and flat overburden at the SNO detector site greatly affect the atmospheric muon rate. While Super-Kamiokande observes about 2.2 atmospheric muons per second [30] SNO sees only about 3 per hour. In the context of an atmospheric neutrino induced muon flux measurement the suppression of atmospheric muons is extremely useful, because angular misreconstruction and multiple scattering of cosmic ray muons may result in contamination of the neutrino induced muon signal. To avoid contamination by atmospheric muons, experiments to date have been forced to restrict their analyses of neutrino induced muons to upward going muons, or at most angles just above the horizon. Atmospheric muons range out at about 14km w.e., so at SNO's depth under a flat overburden muons up to a zenith angle of $\cos\theta_z$=0.4 can be assumed to be induced by atmospheric muon neutrinos.

As was shown in Figure 1.6, at angles just below the horizon the path length to the neutrino production point changes very rapidly with $\cos\theta_z$. This makes determination of the oscillation parameter L/E extremely difficult right at the point at which oscillation begins to turn on. Above the horizon the path length to the atmospheric neutrino production point is so short that under the hypothesis of vacuum neutrino oscillation (with Super-K's best fit Δm^2) muon neutrinos will not have had time to oscillate into tau neutrinos. SNO is therefore able to make the first high energy measurement of the unoscillated atmospheric neutrino flux.

The rate of neutrino induced muons is dependent on detector area rather than detector volume, and in this respect SNO is only $\sim$1/5 the size of Super-K, rather than 1/20. The ability of SNO to observe neutrino induced muons in the range $-1 < \cos\theta_z < 0.4$ (rather than just $-1 < \cos\theta_z < 0$) adds 40% more phase space, further reducing the statistical difference. The major impact of SNO, however, will be the observation of neutrino induced muons above $\cos\theta_z = 0$. These events provide qualitatively different information, as the distance to their production point is small enough that they will not yet have had time to oscillate. As was discussed in the last section, the Super-K neutrino induced through-going muon data provides the most stringent constraint on Δm^2. This constraint on Δm^2 allows the low-energy data to more strongly constrain $\sin^2 2\theta$, so the addition of the SNO "above the horizon events" in a global fit should help constrain the allowed parameter space.

1.4.1 Neutrino Telescopes and Neutrino Beams

Over the next decade neutrinos produced at particle accelerators will be used to make precision measurements of $\cos\theta_{23}$ and Δm^2_{23}. In the meantime the km^3 "neutrino telescopes" designed to search for cosmic neutrino sources will be constructed and come online. Atmospheric neutrinos are the primary background to neutrino telescopes, so in order to understand a telescope's sensitivity it is important to understand the atmospheric neutrino flux. In the words of Gonzalez-Garcia, "a precise knowledge of atmospheric neutrino flux is of paramount importance for high energy neutrino telescopes" [31].

If the oscillation parameters are measured precisely at an accelerator neutrino experiment, the Super-K data could be used to measure the atmospheric neutrino flux. In [31] Gonzalez-Garcia *et al.* performed this analysis using the Super-K best fit oscillation parameters. A measurement of the atmospheric neutrino flux at SNO, due to the fact that it includes the new "above the horizon" data, would help to reduce the necessary circularity of [31].

1.5 Organization of This Thesis

The organization of this thesis is as follows. This introductory chapter briefly discussed neutrino masses, oscillations, the production and detection of atmospheric neutrinos, cosmic ray muons, and neutrino telescopes. Chapter 2 provides a brief introduction to the Sudbury Neutrino Observatory (SNO). Chapter 3 discusses the electronics channel-level calibrations necessary to perform atmospheric neutrino analyses. Chapter 4 concerns the reconstruction of muons which pass all the way through the SNO detector. Chapter 5 deals with the separation of through-going muon signals from instrumental backgrounds and physics backgrounds. In Chapter 6, we consider the effects that systematic errors in our model of the detector could have upon reconstruction and data selection. A set of physics measurements and conclusions are presented in Chapter 7. A glossary of terms is provided as a reference in Appendix A. A discussion of possible future improvements to the analysis presented in this thesis is given in Appendix B.

1.5.1 Point of Clarification Regarding Terminology

The new measurement that is presented in this thesis requires a change in the standard terminology used by atmospheric neutrino experiments. In the past, the term "upgoing muon" has been used to describe muons that originate from neutrino interactions in the rock. The SNO detector is the first detector able to cleanly separate neutrino induced muons from cosmic ray muons for a large fraction of zenith angle above $\cos\theta_z = 0$. As these "above the horizon" muons are traveling downward the term "upgoing muon" would be misleading. In this thesis the term "neutrino induced muon" will always refer to muons that are produced outside of the SNO photomultiplier tube array (typically in the rock). Neutrino interactions which occur inside of the photomultiplier tube support structure (PSUP) will be referred to as "internal neutrino interactions" or "neutrino interactions inside the PSUP"; these events are the union of the "fully contained" and "partially contained" event classes of Super-Kamiokande.

Chapter 2

The Sudbury Neutrino Observatory

The previous chapter discussed the theory behind neutrino oscillation, cosmic rays, and atmospheric neutrinos. The goal of this chapter is to provide the reader with a description of the detector sufficient to allow understanding of the analyses in the chapters that follow. The general properties of the detector will be introduced first, followed by an introduction to the front end and trigger electronics. A brief description of calibrations on which the thesis analyses depend is presented next, and the chapter finishes with an explanation of the different "phases" of the experiment.

The SNO detector has previously been described in detail in [32], and also by the graduate students responsible for its construction. Readers interested in a detailed description of the detector are referred to the cited paper and the early SNO PhD theses. In particular, readers interested in the SNO trigger and front-end electronics are referred to [33] and [34] respectively.

2.1 Detector Physical Geometry

The Sudbury Neutrino Observatory (SNO) detector is located in a 30 meter high barrel-shaped cavity on the 6800ft level of INCO's Creighton Mine near Sudbury, Ontario. SNO is a water Cherenkov detector, and was designed to search for evidence of the flavor transformation of neutrinos produced in the sun. The solar neutrino target is a kiloton of heavy water (D_2O) contained within a 12 meter diameter acrylic vessel (AV). The otherwise full rotational symmetry of the detector is broken at the top of the detector, where the detector "neck" allows access to the D_2O volume. The Cherenkov light produced by relativistic charged particles is imaged by 9438 photomultiplier tubes (PMTs) arranged on the Photomultiplier Support structure (PSUP), a geodesic sphere 17.8 meters in diameter. The PSUP and AV are immersed in ultra-pure light water (H_2O), which acts as a radioactive shield and provides buoyant support for the filled AV. Distributed on the outer surface of the PSUP are 91 outward looking (OWL) PMTs which are used by the solar analysis as a tag for cosmic ray muons. During the detector commissioning a pair of PMTs known as "neck tubes" were installed in the detector neck to veto light originating from discharges of charge that accumulated on insulators in the neck region. A representation of the SNO detector is shown in Figure 2.1.

Measured from the center of the detector the total photocathode coverage is 31%. Light concentrators which surround the PMT increase this coverage to 59%. These concentrators collect light that originates up to seven meters from the center of the detector and block light from larger radii. The purpose of the blocking is to reduce sensitivity to radioactive background events in adjacent PMTs and the light water. This "concentrator cut off" is unintentionally responsible for causing muon events in the light water to have a characteristic geometry. This event geometry will be shown in chapter 4.

The detector is located 1.3 kilometers from the vertical mine shaft. All detector com-

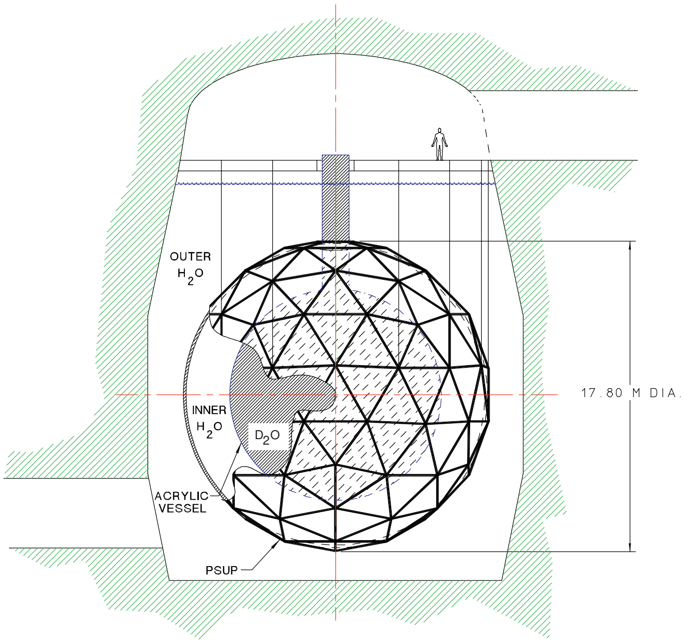

Figure 2.1: The SNO detector is shown in this diagram, which was taken from [32]. The detector "neck" allows access to the D_2O volume for calibrations.

ponents were lowered down this access shaft during the construction of SNO [35]. As can be seen in Figure 2.2, the SNO site is about 70 meters into the mine's "hanging wall" of norite rock (ρ=2.85$\pm$0.02g/cm^3). The hanging wall rests on a "foot wall" of granite (ρ=2.83$\pm$0.10 g/cm^3) and between the two runs a seam of nickel ore which was actively mined during the course of the experiment. The fault between the hanging wall and foot wall slopes at an angle of about 25 degrees from vertical.

2.2 PMTs and Front End Electronics

The SNO photomultiplier tubes are 20cm diameter Hamamatsu R1408 tubes. The PMTs are run at approximately 2kV, at a gain of about 10^7. The PMT noise rate is measured during each run by means of a pulsed global trigger (PGT). The PGT trigger fires every 200ms in order to sample the detector activity at pseudo-random times. The same activity is presumed to underly events triggered by the other triggers. Over the course of the ex-

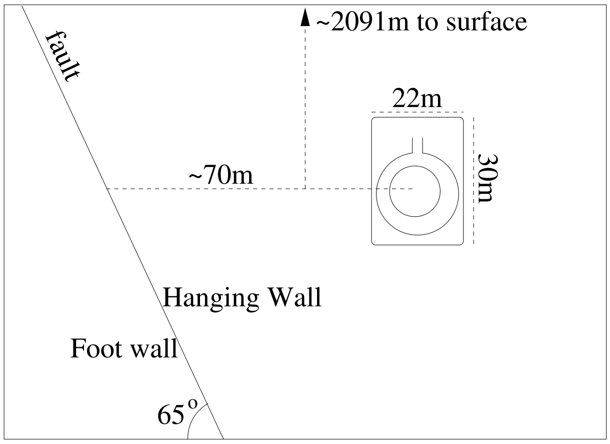

Figure 2.2: This figure shows the position of the SNO cavity relative to the nearest geological feature. The region of ore near where the hanging wall meets the foot wall was actively mined at some points during the course of the experiment.

periment the average PGT event has decreased in size from about 3 tubes hit to about 1.8 tubes hit. The PMT dark firing rate that corresponds to 1.8 tube hits is somewhat less than 500Hz. This exceptionally low rate is partially due to the fact that the PMTs are in a cool (10° Celsius) environment

A plot of the path-corrected PMT hit times (time residuals) from a calibration source known as the "laserball" is shown in Figure 2.3. The laserball source will be described later in this chapter. Briefly, it is an isotropic light source that produces short ($\sim$ 200ps) bursts of monochromatic light. To first order the transit time distribution of SNO PMTs is a Gaussian with RMS of about 1.5ns; these hits make up the "prompt peak". About 5% of PMT firings produce "late pulses", the origin of which is uncertain. The light from the laserball produces reflection peaks which correspond to reflections from the PMT light concentrators and the PMT faces themselves. When the PMTs are illuminated face on roughly 1 in 1500 hits are "prepulses", hits which come about 20ns earlier than the average time of the prompt hits. Prepulse hits are due to photons that make it past the photocathode and eject an electron when they strike the PMT's first dynode.

The SNO front end electronics are housed in 19 "SNO crates". Each of these crates

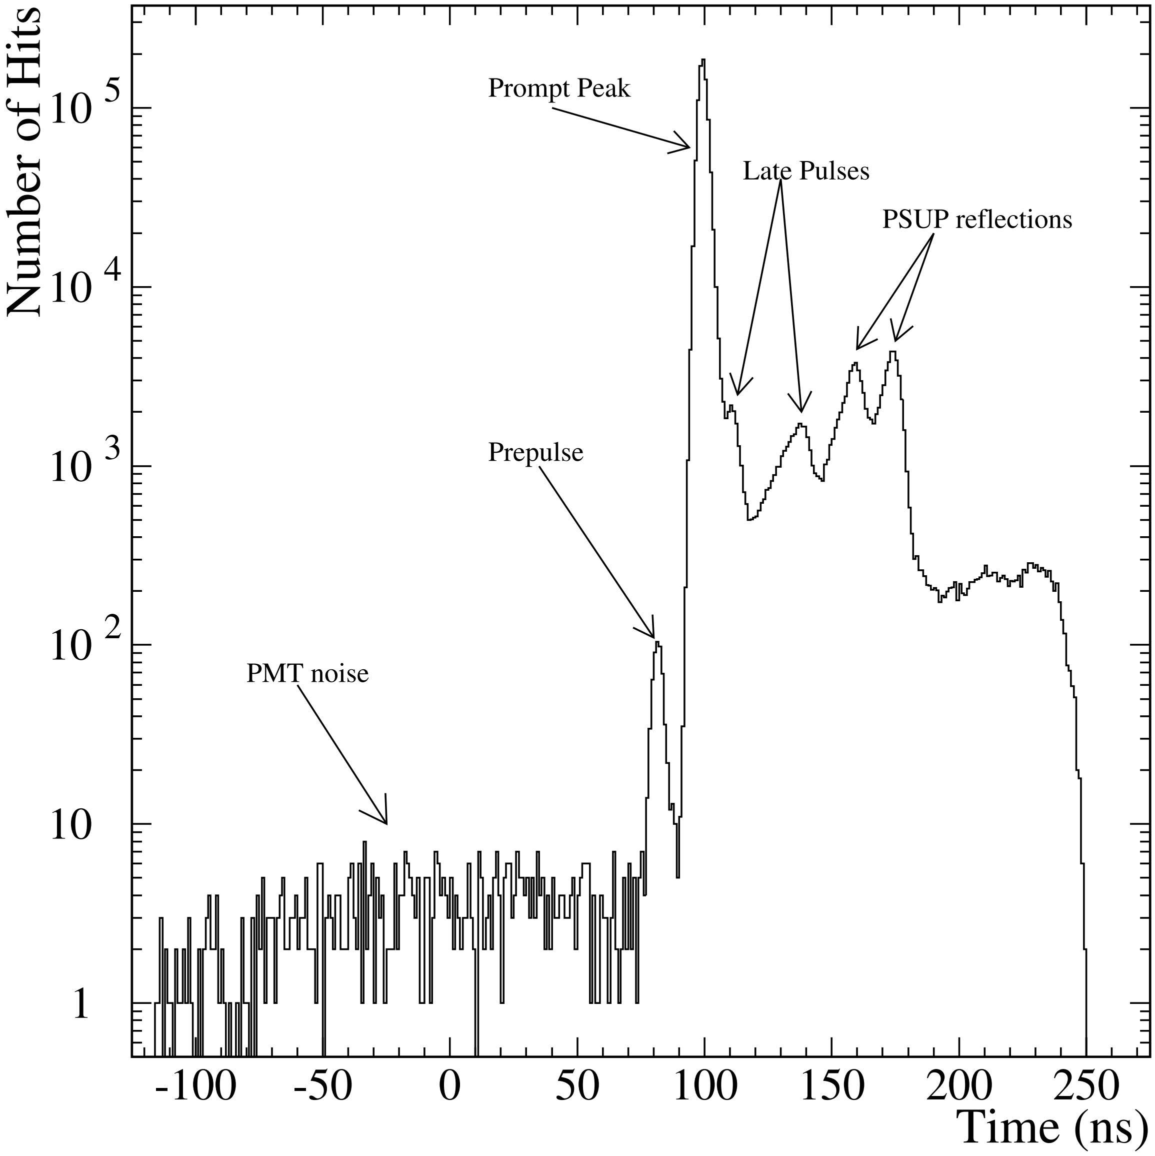

Figure 2.3: Time residual plot for SNO tubes illuminated by a short duration light source positioned near the center of the detector. The presence of detector effects such as Rayleigh scattering and reflections means the time residual distribution shown here is not the same as the intrinsic PMT transit time distribution. The intrinsic transit time distribution was measured in a light box and will be shown later in this thesis as a part of Figure 4.11

is responsible for supplying HV to, and processing signals from 512 channels. Each PMT is associated with a single channel, but because there are more channels than PMTs each crate is responsible for fewer than 512 phototubes. PMTs are connected to the front end electronics with approximately 32m of custom 75Ω RG59-like waterproof coaxial cable. The PMT cables are collected and connected in groups of 32 to a PMT Interface Card (PMTIC)[1], which provides high voltage for the tubes. Each SNO rack[2] has a single high voltage supply, so PMT voltages are customized by individual through-hole feed resistors between the HV source and the PMT. High voltage can therefore be removed from a single tube by physically removing one of these resistors. Each PMTIC also has four high voltage relays to allow the HV to be disabled for groups of 8 channels by the GUI hardware interface program.

Each of the 16 PMTICs in a SNO crate has a PMT signal connection to a corresponding front end card (FEC or motherboard). Four daughterboards are mounted on each motherboard, so each daughterboard is responsible for 8 individual PMT channels. The PMT signals are passed from the PMTIC to a daughterboard through a short connection on the motherboard. Each crate also has a trigger card which communicates with the global trigger system, and digital communication with the crates is handled by a translator card, which is also the bus master for the SNO crate. Digital and analog signals are passed between these boards on a custom backplane.

The daughterboard contains three sets of custom designed chips that are used in signal identification, processing, and storage. Two discriminator chips (SNOD) reside on each daughterboard; each SNOD contains four threshold crossing discriminators to discriminate the individual pulses from four PMTs. The discriminator thresholds are set by an eight bit DAC. The average discriminator threshold at the start of the experiment was 8.6 DAC

[1] The cables are actually connected in groups of 8 to ”paddle cards” which provide a mechanically stable connection to the PMTIC.

[2] The SNO electronics racks contain either one or two SNO crates.

counts above discriminator zero [36], which corresponds to roughly 1/4 of a photoelectron. When a PMT pulse crosses threshold the SNOD sends a logical signal to the daughterboard's pulse integrator, the SNOInt. The PMT signal itself is split before it is passed to the SNOD and SNOInt, and the signal to the SNOInt is delayed. This allows the discriminator to send the SNOInt notification that the PMT has fired before the delayed copy arrives.

The PMT signal goes to the SNOInt, where it is split further into high and low gain copies. The SNOInt integrates these signals on capacitors external to the chip. The charge on these capacitors is made available to an analog storage chip (the QUSN7) at two time intervals after the threshold crossing occurred. This is done to allow separate measurements of the charge at early and late times. There is one QUSN7 chip for each channel, and it is able to buffer analog charge and time data for up to 16 PMT hits without being read out. The QUSN7 has two internal time to amplitude converters (TACs) which are used to convert the time between the discriminator and global trigger signals to an analog charge. In the event that a global trigger has been issued, the QUSN7 stores four analog variables for each PMT hit: The TAC, two high gain charge measurements taken at Short and Long times (QHS, QHL), and a Low gain Long time charge measurement (QLL). The QUSN7 also stores the lowest 16 bits of the global trigger ID (GTID), which is used to build asynchronously read out PMT data into coherent events.

The SNO motherboard contains a field programmable gate array that is responsible for moving analog data from the QUSN7 onto the motherboard for digitization, and the resulting digital information into a digital storage buffer. The analog data from the QUSN7 is digitized by four commercial 12 bit analog to digital converters (ADCs), one for each of QHS, QHL, QLL, and TAC. Once the charge and time data are digitized, three 32 bit words are written to a memory board attached to the motherboard. This on board memory storage is large enough to hold many hours of data taken at normal rates without being read out.

2.3 The Detector Trigger

As was mentioned earlier, the SNO PMTs fire at a rate of several hundred Hertz. This rate is made up of the PMT dark current and detection of Cherenkov photons from low energy radioactive backgrounds in the detector. Below about 4MeV the solar neutrino signal that SNO is primarily interested in is swamped by radioactive backgrounds. In order to reduce the data rate to a reasonable level the SNO detector is designed to trigger on signals with an energy large enough for the event to be considered interesting.

The SNO detector actually has several different triggering modes. The most important mode for the analyses in this thesis is the energy threshold trigger. Below several hundred MeV the number of PMTs hit may be used as a rough estimator of the event energy. The number of PMTs hit and their observed charge are separately summed using analog electronics, and the detector can be made to trigger on either or both of the summed number of hits and the summed charge. As was mentioned earlier, the detector is triggered at a constant 5Hz to sample the background intrinsic to all events. It is also possible for calibration sources to send a message to the detector trigger electronics instructing it to issue a trigger.

In normal "solar neutrino" running the detector is triggered for events with greater than roughly 20 PMTs hit. With this setting the trigger is 100% efficient at detecting events in the D_2O with energies greater than 3.5MeV [37]. The rate of cosmic muons is low enough at SNO that all events with a number of hits greater than this threshold are saved to tape. The deposited energy of events associated with atmospheric neutrinos is orders of magnitude greater than the solar neutrino signal, so for the purposes of the analyses in this thesis the trigger threshold may effectively be ignored.

2.4 Calibration of the SNO Detector

A variety of calibration techniques are required in order to understand the data produced by the SNO detector. Many of the more complex calibrations rely on lower level calibrations in order to be performed at all. For example, each of the calibration variables in the following list depends upon all the calibrations that come before it: the pedestal value of each of the 466,944 analog storage buffers, the charge response and timing offsets of each individual PMT, the attenuation lengths of the H_2O and D_2O at multiple wavelengths, and the energy response of the detector. In addition, since each of these variables may vary with time, the calibration of the detector was an ongoing task during the entire course of the experiment. The amount of physical and mental work that has gone into the calibration of the SNO detector is immense, and it has been described in a large body of internal documents and PhD theses. It is therefore not possible to describe the entire calibration procedure here, so this section will outline the calibrations that are essential to analyses presented later in this thesis. A set of calibrations that were developed by the author and are specific to muon analyses will be described in detail in Chapter 3.

2.4.1 Electronics Calibration

The calibration of the electronics channels requires the measurement of several million variables. As was mentioned earlier, there are 466,944 analog storage buffers, and each of these has its own charge pedestal[3]. The charge pedestals are measured by sending a "pedestal" signal to each channel's discriminator, which forces it to fire. The timing profile of the 19,456 time-to-amplitude converters is measured by firing the pedestals at 31 specific time intervals relative to the global trigger. During the course of the experiment these pedestals and "TAC slopes" were performed twice weekly. The calibration is an automated

[3]The "charge pedestal' is the value measured when no charge is actually observed by the attached phototube.

part of the DAQ system, and the detector operator can run the entire calibration with only a few mouse clicks. It was observed that the value of individual channel pedestals sometimes changed after electronics racks were power cycled, so the electronics calibration was also performed after any planned or unplanned power outage.

2.4.2 PMT Calibration

The SNO PMTs are not all identical, and some efforts were taken to attempt to reduce the effects of this variation. For example, after the supplied high voltage increases past a particular value the single photoelectron detection efficiency and charge response of PMTs increases with increasing PMT high voltage. The voltage at which the detection efficiency begins to increase from zero varies from PMT to PMT by several hundred volts. Before installation the PMTs were coarsely sorted into groups based upon their turn-on voltage. This allows the coarse differences to be taken care of by setting the high voltage supply in each electronics rack to different levels. The HV delivered to each PMT is further individually customized using the HV feed resistors mentioned earlier. This reduces the variation in single photon detection efficiency, but despite this some differences remain.

An additional source of variation in the single photon detection efficiency is the threshold setting of the discriminator attached to each PMT. The amount of "noise" on each channel and its susceptibility to electronic "pick-up"[4] is dependent both upon the individual PMT and the position of it's corresponding channel in the SNO crate. It has been observed that pick-up effects are worse in the channels at the edges of the crate, especially in the front-end card nearest to the digital signal translator board.

The single photon detection efficiencies, time offsets, gains, and discriminator walks[5]

[4] Pick-up is unintentional induction of high frequency signals onto a PMT cable, or onto the traces on the electronics boards.

[5] Threshold-crossing discriminators have a well known timing effect known as the "discriminator walk". Pulses with the same shape but different amplitudes cross threshold at different times, the larger the pulse amplitude the earlier the discriminated time.

are all measured using a calibration source known as the "**laserball**". The "diffuser ball" from which the laserball derives its name is a 10cm diameter quartz sphere that is filled with 50μm glass beads suspended in a silicon gel [38]. A nitrogen laser situated in the clean room on the detector deck produces 337.1nm light pulses[6] that pass through a pair of neutral density filters and into a set of fiber optic cables. The neutral density filters allow the light intensity to be varied over several orders of magnitude. The fiber optic cables direct the light into the laser diffuser ball, where the light scatters from the glass beads. The scattering length inside the ball is roughly 1cm [39].

The laserball source is the primary source used for the calibration of the PMTs, and also for the calibration of the D_2O and H_2O optics. The software that measures the PMT gains and discriminator walks was created by Cameron [40], and is referred to by the abbreviation "PCA" for **P**hototube **CA**libration. The PCA calibration is performed by positioning the laserball near the center of the detector and pulsing it at 45Hz for a little more than an hour. The laserball intensity is adjusted using the neutral density filters to produce events with PMT occupancies in the range of 5-10%. This low value of the occupancy was chosen to reduce the probability of any PMT detecting more than one photoelectron in an individual event.

The discriminator walk for a single PMT is shown in Figure 2.4. The scatterplot at the top of the figure shows the electronics calibrated (ECA) PMT hit time as a function of measured high gain charge in pedestal subtracted ADC counts. The charge distribution is shown in the histogram below the scatterplot. The well understood tendency for the discriminator to fire earlier for large pulses is clearly visible. The blue line represents the discriminator threshold; the hits to the left of it are "crosstalk" hits, in which the discriminator fired on noise induced by a real hit on a nearby (usually adjacent) channel. The discriminator walk for each channel is parameterized using a set of 10 interpolation points

[6] The laser may also be used to pump dye cell units at 365, 386, 420, 500, or 620nm.

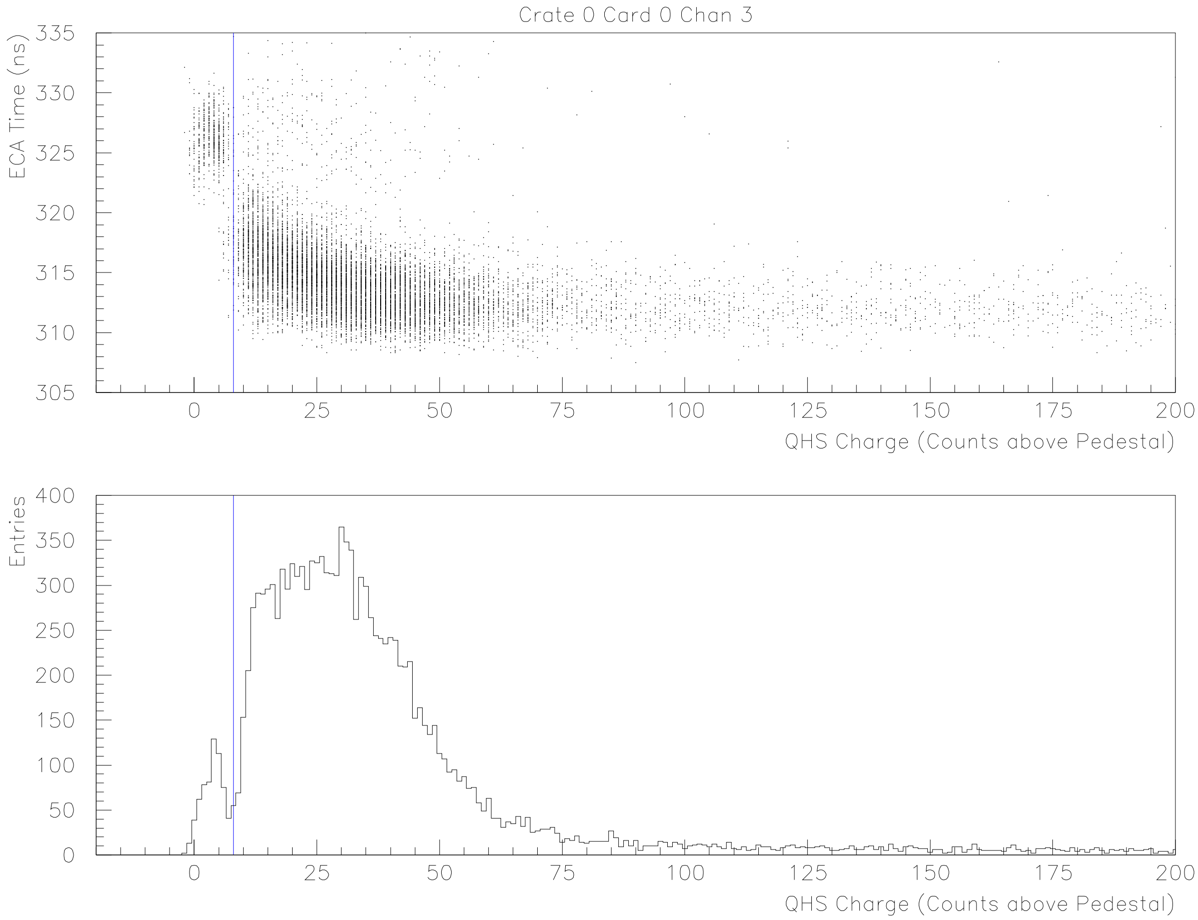

Figure 2.4: The PMT time as a function of charge, and the PMT charge distribution is shown for a single channel. The time has been converted from ADC counts to ns using the electronics calibration, the charge is shown in pedestal subtracted ADC counts. The vertical line shows the discriminator threshold, hits to the left of it are due to crosstalk. Plot was taken from [40].

in charge. The laserball time residual distribution obtained after applying the discriminator walk correction to each of the channels was already shown in Figure 2.3.

It was found that a robust estimate of each PMT's gain could be obtained by measuring the "high half point" of the PMT charge distribution. The high half point is defined as the point at which the charge distribution has fallen to half of its maximum value. This technique for measuring the PMT gains was suggested by Frati, and it was found to be insensitive to the discriminator threshold settings for normal channels. The discriminator

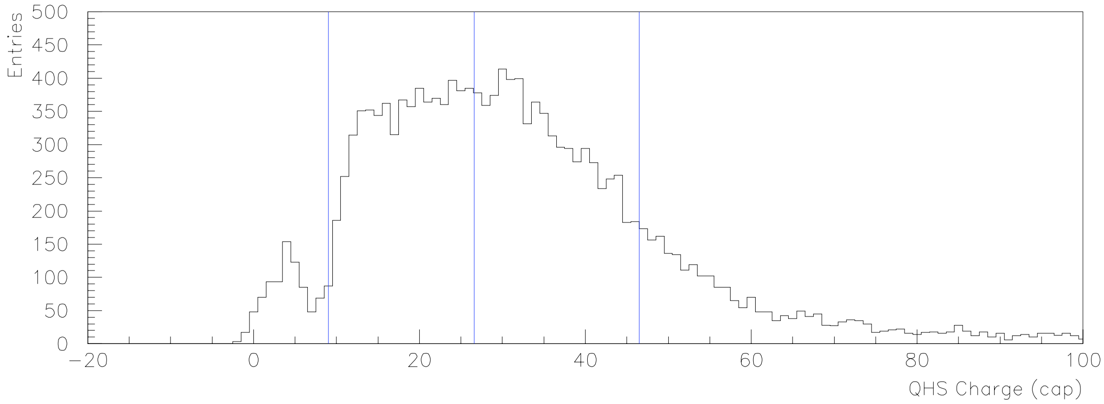

Figure 2.5: The position of the discriminator threshold, peak, and high half point are shown overlaid on the charge distribution for a single channel. Plot was taken from [40].

threshold, peak, and high half point are shown for a single channel in Figure 2.5.

Later chapters in this thesis will make use of "**gain scaled units**" or GSU for short. The charge observed by a phototube is converted from pedestal subtracted ADC counts to gain scaled units by simply dividing the charge in ADC counts by the channel's measured high half point. The channel to channel variation in mean observed charge can vary by up to a factor of two in pedestal subtracted ADC counts, and this variation is mostly removed by using GSU [41]. Studies performed using Klein's simulated charge distribution indicate that one gain scaled unit is roughly equivalent to 1.2 photoelectrons.

2.4.3 Optical Calibration

Measurements of water and acrylic attenuation lengths, Rayleigh scattering, and reflections are also performed using the laserball calibration source. To do these calibrations the laserball is placed in a variety of different positions in the D_2O, thereby varying the distance photons must travel through different media to reach the PMTs. A global fit to all optical parameters is performed using all of the data from a laserball "optics scan". Obtaining the necessary calibration data takes several days of 24 hour calibration running.

2.4.4 Energy Calibration

The main energy calibration source for solar neutrinos is a radioactive ^{16}N source. Nitrogen-16 undergoes a beta decay into ^{16}O, with 74% of decays resulting in an excited state. The excited states of oxygen drop to the ground state through the emission of a 6.13MeV (69%), or 7.12MeV (4.8%) gamma. A commercial deuterium-tritium (DT) generator is used to produce ^{16}N from a (n,p) reaction on ^{16}O in CO_2 gas. The half life of ^{16}N is 7.13s [42], so the radioactive gas must be flowed quickly from the DT generator into the detector. The gas flows through pipes into (and subsequently out of) an 11.43cm diameter cylindrical source canister, which can be positioned at various points in the detector. A plastic scintillator is used to tag β decays that occur inside the chamber, in order to separate them from decays in the tubes leading into and out of the source canister. In order to track variations in the energy scale with time, the energy scale was measured at many points during each of the phases.

Neutrino induced muons have energies in the GeV range, but all of the SNO energy calibration sources have energies of a few MeV. In the past, atmospheric neutrino experiments that are not as deep as SNO have used stopping cosmic muons and electrons from the decay of stopped muons as calibration sources. Due to the depth of SNO, however, cosmic sources are not abundant enough to produce the statistics necessary for an energy scale measurement. In chapter 5 we will see that the SNO simulation, which is calibrated to match ^{16}N, does a good job at estimating the light output of muon events.

2.5 Phases of the SNO Experiment

The SNO detector was built and optimized to study solar neutrino flavor change. The D_2O in the SNO detector makes the detector sensitive to the neutrino-deuteron neutral-current interaction, in which a neutron is released from the deuterium nucleus. The sensitivity of the detector to solar neutrino flavor change depends upon its ability to separate electrons produced in the neutrino-deuteron charged-current (CC) interaction from neutrons produced in the neutral-current interaction. The experiment was run in three different phases, each with different techniques for detecting neutrons.

The first phase of the SNO experiment is known as the "D_2O phase". In the D_2O phase neutrons were detected through their capture on deuterium, which results in the production of a 6.25MeV gamma. The distribution of estimated energies of events was used to estimate the number of charged current and neutron capture events. The D_2O phase ran from November 2, 1999 to May 31, 2001.

In the second phase of the experiment, two tonnes of NaCl salt was added to the D_2O and used for neutron detection. The cross section of neutrons on Cl is 0.43barns, which is three orders of magnitude larger than the cross section on D_2O (5.2×10^{-4}barns). For this reason the addition of the comparatively small amount of salt increased the neutron detection efficiency. In addition to the increased capture efficiency the capture of neutrons on Cl results in the production of multiple gammas, with a total energy of 8.6MeV. The presence of multiple gammas allows the use of isotropy parameters to further distinguish neutron capture events from CC electron events. The "salt phase" of the experiment ran from July 26, 2001 to August 29, 2003.

After the salt phase was completed, data was taken for roughly one month to verify that observed optical changes were reversed with the removal of the salt. After the short "$D_2O - 2$" phase, the detector was turned off for an electronics upgrade and the installation of proportional counters. In the "NCD phase" of the experiment, neutrons are detected by their capture on ^{3}He in wire proportional counters suspended in the D_2O. The NCD phase started production data taking on January 3, 2005, and is expected to be completed on November 20, 2006. Upon completion, the D_2O will be removed and returned to Atomic Energy of Canada Limited.

The data used for the neutrino induced flux analysis in this thesis are from the D_2O phase only. Some data from the salt phase are used for calibration and estimation of instrumental background contamination. It is anticipated that the results of this thesis will be used by the SNO collaboration to perform a three phase analysis. The final live time of the detector for muon analyses is expected to be close to 1000 days, which will increase the available statistics by slightly more than a factor of three from those used here.

2.6 Reminder

With the introduction of the SNO detector complete we will now begin to focus on the analyses necessary to measure the neutrino induced muon flux. The following chapter will deal with individual channel level calibrations. Readers are reminded that a guide to the thesis was presented on page 28, and that if the meaning of a term cannot be recalled it may be helpful to check the glossary, which starts on page 225.

Chapter 3

Channel Level Calibrations

The previous chapter introduced the reader to the SNO detector. To make physics measurement using the detector it is necessary both to understand the data the detector produces and to create a robust detector Monte Carlo model. To do this the SNO collaboration employs a variety of calibrations and Monte Carlo studies. These calibration measurements differ greatly in complexity, from simple automated routines which choose electronics hardware settings to measurements of the light transmission properties of the water. This chapter deals with the electronics channel level calibrations which are necessary to do a through-going muon analysis. Readers interested in detailed calibration information should see: For the calibration of the phototubes [40], for the electronics calibration [43], for the optical calibration [39], for the solar neutrino energy calibration [44] and [45], for the collaboration analysis program SNOMAN[1] [46], and for the handling of geometry in SNOMAN [47].

[1] SNO Monte Carlo and ANalysis

3.1 Introduction

In order for high level analysis tasks such as the reconstruction of events inside of the SNO detector to be successful it is necessary to calibrate the PMT information read out by the front-end electronics. The measured charges must have their baseline digitization (pedestals) subtracted, and the measured times must be converted into nanoseconds relative to global trigger. The charge pedestals are measured by finding the median digitized ADC value when the discriminator has been forced to fire, and the relationship between ADC counts and nanoseconds is measured by forcing the discriminator to fire at varying times relative to the global trigger. A complete description of the generation and use of electronics calibration constants is given in [43].

The next step after the electronics calibration is the phototube calibration (PCA). The time offset between when a tube is hit and when the discriminator fired is measured by suspending the laserball in the center of the detector. The laserball illuminates the PMTs at approximately single photoelectron levels (about 8% occupancy) with a narrow pulse of light at a single wavelength. The detector is triggered at a constant known time after the laserball fires, and this allows the channel time offsets to be measured. The SNOD is a threshold crossing discriminator, and thus the threshold crossing time is dependent upon the amplitude of the PMT pulse. This effect is measured at the same time as the channel timing offsets, and a walk correction based upon the PMT charge is applied when the tubes are calibrated. Finally, the charge response (gain) is not the same for all PMTs, so the PCA is responsible for characterizing the charge distribution of each tube.

As the detector ages individual PMTs and electronics channels occasionally fail. Phototube channels that develop a failure at their wet end are not recoverable and are disabled, although their associated electronics channels are usually left active for use in crosstalk studies. Front-end electronics failures are repairable, but there is usually a period after a front-end failure during which the channels are temporarily disabled. Due to these failures the number of online channels has continuously varied over the course of the experiment, usually declining by a few hundred before repairs can be carried out, and with a steady downward trend as wet-end failures accrue. While offline channels are tracked by the data acquisition (DAQ) system a set of analysis tests to determine which channels are working properly has been devised. This has the advantage that it can identify and flag failures that occur during runs, it identifies offline channels even if they have not been set in the DAQ due to operator error, and it also identifies channels with abnormal behavior. These analysis online tube (ANxx) checks are described in 3.1.1 below.

The calibrations described above are sufficient for analyses in which the PMTs are illuminated at single photo-electron light levels. As the light level increases these calibrations begin to break down because the single-photon hypothesis which they are founded upon is no longer valid. This chapter discusses three measurements and calibrations that are important for muon analyses. The first (3.1.2) is the calibration of the low gain charge measurement (QLL), the second (3.1.3) is the Monte Carlo PMT transit time distribution, and the last (3.1.4, 3.1.5) is the correction of the multiphoton timing effect[2] in Monte Carlo and data respectively.

3.1.1 Identification of Analysis Online Channels

There are many failure mechanisms that may render an electronics channel temporarily or permanently unusable, and channels may undergo gradual degradation or sudden spectacular death. One of the most common failures is believed to be high voltage breakdown due to water leakage into the PMT wet end connectors. When this occurs it is common for one or more of the feed resistors on the board which supplies high voltage to burn out. Occasionally only the feed resistor for the individual PMT burns out, but sometimes all

[2] The tendency for tubes hit by multiple photons to fire earlier than tubes hit by single photons.

eight resistors supplying PMTs on the same "paddle card" also burn. It is also common for the main feed resistor for all 32 channels on the high voltage card supplying the bad tube to fail. If the detector operator fails to start a new run after one of these events, then for some amount of time the DAQ accounting of the number of online tubes will be incorrect. The "ANxx banks" are a set of run-specific lists of problem channels which have been developed in order identify channels that are offline, or are otherwise dead to the analysis.

Two sets of ANxx banks are created, one for neutrino and calibration runs, and another for the phototube calibration (PCA) runs. Both types of ANxx banks contain information about whether the tubes are online or not. The PCA ANxx banks also check that the distributions of charge and time for single-pe laserball events are normal. The ANxx checks that are common to both PCA and neutrino are the following:

- **Zero Occupancy:** Channel does not fire during run.
- **Low Occupancy:** Channel has much lower than normal occupancy for NHit triggered events in at least one integration period during the run. The integration period is 30 minutes for neutrino runs. This check identifies channels that fail part-way through a run.
- **High Occupancy:** Channel fires in more than 1% of Pulsed Global Trigger events.
- **Sync Clear:** Channel GTID did not roll over at the proper time, so the channel will appear in the wrong events.

The checks that are specific to PCA runs are:

- **TAC:** Large fraction of channel hits fall outside of a 20ns time window relative to the laserball firing.
- **QHS:** Large fraction of channel hits fall outside of the range 5 to 80 ADC counts above pedestal.

- **QLL:** Large fraction of channel hits fall outside of the range -5 to 20 ADC counts above pedestal.

- **Gain Ratio:** Ratio of high gain to low gain charge measurement for individual hits does not fall within typical range.

- **Hardware:** In a small number of cases a problem with the hardware configuration of a set of tubes was known, this bit allows those tubes to be removed by hand.

- **Former T/Q Failure:** Set if the channel failed the time, charge, or gain ratio cut when the last PCA run was taken. Each type of former failure is flagged individually.

Finally, one check is specific to neutrino runs:

- **Prompt Occupancy:** Low NHit data are reconstructed and time residuals found for each tube in each event. The occupancy of each tube inside of a time window of $\pm$10ns from the fit time is determined, and events with a small "prompt occupancy" are removed.

The "Former Failure" tests are used to check that the front end electronics was working properly in both the PCA runs immediately prior to and after the neutrino or calibration run in question. In this way if a channel is fixed prior to a PCA run being taken it will still be flagged as bad for the entire period. The PCA runs are taken at a monthly frequency.

Individual analyses are able to select which of the ANxx checks they will apply. The solar neutrino analyses, for example, make no use of the low gain charge measurement so the QLL and Gain Ratio tests are not applied. The number of channels used by the muon analysis is shown as a function of run number in Figure 3.1.

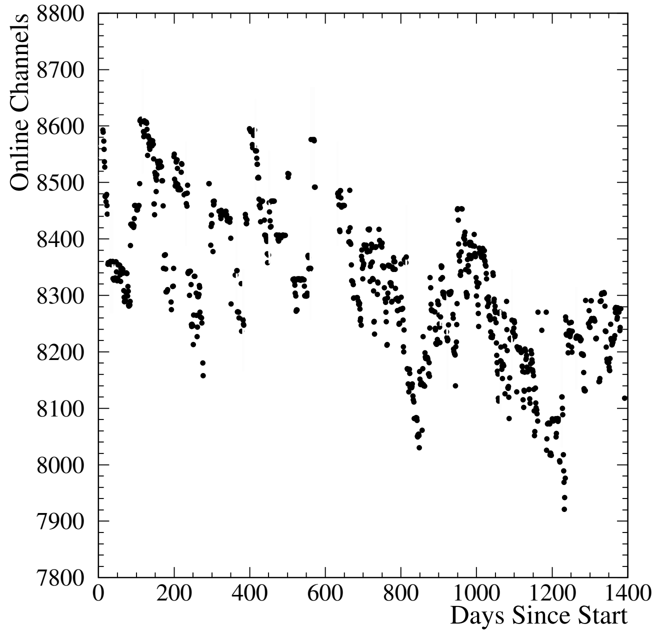

Figure 3.1: The number of channels used in the muon analysis is shown against number of days since November 2, 1999. The zero has been suppressed in order to exaggerate the declining number of channels. The runs used in the salt phase analysis begin at day 633.

3.1.2 Calibration of the Low Gain Charge Measurement (QLL)

In section 2.2 it was mentioned that the SNO electronics perform two charge measurements for each PMT signal. These are QHS, a high gain ($\sim$ 40 DAC counts/pe) short time span measurement and QLX, a low gain ($\sim$ 2 DAC counts/pe) long time span measurement. The SNO experiment was designed with this feature because the Kamiokande-II experiment had observed events which had saturated the electronics [48]. Despite this ability to measure charge over a large dynamic range the charge deposited in an event has been mostly unused in solar neutrino analyses because only a few percent of the PMTs are hit by multiple photons in low energy events. The charge resolution of the SNO phototubes is not fine enough to differentiate single photoelectron hits from double photoelectron hits. For low energy events a better energy resolution may be obtained by using the number of tubes hit rather than the sum of the PMT charges. The charge is used in instrumental background rejection routines (the phototubes can resolve the difference between zero or one photon well), and in the discriminator walk correction mentioned earlier. Neither of those analyses make use of the low gain charge measurement (QLL), so until the development of muon analyses there was no compelling reason to calibrate the QLL measurement.

The QLL measurement of charge differs from the commonly used QHS measurement in two ways. The QLL signal is about an order of magnitude smaller than the QHS signal, and the charge measurement is taken at a much later time after the discriminator fires ($\sim$300ns) than QHS is (63ns). It is only necessary to use the QLL measurement when the charge exceeds the dynamic range of QHS. This happens for only a few percent of the tubes in a typical muon event, but the ability to use QLL is necessary in order to extend the multiphoton timing correction discussed in 3.1.5 up to and slightly past the average QHS railing point.

It was found that the QHS to QLL ratio for every channel in the detector varied at about the 10% level, but that the variation in ratio for channels with the same daughterboard channel number (1-8) is only about 4%. The ratio was found to increase slightly with increasing daughterboard position from the bottom of the crate, so a calibration based upon the motherboard channel number (0-31) was pursued. The ratio of pedestal subtracted QHS to QLL was measured for each channel in the detector using about three live weeks of muon data, one week from the beginning of the D_2O phase of the experiment, another from the end of the D_2O phase, and the third from the end of the salt phase. Data were accumulated for each electronics channel and then fit to a first order polynomial using linear regression. The data from channels with at least ten hits with a measured QLL of 50-200 DAC counts above pedestal were combined to find the mean QHS to QLL ratio for each motherboard channel number. A more complete description of the procedure is given in [49].

The new QLL gain calibration was tested using a different data set from the one used in its creation. The difference between the high and low gain calibrated charges (PHS and PLX respectively) was tested using large charge laserball data. The difference between the two charges is shown against the average of the two in Figure 3.2. The error bars in the figure represent the RMS spread of the data in each bin. Better agreement is achieved at all charges by using the new motherboard channel number dependent PLX gain calibration rather than the original gain calibration which assumed a fixed QHS to PLX ratio.

3.1.3 Single Photon Transit Time Distribution

There is an interval known as the "transit time" between the time the photo-cathode of a photo-multiplier tube (PMT) emits a photo-electron, and the time at which the resulting pulse leaves the PMT base. For SNO tubes this time is on the order of 30ns, of which about 10ns is spent during amplification inside the dynode structure. The distance from the photo-cathode to the dynode along an electric field line, and hence the transit time, is not the same for each point on the surface of the tube. This variable distance causes a "transit time spread" that is approximately Gaussian distributed with an RMS near 1.5ns for SNO PMTs.

A plot of time residuals for the laserball in the center of the SNO detector was shown in Figure 2.3; its features are described below. Most photo-electrons that make it to the dynode produce pulses within a few σ of the average transit time; these hits make up the "prompt peak". About one in twenty detected photons in SNO produce "late pulses". The origin of the two late pulsing peaks in SNO tubes is uncertain. Approximately one in 1500 PMT hits in SNO is due to "prepulsing". The prepulse peak is due to photons that strike the first dynode after passing through the photo-cathode without being absorbed. Such a photon ejects an electron from the dynode, and the electron then cascades through the dynode chain creating a hit that is earlier than normal hits by the interval of the cathode to first dynode electron drift time. Photomultiplier tubes which are kept in absolute darkness still fire at a rate of a few hundred Hz, and this "dark current" PMT noise is due to electrons becoming thermally excited and ejected from an early dynode. These electrons may then be attracted

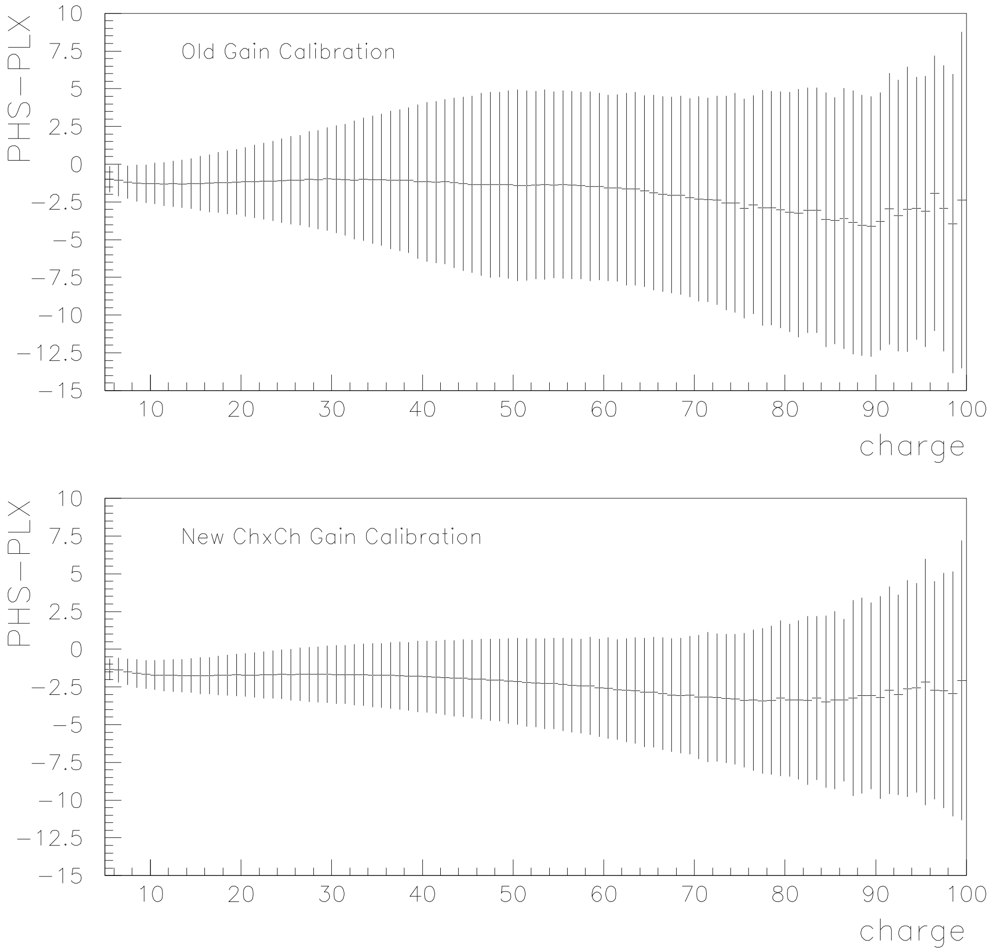

Figure 3.2: Profile histograms of the difference between PHS and PLX for the original (top) and the new (bottom) gain calibrations. The error bars represent the RMS of the data in each bin. The horizontal "charge" axis is the average of the PHS and PLX charges. The new gain calibration has a smaller spread in every bin, and the mean value has less variation with charge.

to the lower potential of the next dynode and initiate a cascade. Finally, as electrons shower in great numbers near the base of the PMT it is possible for diffuse gas inside the PMT to become ionized. These ions are accelerated towards the photo-cathode may eject electrons when they hit it. Since ions are considerably more massive than electrons they take a long time (milliseconds) to transit the tube, and consequently “after-pulsing” is not an important effect when studying the PMT transit time distribution.

The transit time distribution for SNO phototubes is not measured *in situ* due to detector effects such as Rayleigh scattering and reflections. The transit time distribution was measured at Queens University with approximately 20 phototubes inside a dark box. The goal of the experimenters was to measure the transit time distribution at late times, and unfortunately their data does not contain enough precision to see the prepulse accurately. Since prepulses occur in only 1 out of every 1500 direct hits this effect is not particularly important for the low energy solar neutrino analysis. It is an important effect for through-going muons, however, as many tubes are hit by tens or hundreds of pe and the discriminator will usually be fired by the first cascade to reach it. It is also essential to have a realistic model of the transit time distribution when measuring the multiphoton timing effect for Monte Carlo.

The detector effects which make measurement of the transit time distribution at late times difficult are minimal at early times. With this in mind an improved transit time distribution was created by combining SNO laserball data at early times with the Queens transit time distribution at late times. A set of tubes were selected for the analysis based upon having normal occupancy in laserball runs, a flat dark current timing distribution and normal dark current rate. The components of the new transit time distribution are summarized below, a more detailed description is given in [50].

- **Prompt and late times** - The Queens transit time distribution is used.

- **Times between prepulse and prompt peaks** - The Queens transit time distribution is extrapolated to early times with an exponential fit to the rising edge of the prompt peak. The fit is normalized to the Queens data.

- **Prepulse** - Laserball PCA data is used. The contribution of the laserball time spread is assumed to be Gaussian and is measured by comparing the width of the PCA prompt peak to the width of the Queens transit time distribution. The prepulse peak is parameterized with a Gaussian which is then deconvolved by the measured laserball time spread.

- **Prepulse to prompt normalization** - Measured with laserball PCA data. The PMT dark current is fit with the data prior to the prompt peak and removed when calculating the prepulse to prompt ratio.

- **Overall normalization** - After extrapolating the Queens data to early times and patching the prepulse onto it the resulting distribution is not normalized. Since this is intended to be a probability distribution the integral from -23.9ns to 68.1ns is normalized to one.

Many Monte Carlo tests are performed by simulating “photon bombs” rather than charged particles. Photon bombs can be customized to have any position, direction, energy, time, and multiplicity that the user desires. A Monte Carlo simulation of photon bombs at the center of the detector, with isotropic direction, a time spread typical of the laserball, and multiplicities that will lead to almost entirely single photoelectron hits was used to compare the old and new transit time distributions. A comparison of the two is shown in Figure 3.3 with the measured laserball time distribution overlaid. The new transit time distribution is clearly in much better agreement with the data than the old.

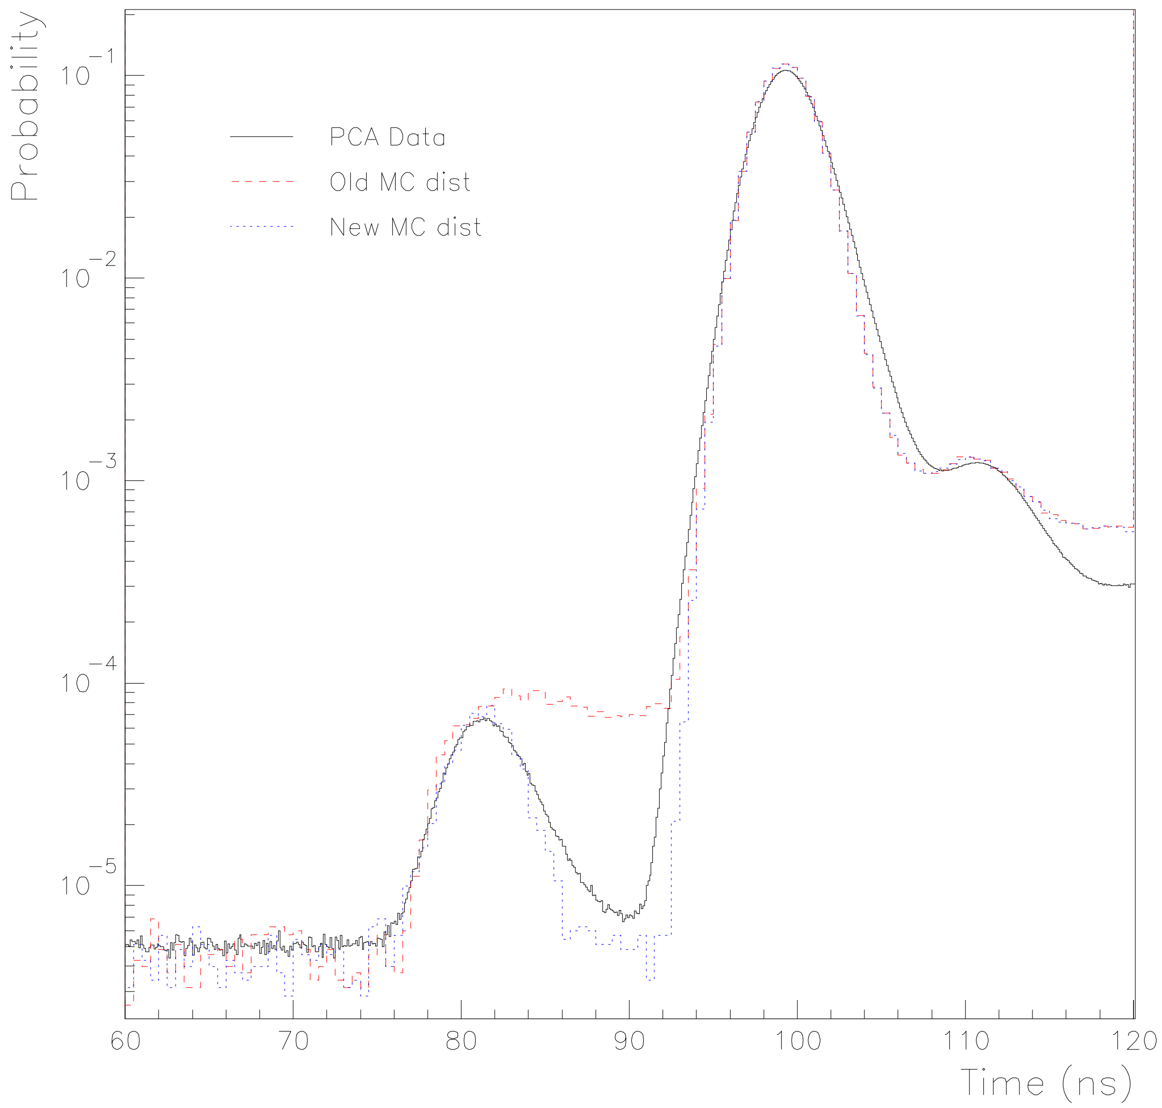

Figure 3.3: Comparison of MC photon bombs to laserball PCA data with the new and old transit time distributions. The PCA data includes all tubes, not just the tubes which were used in deriving the new transit time distribution. A Gaussian smear with an RMS of 0.84ns has been added to the MC to mimic the laserball timing spread.

3.1.4 Multiple Photoelectron Timing Correction for Monte Carlo

If two photons strike the photocathode of a PMT and liberate electrons simultaneously the pulses they produce may not reach the discriminator at the same time due to the transit time spread discussed in 3.1.3. If both of the pulses exceed the discriminator threshold it is the first of the two which will fire the discriminator. Photomultiplier tubes illuminated by large numbers of photons will thus have earlier times on average than tubes illuminated with single photons. The larger the number of photons which strike the photocathode the earlier the time will be shifted. In addition, as the number of photons increases the spread in the time at which the discriminator fires (the "first photon time spread", analogous to the single photon "transit time spread") becomes smaller. This multiphoton timing effect is important for muon analyses, as tubes that are near the point where the muon exits the detector will detect hundreds of photoelectrons and thus appear earlier than tubes that are far away and only detect several photoelectrons.

Due to the poor charge resolution of the SNO phototubes it is not possible to resolve the exact number of detected photoelectrons. In an idealized experiment laserball timing calibration data would be taken at varying intensities from single photon to hundreds of photons periodically throughout the course of the experiment. The primary goal of SNO, however, is to study solar neutrino events for which few tubes are hit by multiple photons. The necessary multiphoton calibration data therefore does not exist. In an idealized analysis the multiphoton timing effect would be considered during reconstruction, and a maximum likelihood fitter would apply different timing corrections based upon the estimate of the number of photoelectrons observed by the PMT. This technique was not pursued largely because at the time the correction was developed it was not expected that such an advanced reconstruction routine would be written.

As the number of photoelectrons increases the correlation between measured charge and number of photoelectrons also increases due to the central limit theorem. This makes it possible to create a multiple photon timing correction based upon the charge measured by the phototube. While such a correction will often slightly over or under estimate the timing shift due to Poisson variations in the number of detected photoelectrons it will on average get it right, and since the first photon time spread is smaller than the 1.6ns transit time spread even an imperfect correction will lead to decreasing time spreads with increasing charge.

Estimating the Multiple Photoelectron Time Shift and First Photon Time Spread

It is possible to estimate the multiple photoelectron time shift as a function of charge by using the SNO PMT transit time and charge distributions. The simplest way of doing this would be to sample the time distribution n times and record the earliest time, then sample the charge distribution n times and sum the charges. By repeating this process many times a table of the charge and time shifts as functions of number of photoelectrons could be created. This straightforward method is not pursued for two reasons. First, the photons that illuminate PMTs from, say, a muon event, will not actually arrive simultaneously. Second, the single photon charge response in both Monte Carlo and data varies significantly from PMT to PMT. In order to get around this problem the full SNOMAN Monte Carlo is used in creating the correction.

The timing correction is measured with simulated events that are somewhat similar to laserball calibration data. A "photon bomb" is created by generating a large number of identical wavelength photons simultaneously at the center of the detector isotropically directed radially outward. The photons then propagate to the PMTs and are allowed to be absorbed, reflected, and scattered as in the regular SNO Monte Carlo. If a phototube is hit and records a time that is inside the prompt peak of the single photon transit time distribution then its data is recorded (prepulse hits are discarded). The hit time and time squared are summed in several charge bins in order to measure a first photon mean time and RMS for the data in that charge bin.

The photon bomb events are simulated with a variable number of photons sampled from a flat distribution. A flat distribution was chosen in order to measure the multiple photon time shift over a large range of charge without bias. The mean charge for each event is recorded, and in order to avoid a bias occurring at large charges the timing correction is only recorded for charge bins below the largest mean event charge. The final wrinkle is that because of SNOMAN memory limitations there is a limit to how many photons may be generated. In order to avoid this problem the photon bombs are directed radially outward in 20^o cones, rather than over the full 4π solid angle. The central direction of these photon cones is distributed over the whole solid angle in order to sample all of the PMTs equally, and only PMTs within 15^o of the central cone direction are included in the analysis.

Applying the Correction to Monte Carlo

The multiphoton timing correction is applied shortly after the generation phase in the Monte Carlo. Once the charge and hit time of each hit PMT is available a multiphoton time is created by adding the mean multiphoton time offset for that particular charge to the tube time. The effect of applying the correction can be seen in Figure 3.4. The figure shows profile histograms of PMT time residuals for 100 Monte Carlo cosmic muons with (solid) and without (dashed) the correction applied. As can be seen in the plot, the multiple photon effect has been removed to well within the first photon time spread.

3.1.5 Multiple Photoelectron Timing Correction for Data

The strategy behind the multiple photoelectron timing correction for data is similar to that for the simulation, but several complications require that the measurement be performed somewhat differently. The first problem is that the discriminator walk correction is measured at low intensities, and cannot be reliably applied for tubes with large charges. This

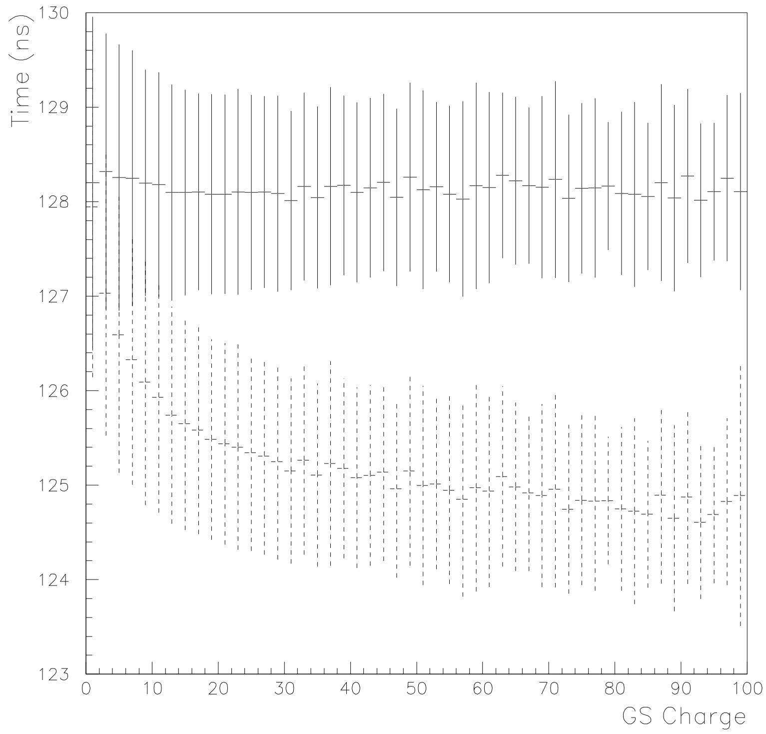

Figure 3.4: Profile histogram showing the time residuals for PMTs hit in 100 simulated cosmic muons as a function of PMT gain scaled charge. The solid lines represent data where the multiphoton timing correction has been applied, dashed lines show data without the correction. The error bars in each case show the RMS of the PMT time residuals at that charge. Note that the actual value of the ordinate to which the time residuals are corrected is arbitrary.

means that the multiple photon correction for data must correct both the multiple photon timing effect and the discriminator walk effect. Since multiple photon calibration data was not taken from the start of the SNO experiment it is not possible to pursue a channel by channel correction (as the front end electronics which reads out the channels may have changed). For this reason a global correction must be pursued.

Another complication is that SNO does not possess a prompt, isotropic light source with a continuously variable intensity. The laserball used for phototube timing calibrations has often had a large up-down asymmetry, and in single photon calibration runs there has been up to a factor of two between the occupancy of tubes at the top and bottom of the detector. The intensity of light emitted by the laserball varies by about a factor of two from event to event, and the average event intensity tends to decrease as the laserball is run. The intensity is controlled by using a set of neutral density filters to attenuate the laser beam, and these filters unfortunately only allow for coarse changes in intensity. As will be seen below, all of these "problems" may be viewed as "opportunities", as they make it easy to cover a large range of intensity with a relatively small number of runs.

The up-down asymmetry of the laserball is exploited by dividing the PMTs into ten bins in which all of the PMTs have similar occupancy in single photon laserball calibration runs. The event to event variability in the intensity of the laserball is exploited by finding the mean and RMS of the PMT gain scaled charges for each bin of PMTs in each event. Only tubes with charges within 0.4σ of the bin mean charge are used in measuring the multiphoton timing effect. The multiphoton timing correction was measured before the ANxx channel checks were developed, so "analysis online" channels were selected using a list of channels provided by the solar neutrino energy calibration group. Channels were further screened for normal occupancy, and ten channels were removed by hand when it was discovered that they had QHS/PLX gain ratios that were very different from normal channels. The runs used in creating the multiphoton timing correction for data are listed in

Run	Rate	Length(min)	$< NHIT >$	$< QHS >$
29810	$25Hz$	100	900	640
29812	$10Hz$	40	1900	650
29814	$4Hz$	10	4000	650
29816	$4Hz$	1	4300	650
29818	$2Hz$	8	5500	660
29820	$2Hz$	5	5800	660
29822	$2Hz$	8	8400	690
29824	$2Hz$	8	9200	810
29826	$2Hz$	6	9300	990
29828	$2Hz$	4	9375	1300
29832	$2Hz$	3	9375	2800

Table 3.1: Multiphoton PCA runs taken on January 31, 2003. For all runs the laserball was at (0., -21.59, 64.14) (cm) and the detector was triggered with EXT ASYNC. The neutral density filter settings used that day are available in the shift report. The mean values for uncalibrated QHS were estimated using the XSnoed event display program.

Table 3.1.

The routine that measures the multiphoton timing correction for data is very similar to that for Monte Carlo. The mean time residual for prompt light (defined as light falling in a 17ns window around the single photon prompt peak) is found in a number of charge bins. About twice as many charge bins are necessary, as the multiphoton timing correction for data must account for both the multiphoton effect and the discriminator walk. In order to avoid contamination from single photon events it is required that at least 30% of the channels in an occupancy bin fire in order for data from that bin to be considered. The timing correction for data extends out to roughly 100 photoelectrons.

Applying the Correction to Data

The multiphoton timing correction is applied during the calibration phase of processing an event in SNOMAN. The timing correction was developed using the laserball calibration source so it is preferable to test the performance of the calibration using another source. Unfortunately no other intense calibration source was run during the D_2O or salt phases of the experiment, so the calibration must be tested using muons themselves. A technique for reconstructing muons without using the PMT timing information has been been developed and will be discussed later. By using a reconstruction routine that does not use charge information the multiphoton correction can be tested without introducing circularity. This reconstruction algorithm was run on fifty-six muons from the start of the D_2O phase, and a profile histogram of the charge-only fitter time residuals as a function of charge are shown in Figure 3.5 spanning a range from roughly half of a photoelectron to 1000 photoelectrons. Two thirds of the PMTs directly illuminated in a muon event have a charge in the range of 2-20 gain scaled charge units, and over this range the mean multiphoton corrected time residual varies by less than 0.2ns. The laserball intensity used in the multiphoton timing calibration allowed a correction to be developed only up to a charge of roughly 100 photoelectrons, but only about one in two hundred PMTs hit in a muon event have a charge larger than this. In the range of 2-100 gain scaled units the time residual RMS for directly hit phototubes with a time residual of less than 5ns is close to 1.1ns, significantly smaller than the single photon time spread of 1.6ns.

3.2 Summary

In this chapter four calibrations necessary for atmospheric neutrino analysis were discussed. First, good working channels are identified using the ANxx processor in SNOMAN. Second, the low gain charge measurement, QLL, was calibrated to the well understood high gain measurement, QHS. Third, the Monte Carlo PMT transit time distribution was altered to include the observed prepulse hits. Finally, the earlier calibrations were used to create a charge based correction for the multiple photon timing effect. The next chapter concerns the reconstruction of through-going muon events. The reconstruction algorithms make use the PMT charge and time information from working PMTs, so the calibrations

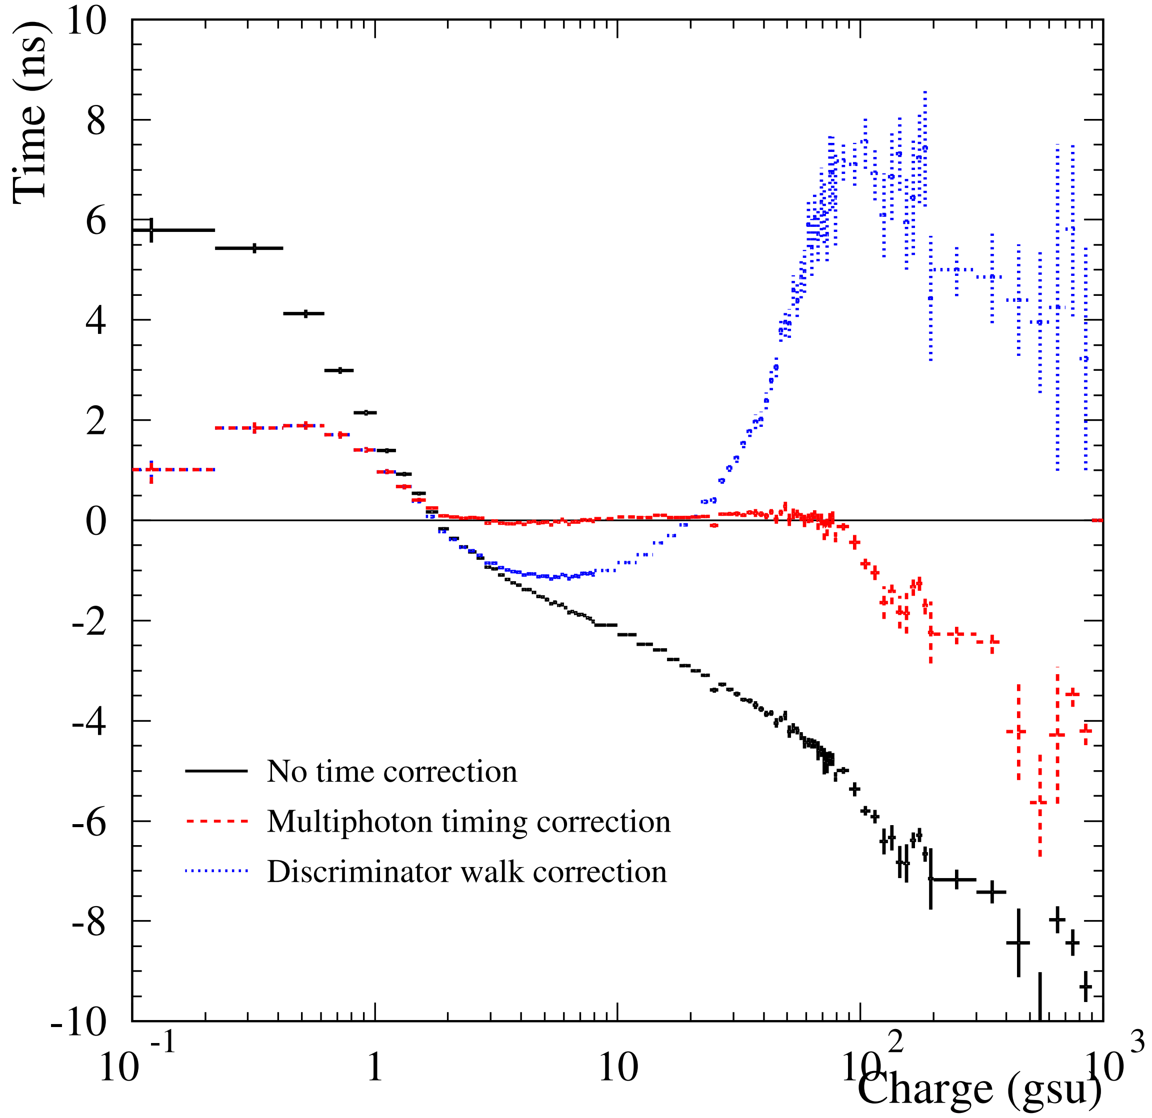

Figure 3.5: The mean time residual for a set of D_2O data muons reconstructed without using timing information is shown as a function of charge for three different timing calibrations. The solid black points show the time residuals if no timing correction other than a phototube cable length correction is applied. The dashed points show the time residuals with the multiphoton timing correction applied, and the dotted points show the time residuals if the regular SNO discriminator walk correction is applied. The walk correction is intended for single photon data and clearly over-compensates the time at large charges. Most analyses using the multiphoton timing correction only use tubes with charges greater than two gain scaled units in order to avoid the problematic region in which the data switches from single to multiple photons.

discussed in this chapter were necessary to make reconstruction possible.

Chapter 4

Reconstruction

The preceding chapter discussed the calibration of information from individual PMTs. In this chapter the PMT information is used to reconstruct through-going muon events. Reconstruction is the process in which the parameters that describe a given muon track, entrance position and direction, are estimated using the PMT information. Specifically, the reconstruction algorithms make use of each PMT's spatial position along with the charge and time it observed. Two algorithms are discussed in this chapter. The first is very fast and is used to produce a seed for the second algorithm. The second algorithm is a likelihood fitter which finds the track definition that maximizes its likelihood function.

4.1 Introduction

Cosmic ray muons are the most prominent signal in large water Cherenkov detectors. The rate of cosmic muon events is typically orders of magnitude larger than the neutrino signal or possible proton decay signal the detectors were designed to study, and cosmic muons are thus the most important background to such measurements. The cosmic muon background has in general been easy to separate from other signals because these muons typically produce Cherenkov light over a track several meters in length, and they all propagate downward. A technique for reconstructing muon tracks by minimizing the photomultiplier tube (PMT) time residuals has been developed and used by several experiments [51], [52], [53].

The Sudbury Neutrino Observatory (SNO) is the first approximately spherically symmetric large water Cherenkov detector, and in principle it should be possible to exploit this symmetry to improve event reconstruction and identification. The reconstruction advantage[1] of a spherically symmetric detector is most obvious for through-going muons, muons produced outside the detector that travel through the detector without stopping. In a spherical detector such as SNO the pattern of hit tubes and the distribution of PMT charge and time measurements depends upon a single parameter. This is the impact parameter, the distance from the center of the detector to the point of closest approach along the muon's track. In SNO, muons with different trajectories but the same impact parameter produce similar hit patterns, as is shown in Figure 4.1. This is in stark contrast to the situation in a detector such as Super-Kamiokande which does not have spherical symmetry. In these detectors downward traveling and horizontal muons produce very different hit patterns.

That the impact parameter is of paramount importance can be seen from mathematical considerations. Four parameters are necessary to fully describe the geometry of a line intersecting a sphere, and there are many possible parameterizations. Two fairly obvious choices of parameterization are to either specify the entrance and exit positions in polar co-ordinates, or the entrance position and muon direction in polar co-ordinates[2]. Another way to parameterize the problem is to define a special frame in which the muon track is fully described by the impact parameter, and to then find the rotation matrix required to take any particular event into that special frame. In addition to being a way of parameterizing the

[1]Another advantage of the spherical design is that it allows SNO to use light concentrators, which increase the effective photocathode coverage of the detector. These concentrators are typically not used in non-spherical detectors because the reflections from them would increase the difficulty of calibration.

[2]Note that since the entrance and exit positions are fixed at a radius of 850cm it only requires two parameters to specify each

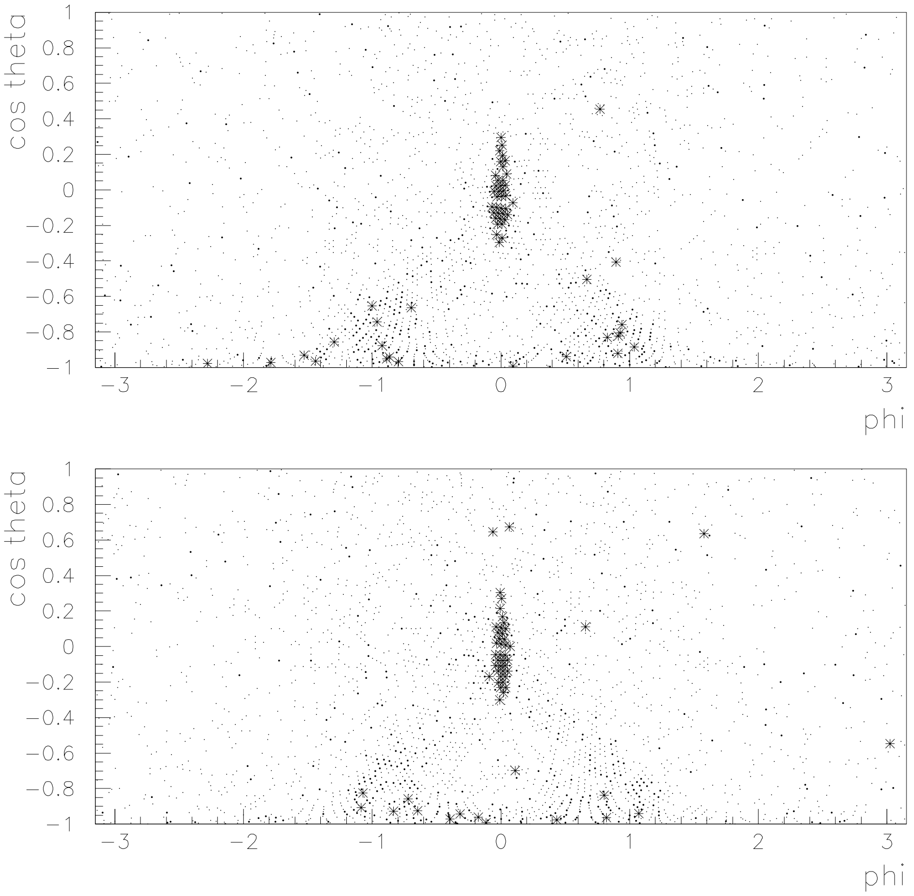

Figure 4.1: The pattern of hit PMTs for two simulated muons, both with an impact parameter of 810cm. Hit PMTs are displayed, with the size of the dot or star representing the charge observed by each hit PMT. Small dots represent PMTs which measured less than 4 photoelectrons (pe), large dots between four and 15pe, and stars above 15pe. In the top figure the neck of the detector is centered at approximately (θ, ϕ) = (-1.3, 0.55), while in the bottom figure it is at approximately (2.3, -0.05).

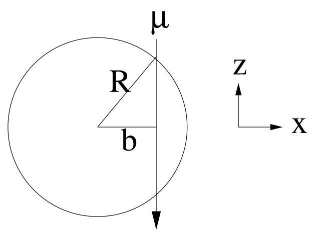

Figure 4.2: Muon trajectory for a muon with impact parameter b in the special frame.

problem defining a special frame has many computational benefits. Designing a routine to calculate PMT time residuals, for example, will be considerably easier if the calculation can be done in a specific rather than a general frame. The reconstruction algorithms described in this chapter make use of such a special frame, in which the muon's equation of motion can be written:

$$\vec{x}(t) = b\hat{x} + (\sqrt{R^2 - b^2} - ct)\hat{z} \tag{4.1}$$

where the origin is at the center of the detector, b is the impact parameter, and the muon enters the detector at t=0. A diagram of a muon track in this frame is shown in Figure 4.2.

This chapter describes two techniques that exploit the SNO detector's spherical symmetry to reconstruct through-going muons. The techniques are coded as SNOMAN processors and are referred to by their three letter processor names, FTR and FTI. The first technique, FTR, is a very fast and simple algorithm that uses data from only a subset of the hit PMTs in an event. It was developed to provide a seed for FTI. The FTI routine reconstructs events by maximizing a likelihood function calculated using the observed hit, time, and charge information of all the PMTs. The heart of FTI is a set of impact-parameter-based lookup tables for muon events in the special frame described above.

4.2 The FTR Algorithm

There is a fairly well established procedure for obtaining a seed for through-going muon fitters for water Cherenkov detectors. The basic principle is to identify two sets of PMTs, those with the earliest hit times and those with the highest charges, and to associate the muon entrance with the first set and the exit with the second. The FTR algorithm is a variation on this basic technique, in which we adjust for the fact that much of the bias in the basic fit depends upon the impact parameter. We will return to this picture of the algorithm at the end of this section; for ease of explanation the algorithm will be described in geometrical terms rather than in terms of a perturbation of the "high charge early time" procedure.

4.2.1 Defining a Specific FTR Co-ordinate System

The FTR algorithm works by reducing the event geometry to two special unit vectors that correspond roughly to the directions from the center of the detector to the muon entrance and exit positions. The two unit vectors are used to rotate the event into a special frame in which FTR has a standard prediction for the muon track parameters. The routine will correct for systematic mistakes in the estimation of the entrance and exit positions, so in establishing these vectors the robustness of their calculation is considered much more important than their correlation with the actual entrance and exit positions.

The search for the entrance and exit position of a muon in an ideal detector would be trivial, but in practice some amount of care is required. Photomultiplier tubes may have fired before the muon entered the detector, due to either the PMT dark current or light from radioactive backgrounds, and for this reason it is not advisable to simply search for the earliest hit PMT. Miscalibrations or failure of the front end electronics may result in incorrect measurement of the PMT times and charges, and this may adversely affect the searches for both early and high charge tubes. After efforts have been taken to remove tubes with bad calibrations, the next step is to use many tubes to estimate the entrance and exit positions.

The FTR routine searches for the muon entrance position by first sorting the PMT hit times to find the 80 earliest hit tubes[3]. The routine then finds the median x, y, and z positions of these 80 earliest PMTs. To remove spurious hits, the angle between each early PMT and the median position is measured from the center of the detector, and early tube selection is then restricted to the 20 earliest tubes for which that angle is smaller than about 26 degrees[4]. The direction from the center of the detector to the mean position of these 20 tubes is the first of the two unit vectors used to describe the event geometry. The vector is called TVec, and in simulations the mean angle between TVec and a vector pointing from the detector center to the true entrance is $\sim 6.5^o$. It was found that this angle could be reduced slightly by limiting the number of early tubes or by further restricting the allowed angle between the PMTs and the median early position. However, both of these changes make the routine more susceptible to early light and channels with electronics problems, so the looser cuts defined above were used in the final version of FTR.

An estimate of the exit position of the muon could be found by a similar prescription using high charge rather than early PMTs, but FTR employs a different scheme to avoid possible problems associated with using small numbers of tubes. The exit position is estimated by calculating the charge weighted average position of all of the hit PMTs, and a second unit vector is defined from the center of the detector to that position.

$$\vec{Q}Vec = \frac{\sum_{PMT}(Q_{PMT} \times \vec{X}_{PMT})}{\| \sum_{PMT}(Q_{PMT} \times \vec{X}_{PMT})\|} \tag{4.2}$$

[3]If fewer than 80 tubes are hit the FTR algorithm fails with a “no data” flag set. In the next chapter we will see that muons typically produce many thousands of PMT hits.

[4]$\cos\theta > 0.9$

This unit vector is called QVec, and in simulations the mean angle between the true muon exit position and QVec was found to be $\sim 5.2^o$. Several alternate definitions of QVec were tested[5], but this simple definition had a smaller average error than almost all of them. A very slight decrease in the angle between QVec and the muon exit position was noted when a charge cut of 2 gain scaled units was applied, but in the interests of simplicity and because of possible differences in the data and Monte Carlo gain scaling, the simple definition of QVec was used.

Once QVec and TVec have been measured for an event a new co-ordinate system is set up using them. In this section (x, y, z) will refer to co-ordinates in the SNO detector frame, and (x_R, y_R, z_R) to the FTR frame. The FTR algorithm calculates the rotation matrix which results in QVec pointing in the $-\hat{z_R}$ direction and TVec in the $x_R z_R$ plane with positive x_R. Diagrams of TVec and QVec are shown in the detector and FTR frames in Figure 4.3 for two different muons.

4.2.2 Muon Direction in the FTR Frame

In Figure 4.3 TVec and QVec are shown pointing to the true entrance and exit positions of the muon track. In reality, TVec and QVec have both a bias and a statistical spread about the true positions. While the cosine of the angle between TVec and QVec ($\cos\alpha_{QT}$) is highly correlated with the impact parameter the relationship is not $\cos\alpha_{QT}= 2(\frac{b}{R})^2 - 1$ as it would be if TVec and QVec pointed to the true positions. Nevertheless, it is relatively easy to find the true entrance and exit positions as a function of the $\cos\alpha_{QT}$.

Several thousand cosmic through-going muons were simulated and the muon entrance position and direction in the FTR frame were calculated. From reflection symmetry the mean entrance position and the muon direction in the FTR frame should have $y_R = 0$, and this was found to be the case. This means that to describe the muon track in the FTR frame

[5] for example using a higher power of Q_{PMT}, or only summing tubes above a certain charge

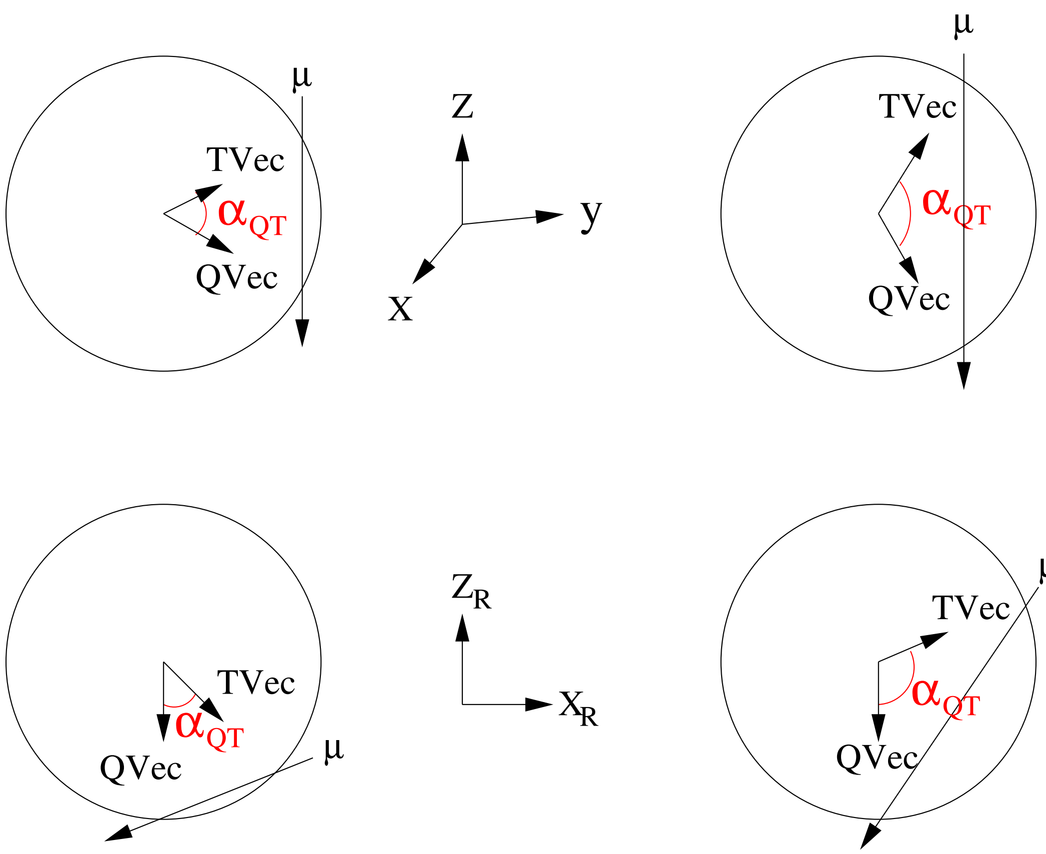

Figure 4.3: Schematic diagram showing QVec and TVec for two muons in the detector (top) and FTR (bottom) frames. One can see from this diagram that the impact parameter is related to the cosine of the angle between QVec and TVec, $\cos\alpha_{QT}$.

only two parameters are needed. The entrance position is described by the z_R coordinate[6], and the muon direction is described by the zenith angle in this rotated frame[7].

The distribution of entrance positions for simulated muons in the FTR frame is shown in Figure 4.4, along with a fifth order polynomial fit. Because the distance between the entrance and exit positions decreases with increasing impact parameter, $\cos\alpha_{_{QT}}$ increases monotonically with increasing impact parameter. In the FTR frame the exit of the muon should always be near z_R=-850cm, and at zero impact parameter ($\cos\alpha_{_{QT}} = -1$) the entrance is at z_R=+850cm. As $\cos\alpha_{_{QT}}$ increases the z_R coordinate of the entrance point becomes smaller. A scatter plot showing muon directions in the rotated frame is shown in Figure 4.5 with an eleventh-order polynomial fit. In the case of zero impact parameter, the muon's direction is directly downward ($\cos\theta_{zenith} = -1$), and at very large impact parameters the muon is traveling horizontally.

The polynomial fits to the two data sets describe the relationships of the average entrance position and muon direction in the FTR frame to $\cos\alpha_{_{QT}}$. The values of QVec and TVec 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 calculate 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 fit parameters in SNO detector co-ordinates.

4.2.3 FTR Performance on Simulated Data

The performance of the FTR algorithm was tested using 19540 simulated cosmic muons, an amount roughly equivalent to 85% of that expected in for the D_2O phase. The muons were generated at a starting radius of 925cm with the energy and direction distributions

[6] y_R=0, and the x_R position of the entrance is constrained by $\sqrt{(x_R^2 + z_R^2)} = 850cm$

[7] Since the average y_R displacement of the entrance and exit is zero in this frame the average muon azimuthal direction is also zero.

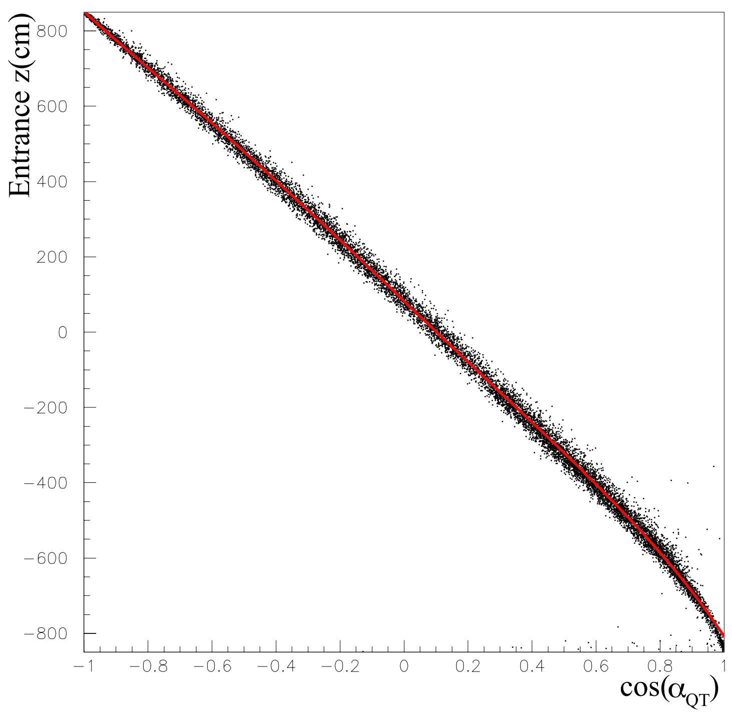

Figure 4.4: Scatterplot of the z_R coordinate of the entrance position vs. $\cos\alpha_{QT}$ for simulated muons in the FTR frame. In this frame the exit position of the muon should always be near z_R=-850cm. A fifth order polynomial fit to the data is shown in red.

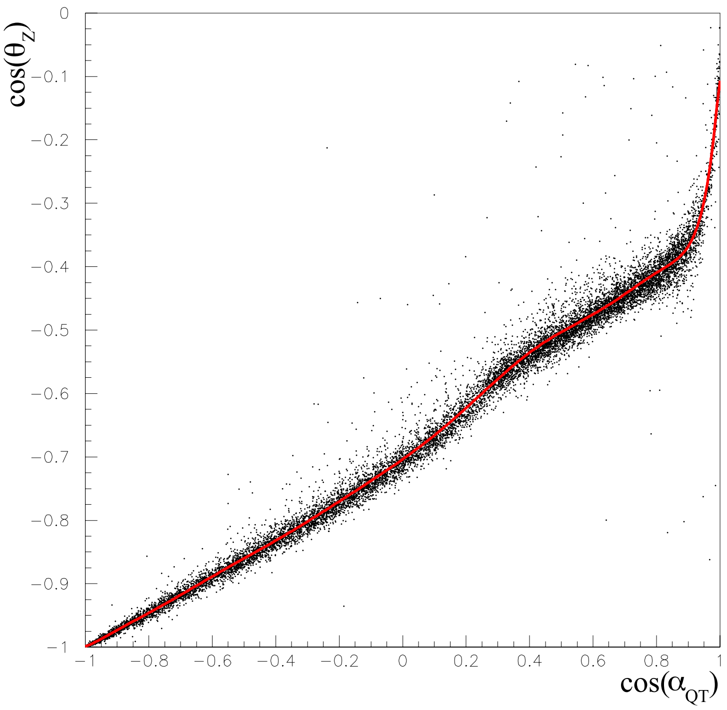

Figure 4.5: Scatterplot of muon direction vs $\cos \alpha_{QT}$ in the FTR frame. An eleventh order polynomial fit to the data is shown in red.

expected for cosmic muons. Comparisons of fit values to Monte Carlo values are shown in Figures 4.6 and 4.7. The first set of figures shows the quality of the impact parameter fit, while the second set shows the quality of the direction reconstruction. The parameter $\cos \theta_{misrec} = \hat{x}_{fit} \cdot \hat{x}_{MC}$ would have a value of 1 for a perfectly reconstructed direction, a value of -1 would indicate the track was reconstructed in exactly the opposite direction.

As can be seen in Figure 4.6, the fitter has a bias in the fitted impact parameter. Since the bias is smaller than the resolution, and since the fitter result is intended to serve only as a seed, this bias is acceptable. At the end of this chapter we will see that the fitted impact parameter distribution for the more advanced fitter (FTI) is flat. In the bottom right hand plot it can be seen that at the edge of the detector FTR has an outward bias, this is the reason for the peak near 1 in the bottom left hand plot.

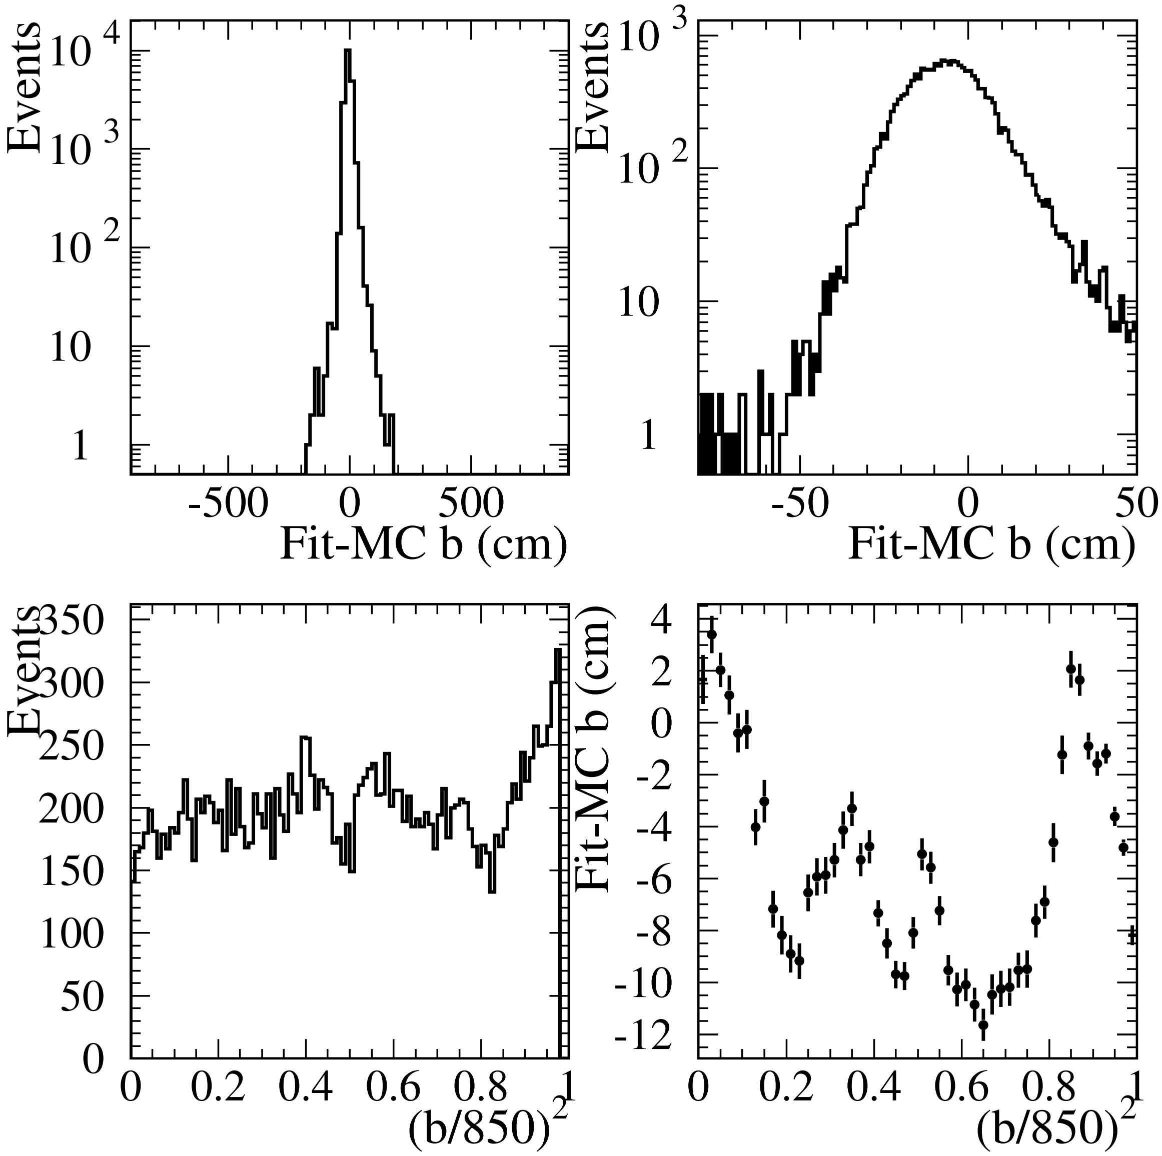

Figure 4.6: These figures show the FTR algorithm's performance on simulated cosmic muons. The top two plots show the error on impact parameter. A Gaussian fit to the peak returns a mean error of -6.57 ± 0.08cm, with an RMS of 11.7cm. The bottom left plot shows the distribution of the square of the reconstructed impact parameter divided by the PSUP radius $(b/850)^2$, with this abscissa the AV is at 0.5. The plot at bottom right is a profile histogram of the fit impact parameter error as a function of Monte Carlo impact parameter. No data selection cuts have been applied, so these plots include both through-going and stopping muons.

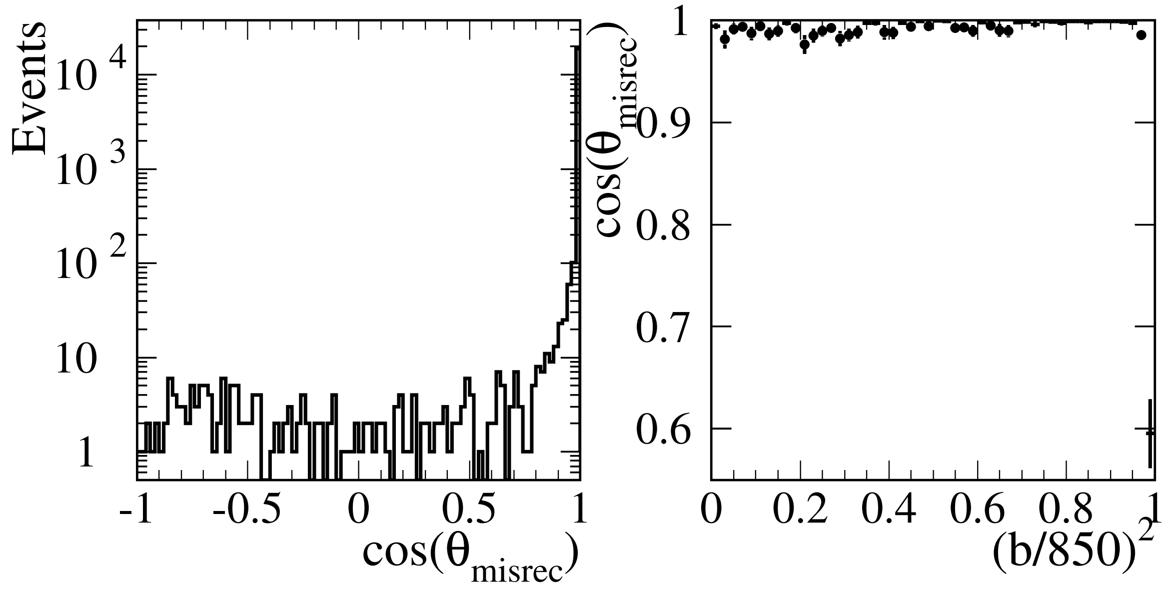

Figure 4.7: Errors in fitted direction for FTR fits to simulated cosmic muons. $\cos\theta_{misrec}$ has a value of 1 for perfectly reconstructed tracks. The left plot shows the distribution of the cosine of the angle between the fitted and Monte Carlo directions. Note that almost all of the data lies in the last bin. The right hand plot is a profile histogram showing the cosine of the angle between the fit and simulated track vs. impact parameter squared. Almost all of the events with large errors, including the bump at -0.7, correspond to either stopping muons or events within a few centimeters of the PSUP.

4.2.4 Muon Analysis Blindness and FTR

The muon analysis cannot claim to be blind, as several people had looked at most of the data before the current analysis was pursued. Nevertheless, in the interest of avoiding bias a blindness scheme was developed and used in this analysis [54]. The scheme consists of two parts: data division and a spoiled acceptance. The first parts of the $\mathrm{D_2O}$ and salt phases, runs 10000 through 12168 and 20684 through 22450 are part of the "open" data set[8]. The open division further had its acceptance spoiled by an exact fraction (near 40%) for events with an FTR estimate of $\cos(\theta_{zenith}) > 0.6$, and a zenith angle dependent function which is unknown to the analysts for events with $\cos(\theta_{zenith}) < 0.6$. Six potential blindness functions were written and installed in SNOMAN by Scott Oser, a SNO collaboration member outside of the muon group. Nick West, the SNOMAN administrator[9], then chose and activated one of those functions.

In order to allow the study of the zenith angle dependence of the cosmic muons before blindness is removed, it is important that FTR does not introduce a large bias in the angular distribution. The generated zenith angle distribution for 800 live days of simulated muons is shown along with the FTR-reconstructed angular distribution in Figure 4.8. The figure indicates that FTR does not introduce an unacceptable bias.

4.2.5 Summary of FTR

The FTR reconstruction algorithm exploits the approximate spherical symmetry of the SNO detector to mitigate the impact parameter dependent bias from a fit based upon searching for early time and large charge tubes. The FTR algorithm was tested with Monte Carlo through-going cosmic muons, and was found to have small reconstruction errors for muons that enter more than a few centimeters into the detector. The cosmic-muon zenith-angle

[8]The open list contains 51% of the $\mathrm{D_2O}$ data and 15% of the salt data
[9]extraordinaire

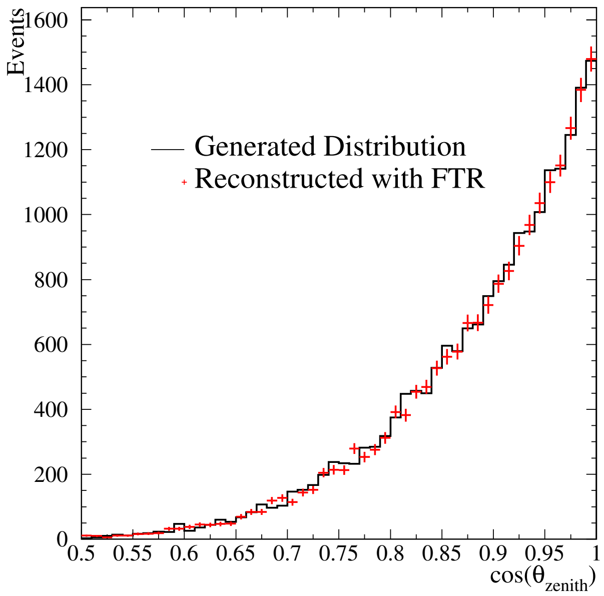

Figure 4.8: Generated zenith angle distribution compared with the distribution measured with FTR. Note the linear scale.

distribution produced by FTR appears to agree well with the input distribution, so the use of the FTR-measured zenith angle in the muon blindness scheme is not anticipated to bias the zenith angle distribution of the non-blind spectrum above $\cos(\theta_{zenith}) = 0.6$.

4.3 The FTI Algorithm

Muon reconstruction in water Cherenkov detectors has traditionally been done by minimizing the χ^2 of the PMT time residuals. The PMT charges have been used in selection of tubes, but the charge information has not typically been used in reconstruction. Also, there is information in which tubes were not hit, but this information is not easy to incorporate in a χ^2 minimization routine. A maximum likelihood technique has been developed that makes use of all the available information. For a hypothesized muon direction, entrance position and entrance time, the routine calculates the probability that each PMT would be hit, and if hit, the probability that it would "observe" its recorded charge and time.

The routine employs a set of lookup tables to get the number of photoelectrons expected to be detected by each tube given the hypothesized event geometry. With this information, it calculates the probability that a tube would have fired and the probability that it would have observed the charge and time recorded. The lookup tables were produced by simulating thousands of muon events at each of a set of impact parameter values. Interpolation is used to obtain predictions between the values of the tables.

4.3.1 Likelihood Function

Depending on the PMT position relative to the hypothesized muon track the number of photoelectrons expected to be detected by any phototube varies from much less than one up to several hundred. With Poisson statistics, if the expected number of detected photoelectrons is λ then the probability that the tube would not fire at all is simply

$$P_N(0|\lambda) = e^{-\lambda} \tag{4.3}$$

and the probability that it would observe *n* photoelectrons is

$$P_N(n|\lambda) = \frac{e^{-\lambda}}{n!}\lambda^n \tag{4.4}$$

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

$$P(hit, Q, t|\lambda) = \sum_{n=1}^{\infty} P_N(n|\lambda) P_Q(Q|n) P_T(t|n) \tag{4.5}$$

where the conditional probabilities, $P_Q(Q|n)$ and $P_T(t|n)$ are yet to be defined.

The likelihood of the observed PMT hits, times, and charges for a given hypothesized muon track is simply the product of the probabilities for the individual PMTs. The following subsections discuss the derivations of $P_Q(Q|n)$ and $P_T(t|n)$, how the lookup table to determine λ was created, and the interpolation scheme for impact parameter values between the table points.

It is useful to point out at this point that it is possible to remove the time (charge) term in Equation 4.5 by integrating over all hit times (possible charges) before performing the sum. This is a useful trick, as it makes it possible to run FTI in a "charge only" or "time only" mode. The "charge only" version of FTI was used to evaluate the multiple photoelectron timing correction, and the result was shown in Figure 3.5 of the last chapter.

4.3.2 Probability of Detecting a Particular Charge Given a Certain Number of Photoelectrons

The SNO experiment was built to detect solar neutrinos, the signature of which is an electron with roughly 5 to 15 MeV of energy. To detect these events it is important that the phototubes have large photocathode coverage, a narrow transit time spread, and as little glass as possible (since many radioactive elements are present in phototube glass). Charge resolution is not particularly important for a solar neutrino analysis, as the probability of a phototube detecting two photons is near 1% and thus almost all PMT hits in a neutrino event are single photoelectron hits. In addition to having a fairly broad peak for single photoelectron hits the charge distribution of the SNO PMTs has a very long tail, extending to the expected charge for 10 photoelectrons.

While the broad charge distribution means that a measured charge can not be used to say with certainty that a phototube was hit by n photoelectrons it is easy to rule out certain possibilities. A phototube that detects a fraction of the mean charge for one photoelectron did certainly not detect 10 photoelectrons, and a phototube that reports a charge equivalent to hundreds of pe almost certainly did not detect only one single photoelectron. The single photoelectron charge spectrum is shown for data and Monte Carlo in Figure 4.9. The

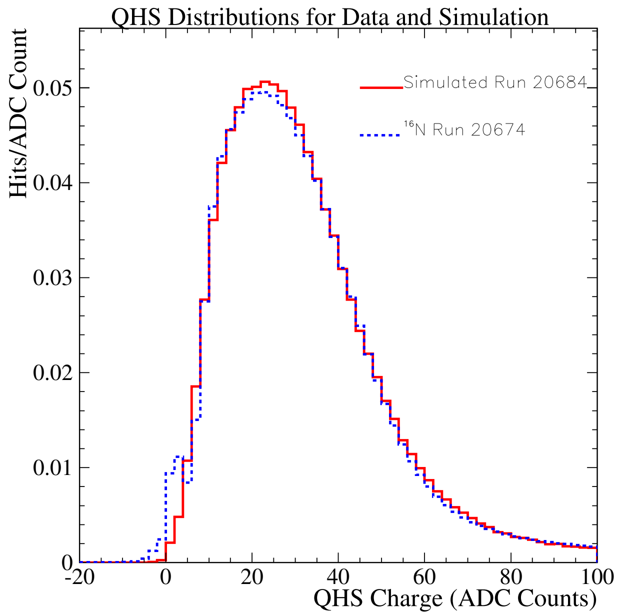

Figure 4.9: Single photoelectron charge spectrum for PMTs, produced using the ^{16}N gamma ray source. Data is shown as dashed blue, the SNOMAN simulation as solid red. Figure courtesy of J. Klein.

agreement between the Monte Carlo charge simulation and the data allows one to actually assign probabilities to the hypotheses that a phototube which observes charge q detected one, two, or *n* photoelectrons.

It is easy to produce a charge distribution for PMTs hit by two, three, or *n* photoelectrons using the SNO phototube Monte Carlo. By summing a large number of pairs of charges selected using the single photoelectron charge distribution one obtains the two photoelectron charge distribution, and an analogous procedure is used for the *n* photoelectron charge distribution. This technique was used to produce charge distributions for each value from one to one hundred photoelectrons. Some of the charge distributions appear in figure 4.10. The entries in these charge distributions are the conditional probability, $P_Q(Q|n)$ of observing charge Q given *n* detected photoelectrons.

The $P_Q(Q|n)$ conditional probability distributions were produced using a combination of the SNOMAN Monte Carlo and some simple independent FORTRAN code. The SNO-

MAN PMT Monte Carlo was used to generate 100 simulated single photoelectron hits on every online tube for run 10000. The charges reported by the MC were histogrammed and this histogram was used to produce the cumulative distribution for the probability of observing charge Q given a single photoelectron. This cumulative distribution was then used in an independent FORTRAN routine to produce charge histograms for *n* detected photoelectrons, where *n* ran from 1 to 100. The single pe cumulative distribution was used to produce the multiple photoelectron histograms instead of the SNOMAN MC for speed. The multi-pe histograms were filled with 100000 entries for each of 1-10 detected photoelectrons, and with 10000 entries for 11-100 detected photoelectrons.

It was found that above 10 photoelectrons the shape of the charge distribution is well parameterized by a two sided Gaussian function. The fit did not parameterize the data well below 10pe, so below 10pe the charge histograms were normalized to produce charge PDFs for FTI. Above 10pe the histograms were fit with the function

$$P(Q|n) = \begin{cases} N \exp\left(\dfrac{-(Q-\bar{Q})^2}{2\sigma_1^2}\right) & Q < \bar{Q} \\ N \exp\left(\dfrac{-(Q-\bar{Q})^2}{2\sigma_2^2}\right) & Q > \bar{Q} \end{cases} \tag{4.6}$$

where the normalization is

$$N = \frac{\sqrt{2}(\text{entries} \times \text{bin size})}{\sqrt{\pi}(\sigma_1 + \sigma_2)} \tag{4.7}$$

and the fit parameters are mode charge $\bar{Q}$ and the left (σ_1) and right (σ_2) RMS of the two sided Gaussian. Two of the fit distributions are shown in Figure 4.10. The value of $\bar{Q}$ was found to be linear as a function of detected number of photoelectrons, and the values of σ_1 and σ_2 were found to be linear as a function of the square root of the number of detected photoelectrons for hits with 10 or more photoelectrons. The parameters of equation 4.6 for PMTs detecting between 10 and 100 photoelectrons have therefore been reduced to three

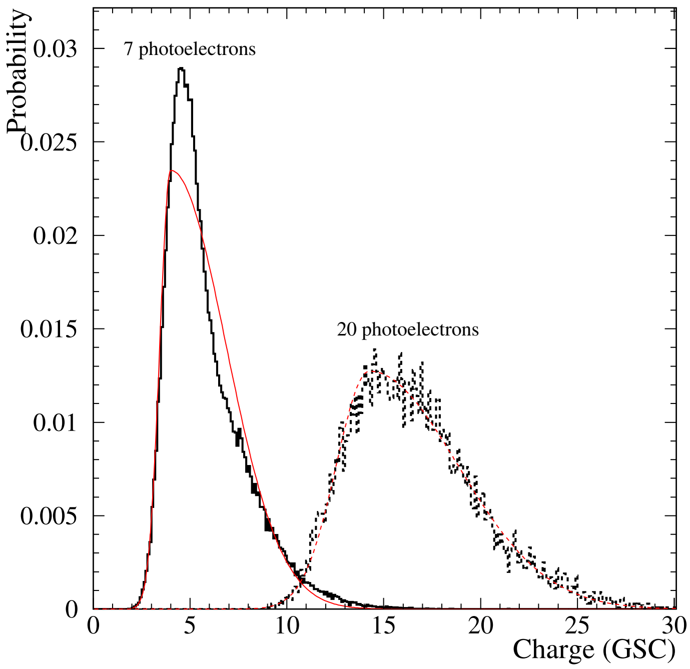

Figure 4.10: Charge distribution for Monte Carlo PMTs which detect 7 (solid) or 20 (dashed) photoelectrons. The histograms are shown with a fit to the double sided Gaussian described by Equation 4.6. The fit models the data well in cases where a PMT detects more than 10 photoelectrons.

linear functions:

$$\bar{Q} = -2.3531 + 0.83075n \tag{4.8}$$

$$\sigma_1 = -1.7565 + 0.7807\sqrt{n} \tag{4.9}$$

$$\sigma_2 = 1.1351 + 0.77393\sqrt{n} \tag{4.10}$$

The agreement between this set of functions and the histograms shown in Figure 4.10 is comparable to the individual fit shown in the figure. It should be mentioned that in generating the n-photoelectron charge distributions the average single-pe charge distribution for a single $\mathrm{D_2O}$ run was used. It has been found that the shape of the PMT charge distribution has some dependence on phototube gain [41], so it is possible that the charge distributions for *n* photoelectrons in PMTs with unusually large or small gains are slightly different from the average. This is not expected to be a major issue for two reasons. The first is that because the charge data for each PMT is scaled using an individually measured gain the mean charge for *n* pe should be fairly consistent across all tubes. The second is that in the D_2O and salt phases the spread in PMT gains is not large, and PMTs with unusually large or small gains are removed from the analysis[10].

As was mentioned earlier, if one integrates over all possible times in 4.5 it is possible to run FTI in a "charge only" mode. Now that we have generated the conditional probabilities for a PMT to detect a particular charge given a particular number of photoelectrons it is possible to combine the first two terms of Equation 4.5 to get the PMT probabilities in the charge only mode:

$$P(Q|\lambda) = \sum_{n=1}^{\infty} P(Q, n|\lambda) \tag{4.11}$$

$$= \sum_{n=1}^{\infty} P_N(n|\lambda) P_Q(Q|n) \tag{4.12}$$

$$= \sum_{n=1}^{\infty} \left(\frac{e^{-\lambda}}{n!}\lambda^n\right) P_Q(Q|n) \tag{4.13}$$

4.3.3 Probability of a PMT Firing at Time t Given *n* Photoelectrons

In any experiment using photomultiplier tubes with threshold crossing discriminators the distribution of observed PMT times depends upon three factors. These are the emission properties of the source, the propagation of the photons from the source to the PMTs, and the timing response of the PMT and electronics to single photoelectrons. To illustrate the effect each of these factors can have on the final PMT time distribution it is useful to

[10]PMTs with high half points outside of the range 30-80 gain scaled units are flagged during initialization and are considered off for the analysis.

consider a simple setup like a single PMT observing a source at the opposite end of a long light tight box. We will start with this simple artificial scenario and increase the complexity to demonstrate the effect each factor has.

As a first step, imagine that we have a single photon light source for which we know the exact photon emission time, that there is a perfect vacuum separating the PMT and source, and a PMT and electronics setup with a delta function timing response. In this artificial scenario the observed PMT time distribution would be a delta function occurring at a time $t = d/c$ after the emission. The first two real world considerations to take into account are the understanding of the source photon emission time and the photoelectron transit time distribution inside the PMTs. It is not possible to know the exact time at which a source emitted its photons but the uncertainty on this can probably be reduced to a sub-nanosecond level, so with a good source setup the photon arrival time distribution will be dominated by the transit time distribution of the SNO PMTs. The intrinsic PMT transit time distribution used in the SNO Monte Carlo is shown as the black dashed line in Figure 4.11.

With our conceptual setup we can now consider the effect of PMTs observing multiple hits. The SNO detector uses threshold crossing discriminators, so if a PMT is hit by two photoelectrons the photoelectron with the shorter transit time will fire the discriminator. This effect can be simulated by sampling from the single photoelectron transit time distribution twice and choosing the earlier of the two times. This procedure was repeated 10^6 times to produce the dotted blue histograms in Figure 4.11. Note in the left hand plot that even with only two photoelectrons the late part of the distribution has been greatly reduced, and with three pe the late times are almost eliminated. As can be seen in the right hand plot, the "prompt peak" also shifts to an earlier time[11]. For this conceptual experiment let us assume that there exists a multiple photoelectron timing correction to take that into

[11] The multiple photoelectron timing correction discussed in subsection 3.1.4 was designed to account for this effect

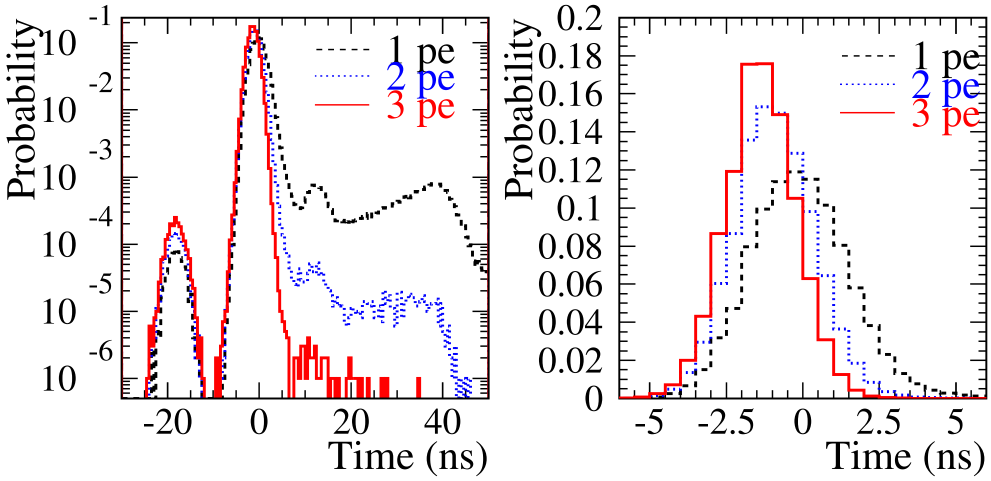

Figure 4.11: These figures show the large effect that multiple photoelectron hits have upon the PMT timing at small numbers of pe. The Monte Carlo PMT time distribution was sampled one million times to generate the dashed black histogram. The same distribution was sampled in pairs and the earlier of the two hits was used to produce the dotted blue histogram, and similarly for three pe. From the logarithmic plot on the left we can conclude that for PMTs hit by multiple photoelectrons, late light (from Rayleigh scattering or reflections) can effectively be ignored when considering the PMT time information.

account.

Let us now consider the flight path of the photons from the source to the PMT. If we were to fill our box with a transparent medium, say water, then the velocity of the photons will no longer be c, but will be scaled by the index of the media. If the photon source does not produce a single wavelength then the photons will no longer arrive at the same time due to velocity dispersion. For PMTs that observe single photoelectrons this will introduce a spread in the time distribution which is proportional to the length of the box. In the case of multiple detected photoelectrons the dispersion will be have a smaller effect, as high wavelength photons (which travel through water at a higher velocity) will tend to arrive earlier. Another effect that adding water would have is that photons could now scatter, increasing the distance they must travel. If we consider the full SNO detector rather than our simple box setup it also becomes possible for photons to refract in (changing their path length) or reflect from the AV (greatly increasing their path length). Once again, in the multiple photoelectron case reflected photons will have almost no impact on the timing because of the threshold crossing discriminator.

Having gone over the factors that affect the PMT timing we can move onto the PMT time distributions expected from light produced from a muon traveling through the SNO detector. The first thing to consider is that for a muon event in SNO the source emits a continuous spectrum over a long time period, and the distance from the source to the PMT is different for each PMT. To make comparisons of the PMT times it is therefore necessary to define a "time residual". For the sake of simplicity we will temporarily assume that the muon is the only source of light in the detector, and that rather than emitting a continuous spectrum it only produces a single wavelength. In this case each photon is emitted at a specific angle from the muon (the Cherenkov angle), and it propagates with a single speed. A PMT time residual can therefore be defined by subtracting the propagation time of the photon and muon from the observed PMT time. The FTI routine assumes a Cherenkov angle of 42^o, and a speed of light of 21.8cm/ns.[12] The PMT time residual is thus:

$$t_{res} = t_{PMT} - t_0 - \frac{d_1}{c} - \frac{d_2}{21.8\text{cm/ns}} \tag{4.14}$$

where t_{PMT} is the PMT hit time, t_0 is the global fit offset time, d_1 is the distance the muon propagated before it emitted the Cherenkov photon, c is the speed of light in vacuum, and d_2 is the distance that the Cherenkov photon is expected to have propagated. This definition of the time residual is shown in Figure 4.12.

With a choice of definition for PMT time residuals in SNO it is now possible to consider what the probability that a PMT illuminated directly by a muon will observe a particular time residual. To make the problem tractable FTI uses three probability distribution func-

[12]The FTI PMT time residual definition ignores the effects of velocity dispersion and refraction when photons pass through the acrylic vessel. In hindsight it is probable that these choices lead to a reconstruction bias due to the extinction of low wavelength light in the acrylic vessel. This bias should, however, be the same in Monte Carlo and data.

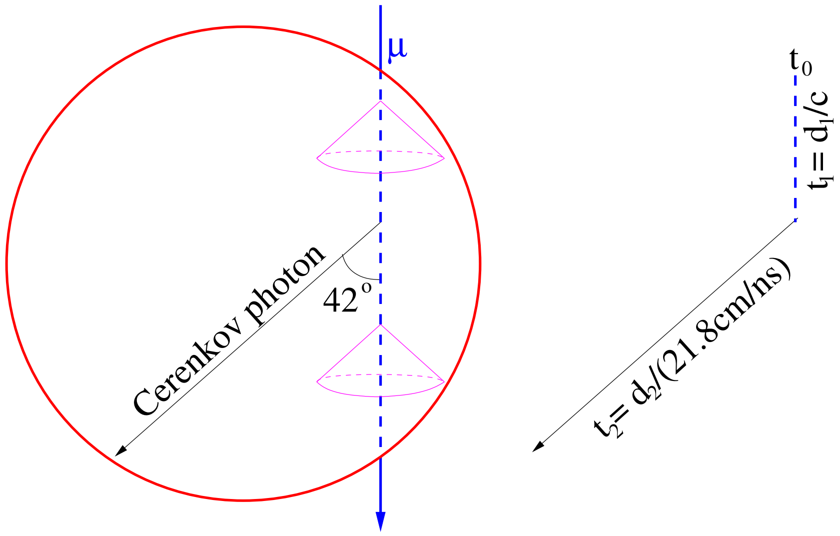

Figure 4.12: Time residual definition for FTI. The PMT time residual is the observed PMT time minus each of the Cherenkov photon propagation time, the muon propagation time, and the event t_0. The Cherenkov angle is defined as exactly 42 degrees, and Cherenkov photons are assumed to propagate at 21.8cm/ns. The two purple cones are included to remind the reader that the Cherenkov cone is produced continuously along the muon track, not just at the discrete points shown.

tions (PDFs) to describe the different parts of the single photoelectron timing distribution.

Prepulse: These hits are modeled by a Gaussian distribution with a mean time of -15.5ns and a width of 2.2ns. The position and width of the distribution come from [50]. The prepulse probability for light produced in the center of the detector at single pe light levels was found to be approximately 1/1500.

Prompt Light: These hits are modeled by a Gaussian distribution with a mean time of 0.0ns and a width of 1.6ns.

Early and Late Hits: All effects such as the PMT dark current, scattering, reflections, and the late part of the transit time distribution are very grossly modeled with a Heaviside function. The function transitions at 1ns, is normalized to 1 over the interval $[-150\text{ns}, 200\text{ns}]$, and the value of the function on either side of 1ns is such that the probability of hits before 1ns corresponds to a dark current of roughly 500Hz.

While these PDFs do not make use of all of the available information they have the virtue of being quickly calculable, a necessary feature given that the code must loop over thousands of hit tubes for each set of hypothesized track parameters. The next step is to decide what probability to assign to each of the PDFs for a given PMT hit. When the FTI lookup tables for the mean number of detected photoelectrons were developed the probability that a PMT hit would be "prompt" was also recorded. For the purposes of the table, prompt hits are defined as having a time residual between -25ns and 5ns. For each hit PMT in an event, FTI is therefore able to use the lookup tables to determine the probability that a detected photon was not prompt. This "not prompt probability" α is used to assign the probability of late hits. For hit PMTs with a time residual later than 1ns the probability assigned to being a late hit is simply 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:

$$
\begin{aligned}
P(t|1\text{ pe}) &= (1-\alpha)\Bigg(P(t|\text{prepulse}) + P(t|\text{prompt})\Bigg) \\
&\quad + P(t|\text{early or late}) \qquad (4.15) \\
P(t|1\text{ pe}) &= (1-\alpha)\Bigg(\rho\frac{1}{\sqrt{2\pi}\sigma_{pp}}e^{\frac{-(t+15.5)^2}{2\sigma_{pp}{}^2}} + (1-\rho)\frac{1}{\sigma_{pr}\sqrt{2\pi}}e^{\frac{-t^2}{2\sigma_{pr}{}^2}}\Bigg) \\
&\quad + \begin{cases} \frac{0.675}{9600}/150\text{ns} & t < 1\text{ns} \\ (\alpha)(1-\frac{0.675}{9600})/200\text{ns} & t > 1\text{ns} \end{cases} \qquad (4.16)
\end{aligned}
$$

where $\sigma_{pp} = 2.2ns$, $\sigma_{pr} = 1.6ns$, and ρ is the fraction of prompt hits that result in a prepulse. A few modifications are required to generalize this equation to *n* detected photoelectrons. The probability of having an early (dark current) hit does not depend on the number of pe detected, so the early part of the Heaviside function has no dependence on the number of detected photoelectrons. As mentioned earlier, PMTs that are hit by multiple pe are almost certain to not at least one photon be prompt, so the probability that the PMT observes a late time has a very strong dependence on the number of detected photoelectrons. The probability of all of the hits being late is simply α^n, and it follows that the probability of at least one prompt hit is $(1-\alpha^n)$. The last revision required to generalize to *n* detected photoelectrons concerns the frequency with which a photon produces a prepulse rather than a prompt pulse, ρ. Since only "prompt" photoelectrons can produce pulses before the prompt peak the probability of being in the prepulse region depends on the number of prompt pe, rather than the total number of detected pe. For this reason the prepulse fraction depends upon the expected number of prompt photoelectrons $\tilde{n} = (1-\alpha)(n)$. For a tube hit by $\tilde{n}$ prompt pe the probability of no prompt hits producing a prepulse is $(1-\rho)^{\tilde{n}}$, so the probability that at least one of the prompt photoelectrons results in a prepulse is therefore

$$\chi = 1 - (1-\rho)^{\tilde{n}} \tag{4.17}$$

We are now ready to write the probability of observing time t given *n* photoelectrons[13]:

$$\begin{aligned} P(t|n) &= (1-\alpha^n)\left(\chi\frac{1}{\sqrt{2\pi}\sigma_{pp}}e^{\frac{-(t+15.5)^2}{2\sigma_{pp}{}^2}} + (1-\chi)\frac{1}{\sigma_{pr}\sqrt{2\pi}}e^{\frac{-t^2}{2\sigma_{pr}{}^2}}\right) \\ &\quad + \begin{cases} \frac{0.675}{9600}/150\text{ns} & t < 1\text{ns} \\ (\alpha^n)(1-\frac{0.675}{9600})/200\text{ns} & t > 1\text{ns} \end{cases} \end{aligned} \tag{4.18}$$

A careful reader may notice that this probability distribution function is not actually normalized, as it has a constant non-zero value at very large and small times. The SNO electronics has a window in which PMTs can participate in an event, so the value t is not

[13]As a reminder, *n* is the number of detected photoelectrons, α is the probability that any given photon is not prompt, $\tilde{n}$ is the expected number of detected pe that are prompt, ρ is the frequency with which prepulses occur in single pe data, and χ is the probability that one of the prompt photoelectrons would have resulted a prepulse.

actually able to have any value. The size of the window is not identical for each electronics channel, so FTI assumes that the window extends roughly 150ns before the prompt peak and 200ns after it. A very careful reader may note that the last term should actually be $\alpha^n - \frac{0.675}{9600}$. This was an error that was discovered after the data had already been processed. It will be fixed in a future release of SNOMAN, but since the magnitude of this error is very small fixing it will not actually affect reconstruction.

4.3.4 Lookup Tables

The backbone of the FTI routine is the set of lookup tables that characterize muon events as a function of impact parameter. The tables were created to answer the question "given a PMT in a particular position relative to a muon track, what is the expected number of photoelectrons that the PMT will detect, and what fraction of those photoelectrons will be prompt?"

The tables describe muons in a frame of reference in which the muon position can be described by Equation 4.1. The surface of a hemi-sphere is divided into 10000 two-dimensional bins in polar angle (cosθ) and azimuthal angle (ϕ). The tables need only describe half of a sphere because in this frame the muon pattern is symmetric under reflection through the x-z plane. The pattern of PMT hits from a muon changes most rapidly in the areas with a large surface density of photons, and the surface density of photons is always greatest near the point where the muon exits. This muon exit point is always at $\phi = 0$ but has a different position in cosθ for each impact parameter. The FTI lookup table bins are equally spaced in cosθ but the are very narrowly spaced near $\phi = 0$ and become much larger as $\phi \to \pi$.

To generate the table for a particular impact parameter hundreds of muons were simulated with that impact parameter, and in each bin a set of counters were incremented event by event. The counters kept track of how many times the center of a phototube was in each solid angle bin, how many photoelectrons (if any) that phototube detected, and how many of those hits were "prompt" (-25ns<t<5ns). Since the goal of the tables is to describe muons which travel on a straight line path through the detector losing energy almost entirely through ionization, muons which deviated too far from their original direction were rejected[14]. For the same reason muons were rejected if they stopped in the detector, if some interaction produced a particle other than an electron or photon, and if their average energy loss as they crossed the detector was greater than 450MeV/m. The data will include muons which stop, multiple scatter through a large angle, or lose energy through radiative processes. The goal of the tables, however, is to describe "standard" muons which travel in a straight line all the way through the detector. The muons were simulated with a starting energy of 50GeV, as this energy was well below the critical energy in water[15] but still large enough that muons did not deviate far from a straight line due to multiple scattering. At this energy roughly half of the simulated events failed the cuts near the center of the detector, and this decreased to roughly 5% of events failing near the PSUP. At each impact parameter being simulated the muons were started isotropically on the surface of a sphere at a radius of 900cm. Once a starting position was chosen a random direction was chosen from the cone of possible directions with the desired impact parameter. The muons for all of the tables were simulated using the run conditions for salt phase run 20906 and with SNOMAN version 5_0086.

It takes a significant amount of cpu time to simulate many thousands of muons, so a choice had to be made between the number of bins in the table, the number of impact parameters recorded, and the number of muon events used to generate the table at each position. The bin size had been set early during the development phase, and bins close to

[14]Muons were rejected if the difference between their average and the generated impact parameter was either greater than 10cm or a factor of 0.002.

[15]The "critical energy" is the energy at which radiative energy loss begins to dominate over ionization. The critical energy for muons in water is roughly 800MeV. Muon energy losses are discussed in greater detail in 5.3.4

$\phi = 0$ only have a PMT centroid inside them in roughly one in nine events. It was decided for statistical reasons that each of the bins should have had a PMT centroid inside itself at least roughly 50 times. Given that near the center of the detector roughly half of the events fail the "non-straight line muon" cuts described above this means that at least 1000 events have to be simulated. As the impact parameter increases the number of detected photoelectrons (and the simulation time) drops rapidly. Given the speed at which they can be generated and the importance of fitting accurately near the edge of the fiducial area it was decided that 2000 events should be simulated at impact parameters greater than 800cm. With these settings it would take roughly two weeks of running on a grid of computers at Penn to generate tables at 100 distinct impact parameter values. Tables were generated at each 1% of the area covered by 850cm except for positions near 0cm and 850cm. Only two percent of the area contains impact parameters less than 120cm, and in the interest of not having to interpolate too far between points two extra table points were assigned inside that range. The most important region for the fitter is near the fiducial radius, so five additional tables were added between 800cm and 850cm, for a total of 108 distinct impact parameter values. The simulation was completed in slightly over two weeks.

A Perl script was used to combine the data from the simulations and convert it to input files for SNOMAN. In order to check that there was indeed enough information in each table the Perl script reported the minimum number of tubes appearing in any bin and the minimum number of detected photoelectrons in any bin. When all of the tables were examined the absolute minimum number of PMTs available in any bin was found to be 32. At 847cm and 850cm there were bins with no detected photoelectrons and at 844cm the minimum was 5 detected pe. For every table inside of 830cm all of the bins in every table had detected more than 10 pe. The Perl script recorded the following piece of data for each bin:

$$\text{Bin entry} = \frac{\text{bin detected photoelectrons}}{\text{total detected photoelectrons}} \Big/ \frac{\text{bin PMTs available}}{\text{total number of events}} \tag{4.19}$$

$$\lambda = (N_{fit})(\text{Bin entry}) \tag{4.20}$$

We can see from this equation that the uninstrumented "neck" of the detector should have little impact on the tables. In each event the neck will occult some position in $(\cos\theta, \phi)$ space, and there will be no contribution to either the numerator or the denominator of Equation 4.19. The value given by equation 4.19 describes the fraction of all detected photoelectrons that a PMT at the center of the bin would observe averaged over all possible muon orientations. If we multiply this value by a prediction for the number of photoelectrons that would be detected for this muon in an "average" detector (N_{fit}) we obtain equation 4.20, which is a prediction for the number of photoelectrons that a PMT located at the bin center is expected to detect for this hypothesized event. The reason it is useful to fit for the expected detected number of photoelectrons in an average detector orientation rather than for the particular orientation is that as the exit point of a muon track gets close to the neck (or a set of offline channels) the total actual detected number of pe drops dramatically. This reduces the correlation of geometry and light levels in the estimation of the likelihood, and in turn reduces the difficulty of minimization.

The second piece of information stored for each bin is the fraction of detected photoelectrons that are prompt, and these two data words are the only information that is recorded for each bin. The data file that contains the 108 tables as well as the bin size definitions is 28 Megabytes large and contains roughly one million data words. This is a relatively modest requirement which is easily accommodated in the memory budget of SNOMAN.

Problems with the Lookup Tables

The FTI lookup tables were generated using SNOMAN version 5_0086. Since that time some problems have been discovered with that version of the Monte Carlo. Two errors in particular are likely to have affected the lookup tables. The SNOMAN Monte Carlo uses a ported version of GEANT3 to handle muon interactions. It was discovered that the outgoing angle of bremsstrahlung photons was too large by a factor of the muon energy in MeV. Events with bremsstrahlung photons therefore often had a very large amount of light produced in the backwards direction. This unphysical backwards light can ruin the hypothesis that the muon is minimum ionizing and spoil the table. In hindsight a cut should have been placed on the maximum allowed gamma energy. The effect of this problem on reconstruction is discussed in section 4.3.8.

The second problem with version 5_0086 of the Monte Carlo has to do with delta-rays, or "knock on electrons". In an apparent effort to save simulation time a cut was placed on the minimum energy at which GEANT would produce electrons. The cut was placed at 2MeV, and removed many hundreds of electrons from typical events. Even with only 0.7MeV of kinetic energy electrons will produce enough Cherenkov photons that an average of about two photoelectrons will be detected. This unphysical cut resulted in a roughly 1% deficit in the amount of light produced. The outgoing angle of knock on electrons below 2MeV is large, so this light also had a disproportionately large effect on the PMT time residuals under the FTI hypothesis of a Cherenkov angle of exactly 42 degrees.

The lookup tables will be regenerated for future analyses now that these problems have been identified. In this thesis the lookup tables are used "as is", and we will see later that this leads to some bias in the reconstructed impact parameter in regions near bad table points. For the purposes of the analyses in this thesis the spacing of the FTI tables is narrow enough that this bias does not present a problem.

4.3.5 Interpolation Scheme

To generate predictions about the number of photoelectrons that a PMT is expected to detect it is necessary to perform two interpolations. First, in order to have a prediction that is continuous under rotation of the PMT positions it is necessary to move from a table with discrete values in each rectangle of solid angle to a continuous distribution. Second, since impact parameter is a continuous variable but the tables exist only at 108 particular impact parameters it is necessary to interpolate between sets of impact parameters.

The first interpolation, smoothing the individual tables to take into account the continuous PMT positions, is accomplished by means of a bi-linear interpolation. For each PMT FTI finds the cosθ and ϕ location of the PMT in the reference frame of the tables and then interpolates between the four bin centroids nearest the center of the PMT. To avoid problems at the edge of the tables the reflection symmetry was exploited by setting up the bins so that bin centroids exist at $\phi = 0$ and $\phi = \pi$. This is not possible, however, for the cosθ bins. In the interest of avoiding an extrapolation and for computational simplicity the bin edges at cosθ=0.99 and cosθ=-0.99 are treated in the bi-linear interpolation as if they were at 1 and -1 respectively. Both of the table values (the fraction of detected photoelectrons expected to be detected in the PMT and the fraction of detected photons that are prompt) are interpolated in the same way.

Now we turn to the more difficult problem of interpolating between different impact parameters. Over most of the solid angle the distribution in number of detected photoelectrons changes slowly with changing impact parameter, so the choice of interpolation scheme should have very little effect on the prediction. In the interest of computational speed a fast scheme should be chosen. In regions near the exit point of the muon, however, the distribution changes very rapidly, increasing greatly with increasing proximity to the muon exit position. This makes the regions where the distribution is changing rapidly easy to identify, as they are the regions with very large numbers of detected photons. The FTI routine first attempts a linear interpolation between the table predictions for the two nearest table impact parameters. If the linear interpolation results in a prediction of more than ten pe detected by the PMT then FTI performs a cubic spline interpolation between the four nearest table impact parameters. It was found that in some very rare cases the spline interpolation returned predictions for the number of detected photoelectrons which were far too small, and sometimes even negative. To avoid these rare cases a decision about which interpolation scheme to use is made after the spline interpolation returns its value for the predicted number of detected photoelectrons. In the case that the cubic spline interpolation returns a value of greater than 11 detected pe FTI uses the cubic spline prediction. In the case of between 9 and 11 pe the spline and linear interpolation are averaged, and in the case of fewer than 9 detected pe the linear interpolation is used.

The interpolation scheme can be tested by generating a prediction for the expected number of detected photoelectrons at an impact parameter for which a table exists, then removing that table and using the nearby tables to generate an interpolation prediction for that point. The interpolation prediction can then be compared to the table, and this gives an upper limit on how bad the error due to interpolation between table positions can be expected to be. This test was performed for the table at 465cm, which is bounded by tables 8cm away on each side at 457cm and 473cm. The distribution of interpolation prediction value divided by the table value for 8479 PMTs was found to have a mean of 1.8×10^{-3} and an RMS of 5.1×10^{-2}. Almost half of the tubes were expected to detect less than an average of one photoelectron, however, so it is important to check the consistency at higher light levels. After applying the requirement that the expected number of detected photoelectrons was at least five the distribution of the 1489 remaining tubes had a mean of 3.9×10^{-3} and an RMS of 3.1×10^{-2}. The variation in prediction due to interpolation is therefore observed to be significantly smaller than the variation in number of photoelectrons expected by Poisson fluctuations.

4.3.6 Likelihood Function

We are now in a position to write the likelihood function used in FTI to reconstruct through-going muon events. Before writing it out we will briefly review each of the elements. The FTI routine fits for six parameters, the entrance position of the muon (in polar co-ordinates), the direction the muon is traveling in (in polar co-ordinates), the time the muon entered the detector, and the number of photoelectrons that the detector would be expected to detect for this muon in an "average" orientation. A set of lookup tables is used to estimate λ, the expected number of detected photoelectrons for each PMT, and this number is used to calculate the probability that the PMT should have fired. With a model of the PMT single photoelectron charge distribution it is possible to generate charge distributions for *n* detected pe, and with these distributions and λ it is possible to calculate the probability a hit PMT would observe the charge it did. The lookup tables also contain the fraction of detected photons that are expected to be prompt. With this fraction and a set of simple timing PDFs for prepulse, prompt, and late or early light it is possible to assign a probability that the PMT would observe the time that it did. The likelihood function used by FTI is simply the product of all of these probabilities plus an exponential punishment to the log likelihood at the very edge of the detector area[16]:

[16]After processing had started the exponential punishment was found to have been applied incorrectly, as $e^{50(850-b)}$. Instead of being zero at 849cm and growing large at 850cm it is large at 849cm and shrinks to zero. A set of Monte Carlo muons were run to see if this had an impact on reconstruction, and it was found that it made no observable difference. In addition, because of the staggered order in which the processing occurred the exponential punishment was only applied to the data for the charge and time only versions of FTI, but not for the nominal version of FTI itself. Again, this is not observed to have any impact on reconstruction.

$$
\begin{aligned}
\mathcal{L} &= \left(\prod_{\text{all PMTs}} \sum_{n=0}^{\infty} P(n|\lambda) P(Q|n) P(t|\vec{x}_{PMT}, n) \right) \times H(b - 849\text{cm}) e^{e^{50(b-849)}} \\
&= \left(\prod_{\text{not hit PMTs}} e^{-\lambda} \right) \left(\prod_{\text{hit PMTs}} \sum_{n=1}^{\infty} \frac{e^{-\lambda}}{n!} \lambda^n P(Q|n) P(t|\vec{x}_{PMT}, n) \right) \\
&\times \quad H(b - 849\text{cm}) e^{e^{50(b-849)}} \qquad (4.21)
\end{aligned}
$$

where $H(x)$ is the Heaviside function, b is the impact parameter of the fit, $P(Q|n)$ is given by equation 4.6 and $P(t|n)$ is given by equation 4.18.

4.3.7 Likelihood Maximization

The FTI routine estimates the muon track parameters by finding the combination of track parameters that maximizes the likelihood function 4.21. The product of many probabilities is inevitably an extremely small number and most extrema finding routines are set up to find minima, so rather than maximizing the likelihood the FTI routine actually minimizes the negative logarithm of the likelihood function. When FTI was first being developed, undetected coding errors resulted in a likelihood space that was not particularly smooth with many local minima. The method of simulated annealing is particularly good at finding the global minimum in a bumpy likelihood space, so FTI was developed using this method. After the code was completed, it was found that the time spent in reconstruction was smaller than the time spent in generating Monte Carlo muon events, and given the small overall number of events there was no reason to switch to a faster minimization routine. The minimization was carried out using the *Numerical Recipes in Fortran* routine for simulated annealing using the downhill simplex method [55].

In the downhill simplex method a set of function arguments is modified until the function value for each set of arguments agrees within some tolerance. An *n*+1 by *n* array of function arguments is known as the simplex, where n is the number of arguments to the function. In the standard downhill simplex method at each step the routine selects the set of arguments that results in the largest (worst) functional value and attempts to replace it with a better set of arguments. The routine chooses the new set by means of reflections, extensions, or contractions of the position in argument space through the n-dimensional plane formed by the other simplex entries. If the value of the function at the new point is still larger than the value for all other simplex entries the routine tries again. Once the routine finds a point that is better than one of the old points it will keep this new point and repeat the above process by trying to find a replacement for the current worst point. The downhill simplex method is not a particularly favored minimization strategy both because it is not very fast and because it usually fails to find the global minimum if the function space has many local minima.

Simulated annealing is a modification of the downhill simplex method that is much better at finding global minima, even in rough function spaces with many local minima. Instead of always discarding the worst point in the simplex the simulated annealing routine may instead discard one of the other points. The probability that it discards another point is large at "high temperature" and decreases to zero at "zero temperature". The generation of the starting simplex is described below, and the "annealing schedule", the process by which the temperature is reduced, follows.

The first task in using the downhill simplex minimization technique is to generate a starting simplex, which generally consists of a seed fit and perturbations of that seed in each of the n function parameters. The FTI routine uses the fitted values from FTR as the starting seed for the entrance position, direction, and time. The initial estimate for the expected total number of photoelectrons detected[17] is approximated as 1.16 times the total charge in gain-scaled units measured in the event. A set of six other simplex points is then

[17] N_{fit} from Equation 4.20

generated by making minor perturbations of this first point. These perturbations can best be described when considering a muon in the special frame in which its equation of motion is given by Equation 4.1, $\vec{x}(t) = b\hat{x} + ct\hat{z}$. The first four perturbations are also shown schematically in Figure 4.13. The other simplex points are generated as follows:

1. Modified impact parameter (b $\rightarrow$ b + Gaussian perturbation with σ=30cm)

2. Rotation of the muon track about the x axis by Gaussian perturbation with σ=5^o

3. Rotation of the muon track about the y axis by Gaussian perturbation with σ=5^o

4. Rotation of the muon track about the z axis by Gaussian perturbation with σ=5^o

5. Alternate estimate for the detected number of photoelectrons (found by doing a quick one dimensional minimization of the negative log likelihood using Brent's method to find N_{fit})

6. Carry out all of 1-4 but with a 2cm impact parameter perturbation and 0.2^o rotations, and perturb N_{fit} by a Gaussian with σ=1000.

When generating new fits that have very different muon entrance positions or directions it is likely that the best fit entrance time is different from that of the seed. To make the search for the best fit position and direction more "productive" the entrance time is effectively taken out of the fit in the first stages of fitting. This is done by finding the mean time residual for PMT hits expected to be inside the Cherenkov cone based on the current fit parameters and then assigning the entrance time in such a way as to make the mean time residual zero. At the "high temperature" stages of the anneal the entrance time is assigned in this way, and at the "low temperature" stages the routine is free to adjust the entrance time on its own.

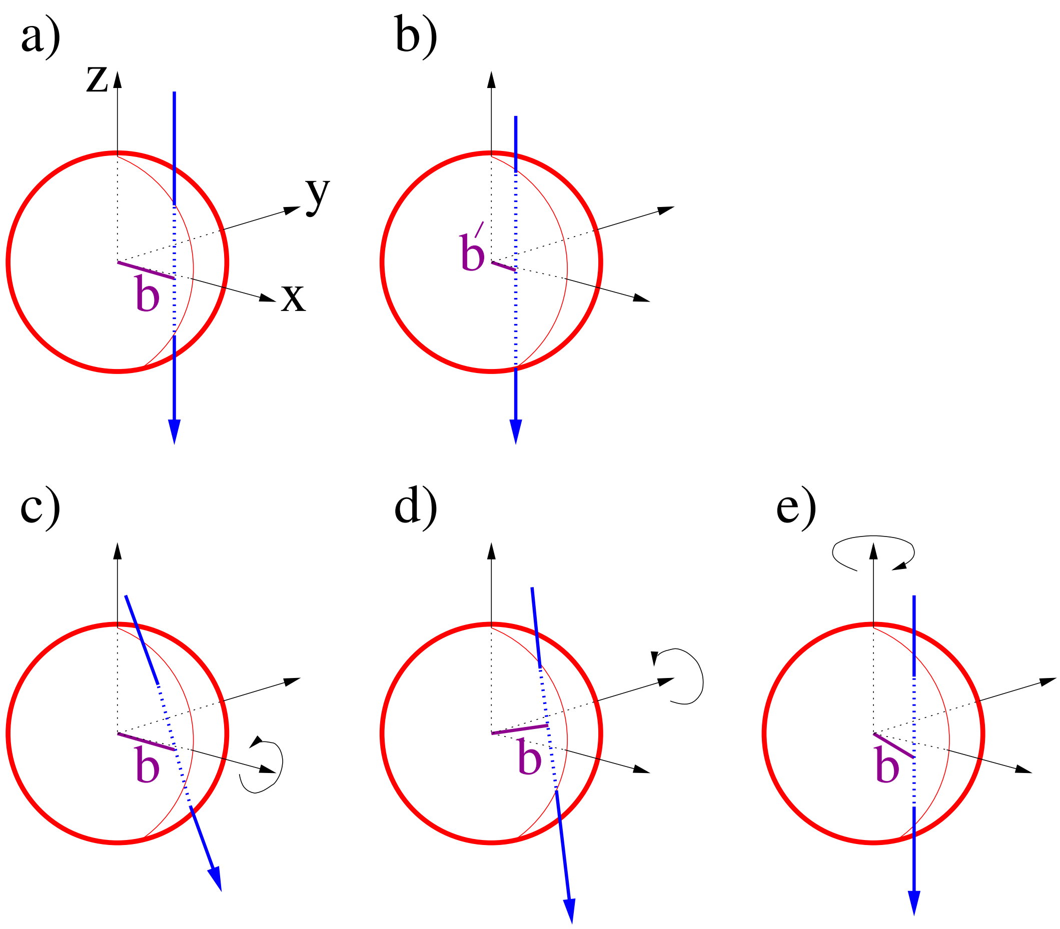

Figure 4.13: Schematic diagrams showing four of the six perturbations used in generating the simplex. In each diagram the muon's trajectory is shown as a thick blue line which is dotted when the muon is inside the SNO detector. The impact parameter is shown as a thick purple line and has the same value, b, in all but the top right diagram. The thin red line shows the great circle for which y=0 at all points.

a) Muon traveling in the special frame. It enters and exits on the great circle and its point of closest approach is at (0, 0, b).

b) Perturbation of the impact parameter from b to b'. The muon is still going directly down so it still enters and exits on the great circle.

c) Rotation of the track about the x axis. The impact parameter is the same as in (a), but the muon now enters to the left of the great circle and exits to the right.

d) Rotation of the track about the y axis. The impact parameter is the same as in (a) and the entrance and exit are on the great circle, but the point of closest approach is no longer at z=0.

e) Rotation of the track about the z axis. The impact parameter is the same as in (a) but both the entrance and exit are to the left of the great circle. The point of closest approach to the center is once again at z=0.

Annealing starts at a high temperature and is decreased linearly in 8 steps, with 30 new fit attempts at each temperature. After this phase the temperature is raised again, and is decreased quickly in 20 steps of 15 fit attempts. If at any time all of the simplex points have a negative log likelihood within a factor of 10^{-4} of each other the attempt at minimization is ended. Once the anneal has finished the entire procedure is repeated using the results from the anneal as a seed, rather than the FTR fit parameters. There is no generally prescribed method for the anneal schedule, the authors of *Numerical Recipes in Fortran* suggest that the user experiment with the technique to find what works best in any particular case. In the case of FTI it was found that as long as some time was spent at both high and low temperatures the chosen anneal schedule did not have much of an effect on the result of the fit.

The great strength of the method of simulated annealing is its ability to find global minima in a function space which is not smooth. During the high temperature phase of the search the simplex should be free to wander far from the starting seed. One upshot of this is that the result of the fit should be insensitive to the starting seed. The FTI minimization routine was tested by performing the fit on a set of simulated muon events, first with the FTR seed and later with the track information reported by the Monte Carlo. In the resultant fits there was no statistically significant difference between the events seeded with FTR and those seeded with the Monte Carlo.

4.3.8 Performance on Simulated Events

The performance of the fitter was tested on a Monte Carlo sample of cosmic muon events from the D_2O phase with statistics equivalent to about 85% of those expected in the D_2O phase. The muons were generated at a starting radius of 925cm with the energy, position, and direction distributions expected for cosmic ray muons. A total of 19598 events propagated successfully and 19574 (99.88%) of these were fitted[18]. A comparison of fit values to Monte Carlo values is shown in Figures 4.14 and 4.15.

The first set of figures shows the error in the fitted impact parameter b, the distribution of impact parameters (in $(b/850)^2$) themselves, and the bias in fit impact parameter as a function of Monte Carlo impact parameter. Although the impact parameter distribution does not have a noticeable slope, it is apparent that a flat fit of the squared impact parameter distribution to a uniform distribution would not return a good chi square. This is due to the unfortunate tendency of the fitter to clump events near the interpolation points and avoid areas between them, indicating that the interpolation scheme does not perform as well as expected. An additional problem is that the fitter avoids certain interpolation points, preferring instead to clump events near one of the adjacent points. The root of this problem is almost certainly the high energy bremsstrahlung photons with unphysical outgoing angles mentioned in subsection 4.3.4. However, these two effects are not expected to be significant in the subsequent analyses, which simply make use of a single cut on impact parameter. Despite the effects of clumping, the impact parameter resolution is actually quite good, and the reconstruction tails are fairly small.

The fitter angular reconstruction is shown in Figure 4.15. The vast majority of events reconstruct within a few degrees of the Monte Carlo direction. As can be seen from the right hand plot in Figure 4.15 most of the badly misreconstructed events have muon tracks that travel along the edge of the detector. Many of the muons that stop in the detector are also badly reconstructed. For through-going events with an impact parameter of less than 840cm only about one in 1000 events has a misreconstruction angle greater than 10^o.

[18]If the sample is reduced by rejecting muons which stop in the detector then 99.97% of the events were fit. With a further requirement that the Monte Carlo impact parameter be less than 845cm all of the events are fit

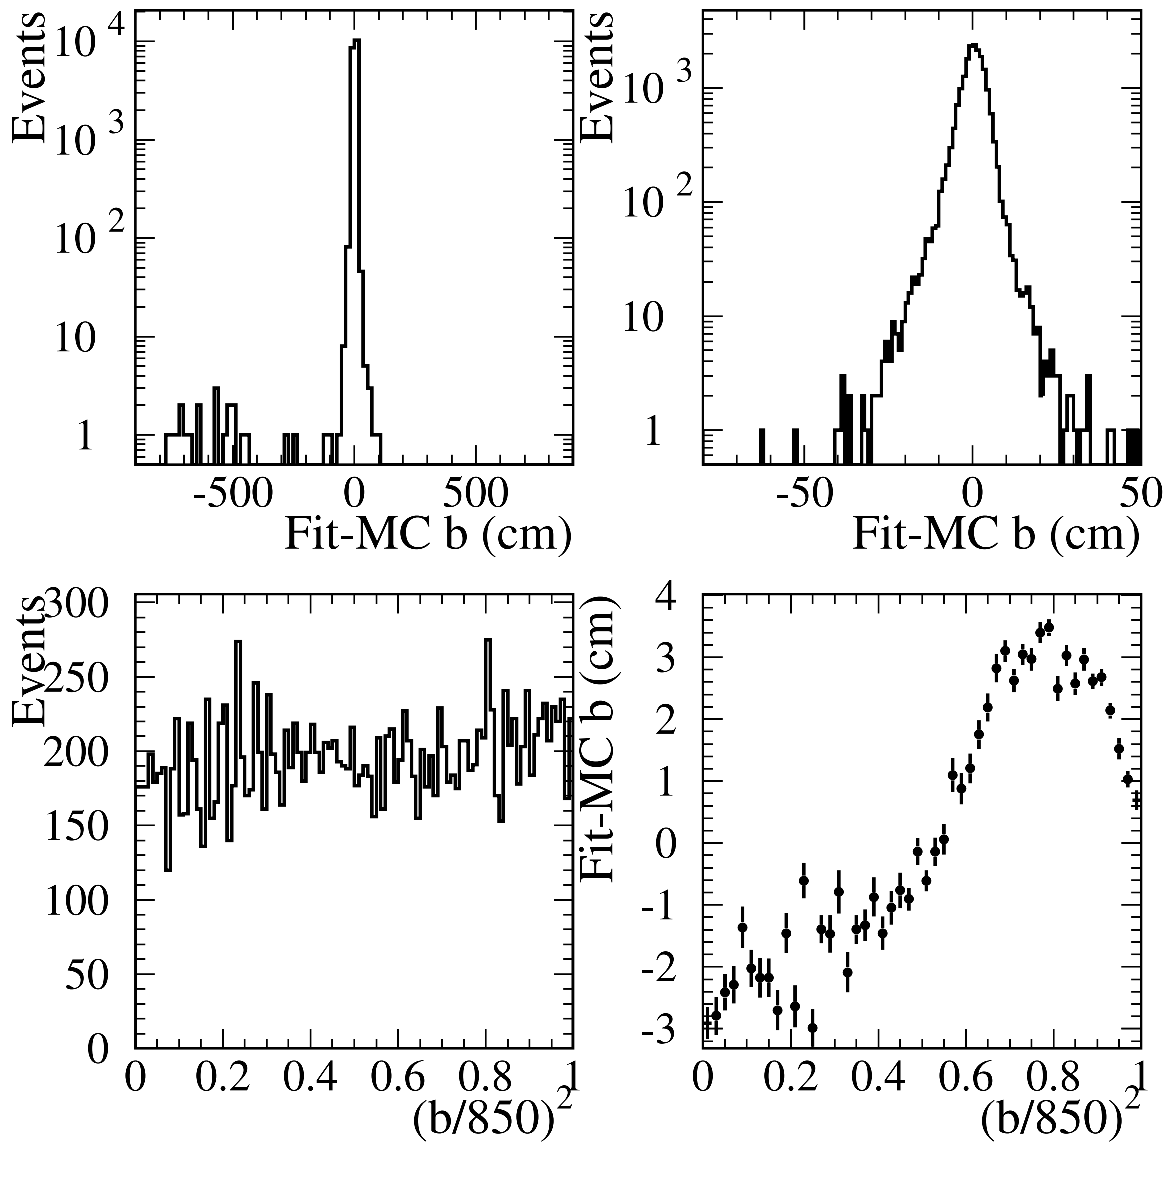

Figure 4.14: These figures show the FTI fitter performance on simulated cosmic muons. The top two plots show the gross and fine misreconstruction tails. The RMS of the peak is roughly 3cm. The bottom left plot shows the distribution of the square of the reconstructed impact parameter divided by the PSUP radius $(b/850)^2$. The large variation in bin height is due to the tendency of the fitter to reconstruct events near the interpolation points in the FTI tables. The plot at bottom right is a profile histogram of the fit impact parameter fit error as a function of Monte Carlo impact parameter. No data selection cuts have been applied, so these plots include both through-going and stopping muons.

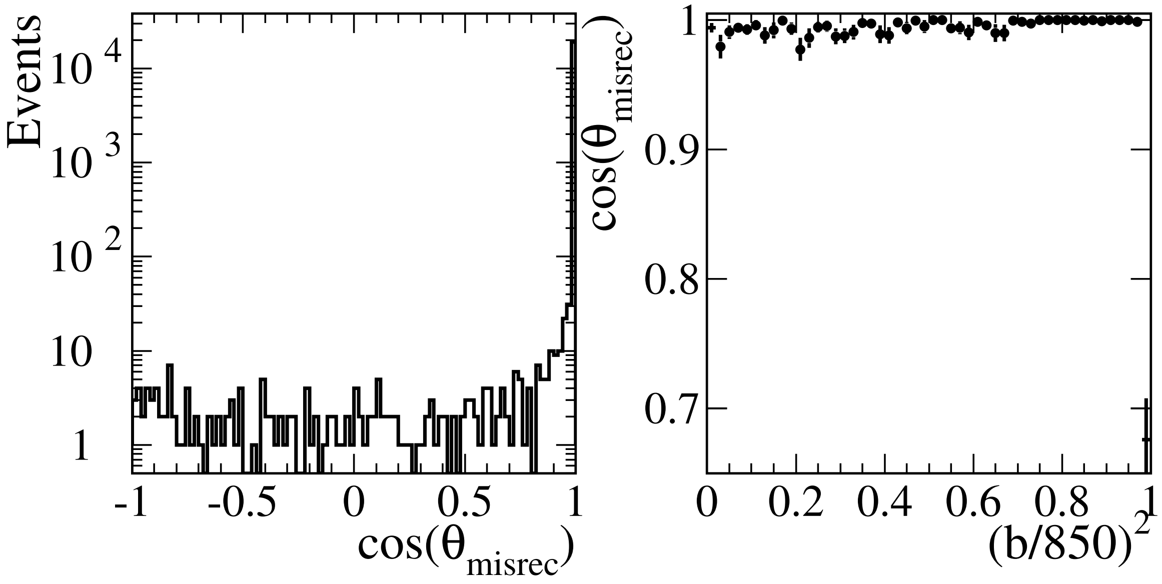

Figure 4.15: Errors in fitted direction for FTI fits to simulated cosmic muons. The left plot shows the distribution of the cosine of the angle between the fitted and Monte Carlo directions. Note that almost all of the data lies in the last bin. More than 80% of the events reconstruct within 1^o of the average direction from Monte Carlo, and 98.5% reconstruct within 10^o. The right hand plot is a profile histogram showing the misreconstruction vs. impact parameter squared. The poor fits correspond to muons which stop and events at the very edge of the detector.

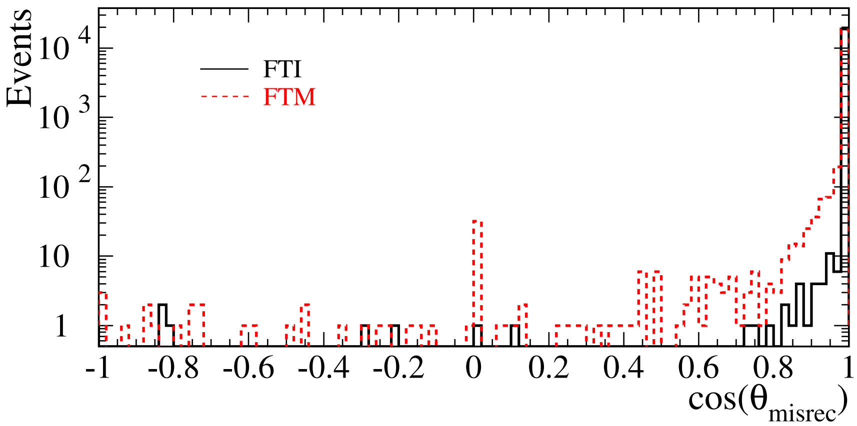

Figure 4.16: This plot compares the direction fit performance of the two muon fitters on simulated cosmic through-going muons out to a radius of 845cm. The peak in FTM at a value of 0 is made up of events in which FTM did not return a fit.

4.3.9 Comparison of FTI to SNO's Time Based Muon Fitter

A time residual χ^2 minimization muon reconstruction algorithm was developed for SNO several years ago by Nathaniel Tagg [52]. The "FTM" through-going muon fitter was found to have a large radial bias that increases with increasing impact parameter in studies of simulated events. A correction for the bias was estimated using Monte Carlo, and at 805cm this correction has a magnitude of roughly 50cm. Some of the radial bias was found to be due to the fact that FTM did not make use of a multiple photoelectron timing correction.

The FTM fitter has been run on the same simulated data as the FTI fitter, and FTM has been modified to use the multiple photoelectron timing correction described in 3.1.4. The FTM radial bias correction has not been updated to reflect the fact that its radial bias has changed, and for this reason the fit impact parameters are not compared here. A comparison of the direction fits for FTI and FTM on simulated muons is shown in Figure 4.16. Most of the data is in the peak near 1 for both fitters. While it is not shown here, the mean misreconstruction angle for events that are not badly misreconstructed is smaller for FTI than for FTM.

While FTI does do a better overall job at reconstruction it is not important from the perspective of the neutrino induced muon analysis. As long as events do not reconstruct more than roughly 20 degrees away from their true direction the zenith angle distributions of neutrino induced muons will not be badly distorted. The two important differences between the two fitters, for the neutrino induced analysis, are the fiducial area and its uncertainty. The recommended maximum impact parameter for the FTM fitter is 750cm, corresponding to 77.9% of the area of the full SNO detector. This compares to 830cm and 95.3% for FTI. The measurement of the zenith angle distribution of neutrino induced muon analysis will be statistics limited, so increasing the fiducial area is an obvious win. In addition, uncertainty on the fiducial area results in a proportional uncertainty on the flux, and the large radial bias observed in FTM would increase this uncertainty. Finally, the ability to run FTI in "charge only" and "time only" modes will help constrain the uncertainty on the fiducial area due to data vs. simulation differences.

4.4 Summary of Reconstruction

This chapter has outlined a maximum likelihood reconstruction algorithm for through-going muon events in the SNO detector. The spherical symmetry of the SNO detector is exploited by recognizing that after applying an appropriate rotation the only variable which distinguishes one muon from another is the impact parameter. Given a hypothesized track the algorithm uses a set of lookup tables to generate a prediction for the number of photoelectrons that each PMT should have observed. Using this prediction the algorithm evaluates the probability that each tube would observe its particular charge and time, and these probabilities are used to calculate the likelihood for the current hypothesis. The fitter uses a simulated annealing minimization to find the best values for the muon entrance position, direction, entrance time, and a detector state independent light output.

The FTI algorithm has been shown to perform extremely well. The mean distance between the fit entrance and the true MC entrance is roughly 10cm, which is smaller than the size of the SNO phototubes. The algorithm reconstructs the event impact parameters accurately all the way to the edge of the detector and reconstructs angle very well out to a radius of at least 840cm. The FTI algorithm uses the FTR fit as a seed, and it has been shown that the resultant FTI fits do not differ substantially from those obtained by using the actual Monte Carlo information as a seed. Large reconstruction errors occur predominantly for muons near the edge of the detector and for events in which the muon stops inside the detector. In the chapter on data selection will show that these misreconstruction tails are easily removed, leaving a very pure sample of muons.

Chapter 5

Data Selection

In the previous chapter the problem of reconstructing through-going muon events was discussed. The maximum likelihood event reconstruction algorithm known as FTI, which fits muon events by maximizing a six parameter function, was discussed. The FTI likelihood function made use of the PMT charge, timing, and hit information to fit for these six parameters: the muon entrance position, direction, time, and the total amount of light. In this chapter we will use the fit information from FTI to make decisions about whether or not any particular event under consideration is a through-going muon or some type of background.

5.1 Introduction

To make a measurement with a particle physics detector it is necessary to define a data set and then to extract the signal one is interested in studying from physics backgrounds and instrumental backgrounds. In establishing the data set one must first identify the periods of time during which the detector is online and in a known and acceptable state. In the SNO experiment the data acquisition system is used to define time periods known as “runs”, in which the detector state is known. Changes in the detector state are marked by changes in the run number. The first section of this chapter deals with the run selection procedure for muon analyses. The second section discusses the problem of separating through-going muon events from instrumental and physics backgrounds by using analysis cuts defined in software.

5.2 Run Selection

Over the course of the experiment the SNO detector was in a variety of different states, and for this reason runs are classified into a variety of types. “Calibration” runs were used to mark the periods in which a source was being used to calibrate some aspect of the detector, such as the attenuation length of the water or the energy response to mono-energetic gamma rays. Runs taken while the detector state was in flux due to electronics repair work or calibration setup activities were flagged as “maintenance” runs. The periods of time in which the detector was believed to be in its nominal state for detecting solar neutrinos were marked as “neutrino” runs.

The run type is set with the data acquisition software by the detector operator on shift at the start of each run. If the detector conditions change the operator is expected to start a new run, potentially with a different run type. For a variety of reasons some runs which were flagged as “neutrino” runs by the operator are not actually acceptable for analysis. Some of these reasons include:

- The run was a “test neutrino run” taken after major hardware work or a power outage to verify that the detector was again in its nominal state.
- The detector had not had its electronics re-calibrated since a major hardware change or power outage.

- The detector deck was at a higher than nominal temperature[1].
- A calibration source was in the detector.
- A member of the calibration group was in the deck clean room doing preparatory work for a calibration run.
- One or more crates were offline due to a hardware problem.
- Planned water system activities are likely to have created large amounts of light inside the detector during the run.
- A serious problem occurred with the synchronization of the detector's clock to the GPS satellites[2].
- The coils in the detector cavity which provide compensation for the Earth's magnetic field were offline.
- The detector operator indicates in the detector log that a serious problem which could affect the analysis occurred during the run.
- A serious problem (such as the HV in an electronics crate turning off) developed during the run but was not identified by the detector operator quickly enough.

In each of the three phases of the SNO experiment a run selection committee was established and charged with producing a list of the runs believed to be acceptable for analysis. In the D_2O phase of the experiment there was a large reluctance among the members of the

[1]When a fire occurs inside the mine all air handlers are turned off and underground personnel proceed to "refuge stations" which have a limited clean air supply. The detector can take data without anyone present, but without cooling the deck temperature quickly increases as heat is transferred in from the surrounding rock. The electronics calibration constants have been found to have some temperature dependence, so periods of time with high deck temperatures must be removed.

[2]While the muon analysis is insensitive to the time of day, a bad clock stamp in the run header information could cause problems when SNOMAN accesses time-specific files (such as calibration constants) from its database.

collaboration to use any of the detector data in making a decision about whether any run was selected, so the run selection committee used only the DAQ run header information and detector and data acquisition log information to make their decisions. In the salt phase of the experiment the run selection committee designed a SNOMAN processor known as SPS[3] (second pass selection) to aide in the selection process. This processor was further expanded by the NCD phase run selection committee. The SPS processor uses very low level information from the data, such as whether or not each PMT is hit during a twenty minute interval, to identify runs with non-optimal detector conditions. The run selection criteria for solar neutrino runs have already been well described by Wilson [56], so this section will simply highlight the differences between the muon and solar neutrino run selection criteria.

Many runs that are not appropriate for use in the solar neutrino analysis are acceptable for muon analyses. One obvious difference between the data typical of the solar and atmospheric signals is the number of channels that fire in an event. The highest energy solar neutrinos produce events that have roughly 150 NHit. The number of PMTs hit in a muon event is far larger. This can be seen in the scatterplot of NHit vs. impact parameter for simulated muons, shown in Figure 5.1. Thus, the solar and muon analyses are in completely different energy regimes, and therefore different run selection criteria may apply.

The energy estimators used in the solar analysis are very sensitive to the number of tubes that fire, so the presence of a small number of channels with much larger than normal occupancy[4] can have a big effect on the number of radioactive background events which leak over the energy threshold. The muon analysis, in contrast, is very insensitive to the number of channels that fire, so some periods in which tens of channels have higher than normal occupancy which were excluded from the solar neutrino analysis are included here.

[3]The SPS processor was written by Monica Dunford and Christopher Kyba at the University of Pennsylvania, and Wan Tseung and Jeanne Wilson at Oxford University

[4]Occasionally after a large seismic event in the mine a handful of channels begin to "scream", firing in a large fraction of events. These channels eventually calm down or are turned off, but on a few occasions several days have passed before this has happened.

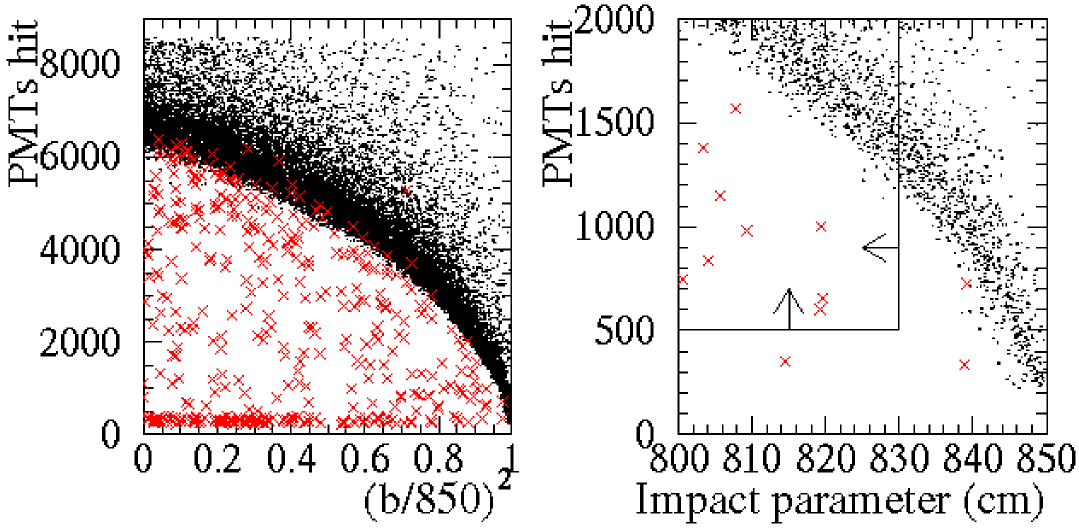

Figure 5.1: The left hand scatterplot shows the number of PMTs which fire in simulated cosmic muon events vs. simulated impact parameter divided by the detector radius squared for all impact parameters. Through-going muons are shown as small black dots, stopped muons and decay electrons are shown as red x's. The plot on the right is a blowup of the region near the edge of the detector, linear in impact parameter. The two lines show the minimum number of hit PMTs and the fiducial radius for the analysis. Note that even at 840cm the muons are still above the cut on the required number of PMTs.

A very closely related issue is the performance of the coils in the detector cavity which provide compensation for the Earth's magnetic field. Photomultiplier tubes lose gain and efficiency when they are in a magnetic field, and therefore changes in the field at the PMTs have an effect on the detector energy response. At a few points in time one or two of the 14 coils were not running at their nominal current, or were off. Runs taken during these time periods were rejected from the solar neutrino runlist, but since small localized changes in PMT gain are dwarfed by the event to event variation in muon light output these runs are included on the muon runlist.

Another problem which affects solar but not atmospheric analyses is the presence of excessive radioactive backgrounds in the H_2O or D_2O. Radioactivity can be introduced unintentionally as a result of a leak which allows mine air into the water, intentionally as a "containerless" calibration source [57], or incidentally as a result of calibration activities[5]. As mentioned above, the solar neutrino analysis is very sensitive to the number of radioactive background events with measured energies above the analysis threshold, so these periods are cut from the solar neutrino analysis. The presence of additional low energy backgrounds has no impact on muon analyses, so these runs need not be rejected.

The events in the solar neutrino analysis are far enough over the trigger threshold that trigger efficiency has thus far been ignored in the analysis. In cases when the trigger system is having problems, however, this assumption may no longer be valid. For this reason runs in which any electronics crate is not contributing to the trigger sum are not used in solar analyses. However, it is clear from Figure 5.1 that the amount of light produced in muon events is large enough that even with the trigger signals from several crates not contributing to the trigger sum the analysis can assume 100% efficiency for muon detection at the trigger threshold of roughly 20 hits.

The previous four paragraphs describe the most common reasons why runs which were not used in past solar neutrino analyses have been included in this analysis. A number of runs were added for additional reasons, and while they are not detailed here it is hoped that the preceding paragraphs give the reader a good indication of the types of considerations that went into making muon run selection decisions. The extra runs which were included in this analysis were subjected to the same scrutinies as the runs used in the solar analyses. With the exception of effects which were attributed to the unusual circumstance for which these runs were under consideration, these runs pass all of the regular selection criteria. The additional runs add 586 hours to the D_2O data set and 2422 hours to the salt data set, bringing the total live time before cuts to 337.3 days and 499.5 days respectively. The runs used in through-going analyses in this thesis are listed in Appendix D.

[5]During the salt phase Na nuclei would became activated by gammas from the ^{16}N calibration source.

5.3 Event Selection

Now that a runlist has been established we turn to the problem of extracting the signal from any potential backgrounds. In general, the difficulty of this problem and the techniques for isolating the signal vary greatly depending upon the detector which is making the measurement. Background events may be rejected with hardware, by setting up a sophisticated detector trigger, or with software, by using analysis cuts. While it is common in particle physics experiments to make use of all techniques, this analysis uses software rejection exclusively.

The combination of very low radioactive backgrounds in SNO and robust analog electronics allowed the SNO experiment to use a trigger which is conceptually very simple. The experiment makes use of several different triggers, but for the purposes of this analysis it is only necessary to consider one, the NHit trigger. The NHit trigger works by keeping track of how many PMTs have fired in a 93ns window. If more than roughly twenty PMTs fire, the event is recorded. As can be seen in Figure 5.1, many thousands of PMTs fire when a muon passes through SNO, so the trigger inefficiency can effectively be ignored for this analysis.

In the next subsection, the relevant backgrounds to through-going muon analyses are discussed. Next, a series of low level analysis cuts and high level analysis parameters used to separate these backgrounds from the muon signal will be described. An estimate for the number of instrumental background events in the neutrino induced muon signal box is presented, and finally the acceptance of the cuts on the cosmic muon signal is estimated using a Monte Carlo technique.

Technical note: Due to a technical glitch which delayed their processing, six of the runs on the muon run list were excluded from the data studies in this chapter and the chapter which follows. The exclusion of these runs is not expected to have any impact on any of the results obtained. For future SNO reference purposes, these runs were 10534, 10536, 10549, 10636, 10675, and 11269.

5.3.1 Instrumental Backgrounds for Through-Going Muon Analyses

The presence of instrumental backgrounds or physics events which mimic the through-going muon signal could cause an overestimation of the neutrino induced muon flux, and a consequent reduction in sensitivity to the neutrino oscillation parameters. There are four instrumental backgrounds in SNO which are known to produce large amounts of light and are therefore a potential background to the muon signal: Wet end PMT connector breakdown (WEB), PMT spontaneous light emission (flashers), light generated above the detector (neck events), and light associated with one of the water systems (bubblers).

During the D_2O and salt phases of the experiment roughly three PMTs suffered wet end failure per week, and during their failure they often produced a very large amount of light. These WEB events usually occurred in bursts, so a simple instrumental cut based on the number of large events in a short amount of time will remove many of them. However, since WEB events do not always occur in bursts a burst cut alone is not adequate. The light from WEB events is produced at the PMTs, whereas the light from muons is generated along a track. It is reasonable, therefore, to expect that a cut based on the PMT time residuals under a fit hypothesis will have a good efficiency for rejecting WEB events. Flashers are quite similar to WEB events, in that they are produced at the PMTs, but flashers generally produce considerably smaller amounts of light. If the flasher energy spectrum extends into the muon signal region, the same algorithms designed to reject WEB should work on flashers.

The so called "neck events" have a spectrum ranging from low NHit events which are a background to the solar neutrino analysis, to events with light levels comparable to muons [58]. In order to reject neck events a set of four PMTs were installed, two in the neck and two in the “glovebox” interface high above the detector. These veto tubes are extremely efficient at removing neck events. In addition, neck events, and especially those events which produce a large amount of light, tend to occur in bursts. Thus they will be easily rejected with a burst cut.

The water levels in the SNO detector are monitored by a system that as an unintended by-product produces air bubbles under water. In the earliest commissioning days of SNO, a detector operator named Nally noticed several events that at first glance looked like identical upward through-going muons in quick succession. The source of the light was quickly identified as being the “bubbler” system, and as a consequence the system was henceforth turned off during data taking. It has been observed that even when the bubblers are turned off there are still occasional single “bubbler events”, so it is important that analysis cuts be developed to reject them. While at first glance the bubbler events resemble muons, a closer look reveals that these events lack the cluster of early time tubes typical of the entrance point of a muon. Furthermore, the air bubble (and light) are produced at the acrylic vessel, so because of solid angle and timing effects the region with the earliest PMTs also has the highest charge PMTs; a very different pattern than that produced by muons. When the bubbler system is operating it produces a great number of events, so it is very easy to produce a data set of bubbler events to check that the cuts are efficient.

The final, and probably the most common instrumental background to muon analyses is electronic pick-up events. The PMT discriminators are sensitive to high frequency noise induced on the PMT cables or on the front end electronics, and may fire in response to induced signals. Two common sources of pick-up events are physical disturbances on the deck, such as the deck door slamming shut, and hardware failures inside the electronics racks. High frequency signals integrate to zero charge, so pick-up events have a very small ratio in the summed charge to the total number of tubes hit. This makes these events easy to reject with charge based cuts, as will be shown in subsections 5.3.3 and 5.3.4[6].

5.3.2 Physics Backgrounds for Through-Going Muon Analyses

The backgrounds discussed in the previous subsection are all produced by point sources of light, with the exception of electronic pick-up. The muon signal is distributed light produced along a particular trajectory, and this difference makes rejecting instrumental backgrounds relatively easy. Some physics events, however, have considerably more in common with the through-going muon signal. Neutrino interactions which take place inside of the PSUP can result in relativistic charged particles which produce Cherenkov light. These events are called "fully contained" and "partially contained" in Super-Kamiokande analyses, and will be referred to as "internal neutrino interactions" or "neutrino interactions inside the PSUP" throughout this thesis. Some of the muons which enter the detector stop inside of it rather than traveling all the way through, and these are by definition background to a through-going muon analysis. It will be shown that most of these events may be removed by estimating the energy loss per unit track length (dE/dx). Muons traveling through water lose a minimum of $\sim$200MeV/m, so events which are measured to have produced less than this are either very poorly reconstructed or are not muons at all.

Cosmic muons are a background to the atmospheric neutrino flux and oscillation analysis. Such muons may be rejected by means of a muon direction cut; the efficiency of such a cut is very closely related to the acceptance of badly misreconstructed events. "Multiple muon events", in which at least two muons from the same air shower enter the detector, may be considered a background to cosmic muon analyses, since FTI is not designed to reconstruct them and, is not able to identify them as double muons.

[6]For reference, the relevant plots are on pages 126 and 136.

5.3.3 Low Level Analysis Tools for Data Selection

Several techniques for the identification of muons have been developed. These techniques are assigned to one of two groups, depending upon their complexity. The simple techniques are referred to as "low level cuts" in this document. More complex techniques which make use of the physical principles associated with muons are referred to as "high level cuts". The low level cuts will be discussed first.

NHit Cut

Muons traveling through the detector produce a great deal of light, so it is reasonable to reject events in which only a small number of PMTs register hits. The NHit distribution of Monte Carlo muons was shown in Figure 5.1, along with the chosen cut value of 500 PMTs hit. For the purposes of this cut only online tubes with good calibrated time and charge measurements are considered when counting PMTs.

Retrigger Cut

Any event that occurs within 5μs of the previous event is tagged as a retrigger. On occasions when PMTs suffer high voltage failures in SNO it is common for a run to end with hundreds or thousands of retrigger events that each produce large amounts of light. A great deal of CPU time would be required to run the analysis software on these obviously bad events, so retrigger events are not reconstructed. Events are required to have an FTI fit to be analyzed so all retrigger events are excluded from all muon analyses. Muons are expected to arrive at random times relative to other events in SNO, so the reduction in muon acceptance due to this cut can be simply thought of as a reduction in the live time of the detector. Note that because most muon events are followed by one or two retrigger events this means that if there was a source which produced muons correlated in time within $\sim 1\mu$s of each other this analysis would be blind to them. Given the few hundred square meters area of SNO and the surface energies which are required for muons to reach SNO such a source would be easily observable by large area surface cosmic ray experiments. The retrigger cut removes 25.7s from the D_2O data set, which is smaller than the uncertainty in the live time itself.

Burst Cut

It was noted earlier that many of the instrumental backgrounds tend to occur in bursts. The open muon data set from the D_2O phase was used to study bursts of large NHit events. On the basis of this study it was decided that if four events with greater than 250 NHit occur within two seconds, the time region should be tagged as a burst. Events which failed the retrigger cut were not considered for the purposes of counting the number of events with over 250 NHit because muons often produce a large number of retrigger events with large NHit. When a burst of events is found, all of the events within the time frame are cut. Once again, the loss of acceptance due to the burst cut will be considered a reduction in the live time rather than a loss of acceptance. It is worth pointing out that the time window cut out by a burst is larger than the time between the first and last event. This is because if an event had occurred before the burst, but within two seconds of the third event, it would have been considered part of the burst. The total amount of time cut out of the D_2O data set due to this burst cut is 7405s [59], or 0.025%.

Neck Tube Cut

To reject the so-called neck events, any event in which all four neck and glovebox tubes fire is rejected.

Raw Timing Distribution Cut

Some instrumental background events have been observed to have very flat PMT hit time distributions. One measure of the flatness of the timing distribution is the RMS of the raw[7] PMT hit times. Based on Monte Carlo studies of cosmic muons, a cut requiring that the RMS of the hit times be less than 38ns was set. A plot of the hit time RMS for simulated cosmic muons is shown in Figure 5.2. Events from the open D_2O data set which failed the burst cut, and simulated events in which an atmospheric neutrino interacts in the water inside of the PSUP are shown for comparison.

Raw Charge Distribution Cut

When muons exit the detector they tend to deposit a very large amount of light in the tubes nearest the exit point. More isotropic events, in contrast, tend to have a more uniform charge distribution. The RMS of the charge distribution is therefore larger for muon events. Based on Monte Carlo studies, events with a charge RMS of less than 4.5 gain scaled units are rejected. A plot of the charge RMS for simulated cosmic muons is shown in Figure 5.2, again with burst events and simulated internal atmospheric neutrino events for comparison.

5.3.4 High Level Analysis Tools for Data Selection

Fiducial Area

While in some detectors the entire area or volume may be useful, it is more common to define a ”fiducial volume” (in the case of a through-going muon analysis fiducial area) outside of which events are excluded from the analysis. This is usually done because the events become systematically different near the detector edge. The solar neutrino analyses in the SNO detector, for example, have typically defined a fiducial volume by including

[7] Ie., simply the calibrated times, not the time residuals found under a track hypothesis.

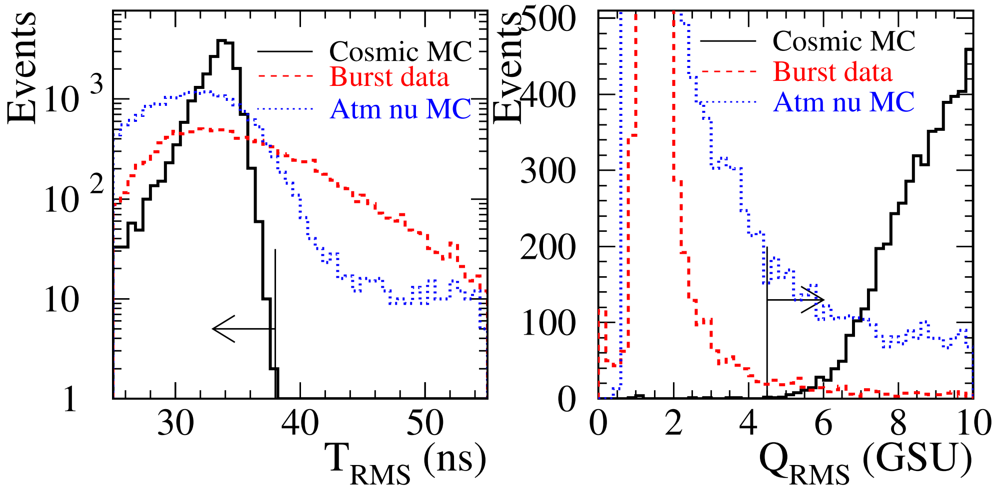

Figure 5.2: The left hand plot shows the distribution of the RMS of the calibrated PMT hit times in simulated cosmic through-going muons. Note that this is the RMS of the raw times, so this is not the same as the distribution of time residual RMS from a fit. Events which fail the burst cut and simulated neutrino interactions internal to the PSUP are shown for comparison. The right hand plot shows the distribution of the RMS of the observed charges (shown in gain scaled units). Note that the distribution for cosmic events continues to the right of the charge plot, and to a smaller extent to the left of the time plot. Both of the cuts are one-sided.

only events within 550cm of the detector center (77% of the D_2O volume). This radius was chosen for several reasons: because the distortion of light by the acrylic vessel is greater near the vessel; to reject events associated with the acrylic vessel itself; to avoid including events produced by radioactivity in the light water; and finally because the calibration sources cannot reach large radii, which would make it necessary to rely upon extrapolation to calibrate the energy in those regions. Muon events are less sensitive to effects associated with the acrylic vessel, but as seen in figure 4.15 the muon reconstruction algorithm begins to have trouble reconstructing muons with very large impact parameters. In addition, some of the high level analyses described below start to fail at large radii, and many of the backgrounds to the muon analyses tend to "pile up" near the PSUP.

The neutrino induced through-going muon analyses will be statistics limited, so it is desirable to set the fiducial area as large as possible without introducing significant systematic uncertainties. The acceptance of muons changes rapidly at impact parameters greater than 840cm, so it is desirable to set the fiducial area so that the uncertainty on the acceptance is not large. Based on Monte Carlo studies of through-going muon reconstruction and the analysis selection tools described below a fiducial radius of 830cm was chosen. This corresponds to a fiducial area of 216.4m^2, more than 95% of the total detector area. For comparison, the fiducial area of the Super Kamiokande detector is $\sim 1200\text{m}^2$ [53]. In SNO the fiducial radius cut at 830cm corresponds to a minimum muon track length of roughly 367cm; Super Kamiokande defines their fiducial area by applying a muon track length cut of 7m.

Fit Number of Detected Photoelectrons

One of the parameters of the FTI fit is the number of photoelectrons the detector would observe in an "average" orientation. This can be thought of as a fit to the number of Cherenkov photons produced by the muon, scaled by the appropriate detection efficiency for photons produced at the given impact parameter. This quantity is used rather than the sum of the PMT charges as a cut for the following reason: A very large fraction of the total light detected in a muon event is detected in the region where the muon exits. If the muon exits in a region with no active PMTs[8] then the charge sum will be suppressed. Further, since the number of online tubes varied throughout the course of the experiment the mean total charge sum also changed. The FTI fit number of photoelectrons should be insensitive to these changes. Based upon a study of simulated muons it was decided to reject all events for which the fit number of photoelectrons is fewer than 2000. The number of detected photoelectrons is highly correlated with impact parameter, and at 830cm simulated through-going muons always result in a fit number of photoelectrons larger than 3000. The cut value of 2000 was chosen so that the acceptance for muons at 840cm would still be

[8]For example the detector neck, or through a region where the PMTs are temporarily offline.

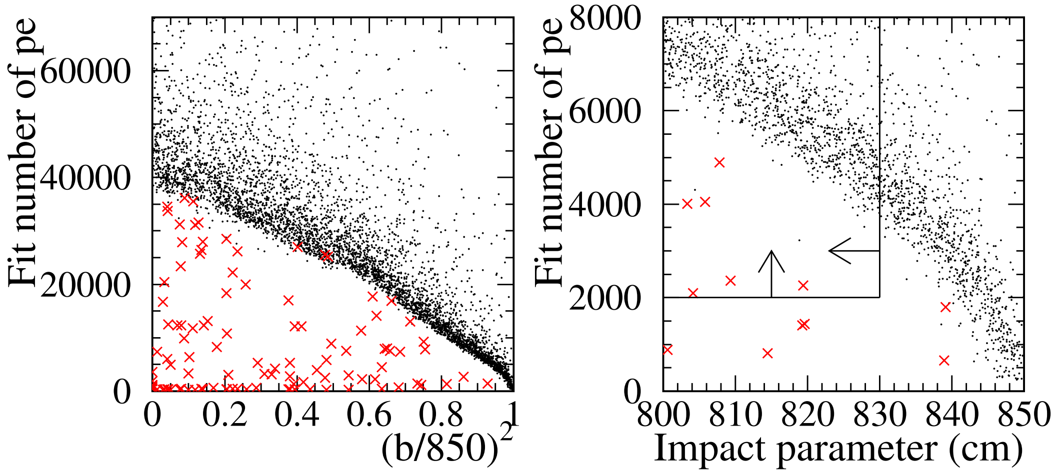

Figure 5.3: The left hand scatterplot shows the number of photoelectrons which the FTI fit returned for simulated cosmic muon events vs. simulated impact parameter squared for all impact parameters. Through-going muons are shown as small black dots, stopped muons and decay electrons are shown as red x's. The data extend well past the maximum value of the ordinate, those data are not shown to make the structure more visible. The plot on the right is a blowup of the region near the edge of the detector, linear in impact parameter. The two lines show the minimum number of fit photoelectrons and the fiducial radius for the analysis. Note that even at 840cm the muons are still above the cut on the required fit number of photoelectrons. Cf. Figure 5.1.

close to one. It is a guiding principle of this thesis that the high level cuts should have a large acceptance out to at least 835cm.

Fraction of PMTs in the Cherenkov Cone That Fire

In through-going events almost all of the PMTs which are inside the muon's Cherenkov cone fire. This makes the fraction in the cone a powerful test statistic, and in this thesis it will be referred to as the RICH parameter. The RICH parameter is sensitive to muons that stop in the detector, as the Cherenkov cone ends earlier than expected. It is also sensitive to most instrumental backgrounds, which tend not to have full occupancy in the expected muon cone[9]. There is a problem with RICH that arises from the "concentrator cut off".

[9]This test statistic is of course completely insensitive to any background event in which every tube in the detector fires.

This cut off reduces the fraction of PMTs in the cone which fire at large muon impact parameters. Due to their design, the concentrators have no noticeable effect on the RICH fraction up to impact parameters of roughly 550cm. From 550 to 700cm the RICH fraction decreases from a median of 0.97 to 0.95, and then begins to fall more rapidly, to 0.7 by 820cm. This drop off is partially mitigated by introducing a new statistic in which PMTs illuminated by the first and last 50cm of the muon track are ignored. The cut on the trimmed RICH statistic, or TRICH, is set at 0.7. A scatterplot showing how TRICH changes as a function of impact parameter is shown along with a histogram showing the sensitivity of TRICH to stopping muons in Figure 5.4. It can be seen from the right hand plot that by using a cut on TRICH the badly misreconstructed stopped muon events can be removed. It should be noted that the RICH parameter was originally developed for SNO by Nathaniel Tagg [52] who called the parameter “FShould”. Tyron Tsui from UBC adapted Tagg’s code and wrote a routine able to calculate RICH and TRICH for any muon reconstruction algorithm.

Cherenkov Cone Time Residual RMS

In many previous experiments, muons have been reconstructed by defining a fit χ^2 based on the difference of each PMT’s time residual from the event t_0. At the end of the fit the χ^2 itself is often used as a figure of merit. While the FTI reconstruction algorithm does not explicitly minimize the PMT time residuals, one can see from Equation 4.16 that FTI assigns a large probability to PMTs with fit time residuals close to zero. For this reason a large time residual RMS will be reflected in the likelihood as a poor fit. Because many of the instrumental backgrounds are point sources of light, the RMS of their fit time residuals tends to be quite large, and they can be rejected by using a cut on the maximum allowed time residual. While the time residual RMS has fairly good discrimination against background events it does not work particularly well as a fit figure of merit. Time residuals

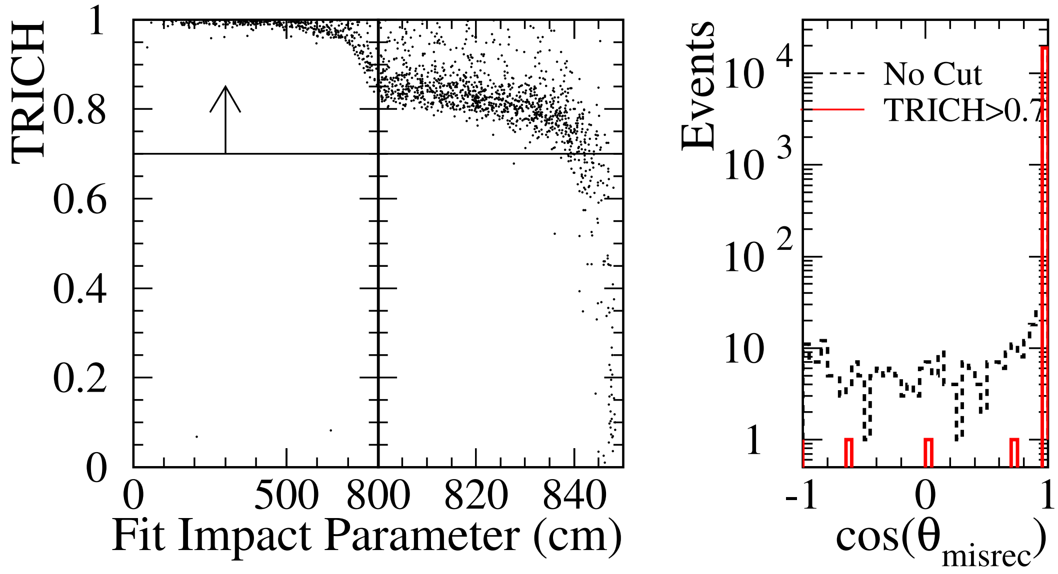

Figure 5.4: The two adjoined scatterplots on the left show the variation of the TRICH parameter with impact parameter for simulated cosmic through-going muon events. Note that the scale of the abscissa changes at 800cm to allow the variation at large impact parameter to be more easily seen. The dots near the bottom left of the figure are two misreconstructed events. The simulated impact parameter for that event was very close to 850cm. The right hand plot shows the cosine of the angle between the true and fit muon directions for simulated muons (cf Figure 4.15). The dashed black histogram shows includes all muons, and the solid red histogram shows only those muons for which TRICH is greater than 0.7.

are only defined inside of the cone, and if the fit is at the edge of the detector and only includes a handful of channels the time residual RMS can easily be small. The in-cone time residual RMS (CRMS) is used in a linear discriminant analysis to reject background events, and will be described in 5.3.7. These two parameters are shown in Figure 5.5.

Cherenkov Cone In-Time Ratio

A statistic which is closely related to the previously discussed Cherenkov cone time residual RMS is the ratio of hit PMTs in the Cherenkov cone that are "in-time" (CITR). The number of PMT time residuals which fall within a time window of 5ns around the fit event time is divided by the total number of hit PMTs inside the Cherenkov cone. Because the time window is quite large compared to the fit time residual RMS for simulated muons, this fraction is close to one for a real muon. It is considerably smaller for instrumental backgrounds which produce light at a point rather than along an extended track. The distribution of the CITR parameter for simulated cosmic muons is compared to that for events which fail the burst cut in Figure 5.5. Based on the simulated muon events it was determined that events with a CITR value of less than 0.85 would be cut.

Estimate of Event dE/dx

Through-going muons produce light as they lose energy along their track. Under the through-going hypothesis muons which stop in the detector, in contrast, have a section of "track" in which no light is produced. Similarly, most internal atmospheric neutrino interactions produce particles which travel short distances compared to through-going muons. The amount of light observed per unit track length is therefore a powerful discriminant against instrumental backgrounds, atmospheric neutrino events with vertices inside of the PSUP, and stopped muons. In this subsection we will pursue an estimate of the muon's energy loss per unit distance (dE/dx) for use as a data selection tool. As the main goal is

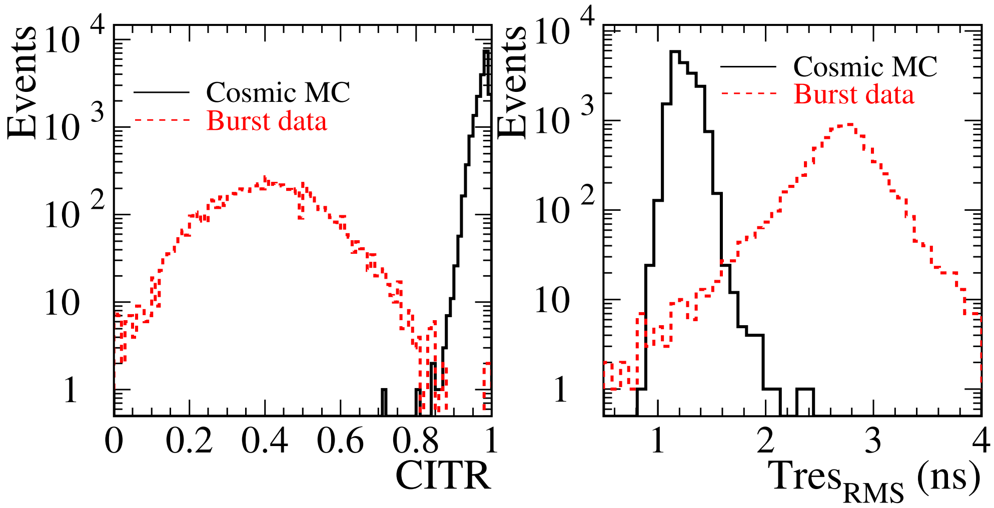

Figure 5.5: These two plots show the distributions of the Cone In-Time Ratio parameter and the RMS of the muon track time residuals for both simulated muons (black solid line) and events which fail the burst cut (dashed red line).

rejection of backgrounds it is not essential to estimate dE/dx with great accuracy; for this purpose an estimate that is roughly proportional to the true dE/dx is sufficient. Before we go into details about how the muon's energy loss per unit track length is estimated, we will briefly review muon energy loss in water.

A nice overview of muon energy loss in matter is given by Gaisser [25], and will briefly be summarized here. Relativistic muons lose energy as they propagate through matter by four main processes. These processes are ionization of the surrounding material, direct pair production of electrons, bremsstrahlung, and photonuclear interactions (hadroproduction). At the low energies typical of neutrino induced muons in SNO the ionization losses dominate, but the ionization loss grows with the log of muon energy, while the other three are all directly proportional to muon energy. The energy loss rate for muons can therefore be written approximately as:

$$\frac{dE}{dx} = \frac{dE_i}{dx} + \frac{dE_r}{dx} \tag{5.1}$$

$$\frac{dE}{dx} = -\alpha_i - E/\xi \tag{5.2}$$

where α_i is the average ionization loss (∼2MeV/cm in water), and ξ depends on the material properties (mainly Z) of the matter the muon is traveling in. The “critical energy” of a material is defined as the energy at which radiative losses begin to dominate over ionization, $\epsilon = \alpha_i \xi$. The critical energy of solids and liquids is given by the particle data group (2006) [60] to within 3-4% as:

$$\epsilon = \frac{5700 \text{ GeV}}{(Z + 1.47)^{0.838}} \tag{5.3}$$

For muons in water ϵ ∼800GeV. About 10% of cosmic muons entering SNO have energies greater than this, and only about one in one hundred of the neutrino induced muons at SNO are above the critical energy.

The rate of Cherenkov photon production ($\frac{dN_c}{dx}$) as a muon passes through SNO is the sum of the rate of Cherenkov photon production from the muon itself ($\frac{dN_{c\mu}}{dx}$) plus the rate of Cherenkov photon production from any charged particles produced by radiative energy loss ($\frac{dN_{cr}}{dx}$):

$$\frac{dN_c}{dx} = \frac{dN_{c\mu}}{dx} + \frac{dN_{cr}}{dx} \tag{5.4}$$

Bremsstrahlung and pair production produce secondary relativistic electrons in SNO, and the total number of Cherenkov photons produced by such secondary particles is roughly proportional to the energy of the electron. This means that the rate of Cherenkov photon production from particles other than the muon is proportional to the muon’s radiative energy loss,

$$\frac{dE_r}{dx} = \kappa \frac{dN_{cr}}{dx} \tag{5.5}$$

where $1/\kappa$ is a constant describing the number of Cherenkov photons produced per MeV of secondary electron energy. This means that if we understood the rate of Cherenkov photon production from muons we could estimate the energy loss of the muon, as

$$\begin{aligned}
\frac{dE}{dx} &= \frac{dE_i}{dx} + \frac{dE_r}{dx} \\
\frac{dE}{dx} &= -\alpha_i - \kappa\frac{dN_{cr}}{dx} && (5.6) \\
\left|\frac{dE}{dx}\right| &= 200\text{MeV/m} + \lambda(\frac{dN_c}{dx} - \frac{dN_{c\mu}}{dx}) && (5.7)
\end{aligned}$$

So if one could estimate how many Cherenkov photons were produced by the average muon event it would be possible to estimate any individual muon's dE/dx.

The FTI routine does not fit for the number of Cherenkov photons produced, but it does fit for a quantity that is close enough to make dE/dx estimation possible. As mentioned in 4.3.1 one of the parameters in the FTI fit for each event is the number of photoelectrons that would be detected with the SNO detector in an "average" orientation. This parameter takes out the effects of offline tubes and the position of the detector neck, so it is related to the number of Cherenkov photons produced multiplied by some impact parameter dependent scale factor. For any given impact parameter it is possible to write an equation similar to 5.7 but using detected photoelectrons rather than produced Cherenkov photons.

$$N(\frac{dE_r}{dx}, b) = M(b) + R(b)\frac{dE_r}{dx} \tag{5.8}$$

where N is the number of photoelectrons detected, and for a given impact parameter, b, M and R are constants that depend only upon b. Specifically, M is the number of photoelectrons that would be detected if no secondary electrons due to radiative losses were present, and R is a term that describes the number of photoelectrons detected from secondary electrons distributed along the muon's track with total deposited energy equal to $E_t = 2\frac{dE_r}{dx}\sqrt{850^2 - b^2}$.

In this thesis a simplification of Equation 5.8 is used to estimate dE/dx. Instead of

Phase	Starting Run	Ending Run
D_2O	10000	12330
D_2O	12330	16014
Salt	20000	22982
Salt	22982	30175
Salt	30175	33535
NCD	50000	58310

Table 5.1: This table shows the six run ranges for which a dE/dx calibration has been performed. The calibration had to be done at several points in time due to changes in water attenuation and concentrator reflectivity, as these act to reduce the FTI fit number of photoelectrons.

estimating $\frac{dE_r}{dx}$ by considering the Cherenkov photons from the secondary electrons and the muon separately it is assumed that the total amount of detected light is a function of the muon's $\frac{dE}{dx}$. In this case we remove the constant term and have:

$$N(\frac{dE}{dx}, b) \approx R(b)\frac{dE}{dx} \tag{5.9}$$

The R term was estimated by making a profile histogram of the FTI fit number of detected photoelectrons divided by the muon path length as a function of impact parameter for simulated muons with an energy loss between 220-230MeV/m. The 220-230MeV/m bin was chosen because it is sufficiently close to the expected minimum muon energy loss that the deviation between Equations 5.8 and 5.9 will be small. Only through-going muons which fit within 5cm and 10 degrees of the true impact parameter and direction were used in the measurement of R. The resulting histogram was fit with a second order polynomial in two parts, from 0-600cm and 600-850cm. This was done because the fit number of photoelectrons is not continuous across the AV. Because the attenuation lengths have changed over the course of the experiment, several fits were done over different time periods. The six time periods are given in Table 5.1.

By construction, equations 5.8 and 5.9 will result in roughly the same estimate of

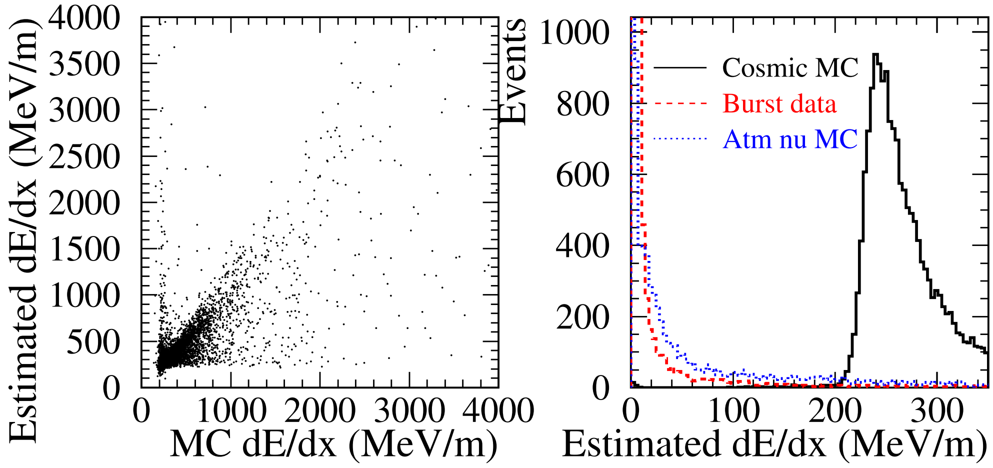

Figure 5.6: The left hand scatterplot shows the estimated vs. true value of dE/dx for simulated cosmic muons. The events with large estimated dE/dx and small true dE/dx occur at large impact parameters. The shorter the path length of the muon the larger the effect fluctuations in the number of detected photoelectrons will have upon the estimate of dE/dx. The events in which the true dE/dx is considerably larger than the estimated value are presumably due to photonuclear interactions which did not result in the production of Cherenkov radiation.

dE/dx at the point where Equation 5.9 was calibrated, in this case 225MeV/m. Away from 225MeV/m the technique used in this thesis was found to have a bias of roughly $0.25(\frac{dE}{dx} - 225MeV/m)$. This bias is acceptable, as this estimate of dE/dx is primarily intended to be used as a cut, and the cut will be set near the 225MeV/m calibration point. Figure 5.6 shows the performance of the estimate on simulated cosmic muons. The scatterplot on the left shows that for most events the estimated dE/dx is comparable to that reported by the Monte Carlo. The histograms on the right show the estimate of the dE/dx for the three types of events discussed earlier, simulated muons, events which fail the burst cut, and simulated atmospheric neutrino interactions inside the PSUP. Based on simulated cosmic muon events a cut for accepting through-going muons was placed at 200MeV/m. It is clear that although the estimate of dE/dx does not perfectly track the true value it is nevertheless a very powerful cut.

Level	Name	Cut Value	Events Remaining
Low	All fit open D_2O events	NHit>250, retrigger	50272
	Burst	See text	34861
	Calibrated PMTs	> 500	9001
	Raw Q_{RMS}	> 4.5 ns	7016
	Raw T_{RMS}	< 38 GSU	6946
	Neck tubes	< 4 fire	6911
High	Impact parameter	< 830 cm	6619
	Fit number of pe	> 2000 pe	6612
	dE/dx	> 200 MeV/m	6559
	CITR	> 0.85	6536
	TRICH	> 0.7	6534
	Linear discriminant	> 0.6	6426

Table 5.2: This table summarizes the cuts that are used to select through-going muon events. The linear discriminant cut will be described in subsection 5.3.7. The cut values were set by studying cosmic muons simulated with SNOMAN 5_0184.

5.3.5 Data Reduction

The effect of the progression of cuts on the open D_2O data set is shown in Figure 5.7. Five histograms are shown with successive cuts applied. By the time the burst cut and the calibrated hit PMT cut are applied the signal to background ratio is already quite large. Note that the requirement of 500 channels with good calibrated charge and time removes a very large number of pickup events, including the peak at 2000 channels (four crates with all channels firing). The remaining data are made up almost entirely of cosmic ray muons. The cuts used are summarized in Table 5.2.

5.3.6 Comparison of Data to Simulated Muons

There is no calibration source for through-going muons, so in the tuning and evaluation of the data selection cuts we are forced to use simulated muons. If there are differences between the simulation and the actual detector the evaluation of the selection cut acceptance could be wrong. The main calibration source for the SNO detector is the radioactive ^{16}N

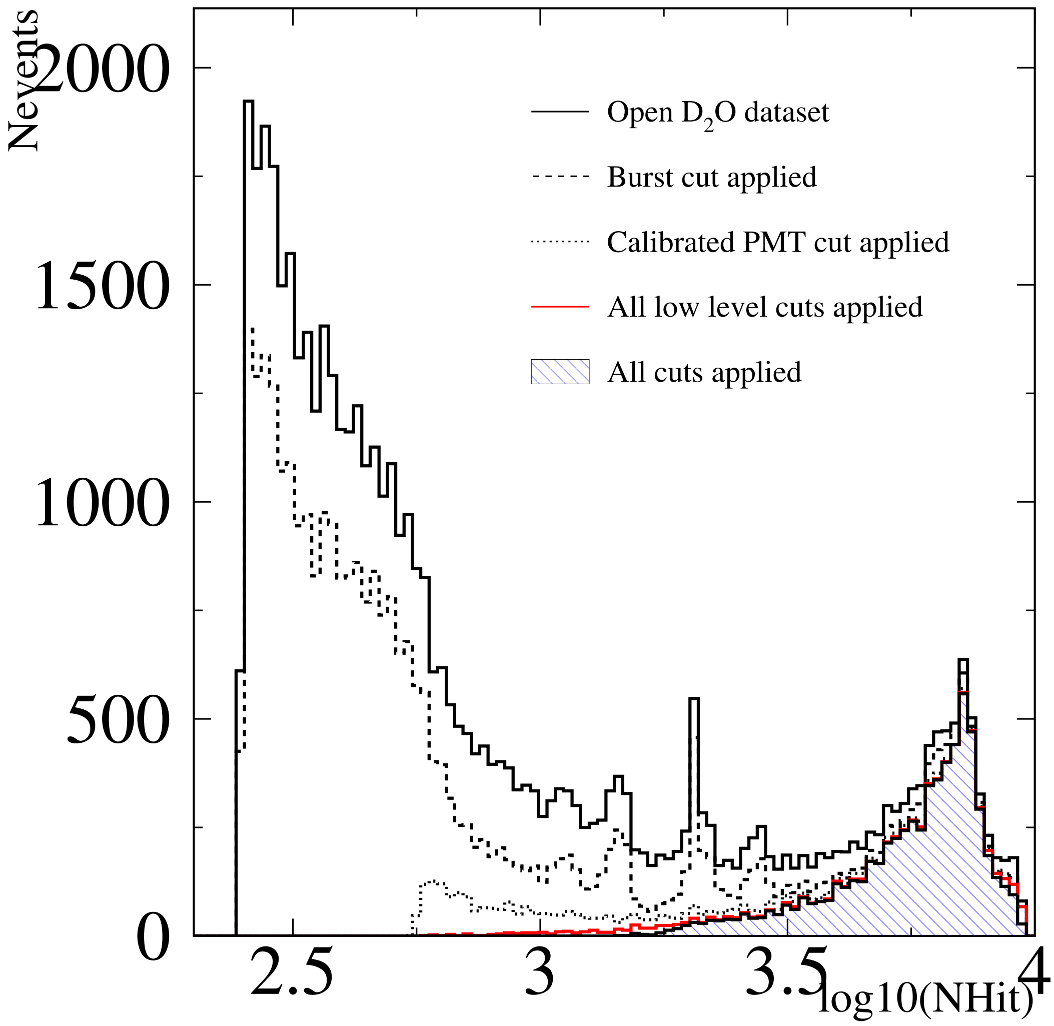

Figure 5.7: This data reduction plot shows the events remaining after a series of low and high level cuts. The abscissa is the total number of channels hit, a number that can be considerably larger than the number of calibrated online PMTs which observe a good time and charge.

source that produces 6.13 MeV gamma rays. The source is run regularly in the D_2O, but only rarely in the H_2O. The ^{16}N source is a good match for the solar neutrino signal of $\sim$few MeV electrons produced in the D_2O through the neutrino deuteron charged current interaction. Muons, in contrast, tend to have energies in the tens or hundreds of GeV, and they produce much of their light in the H_2O. The Monte Carlo simulation is calibrated to result in good agreement between data and Monte Carlo for events in the D_2O with energies of a few MeV, so it is not surprising that at considerably larger energies and in the light water the simulation does not agree with the data to the same extent.

The comparisons between data and simulated events shown in this section were all produced by comparing data and simulated cosmic muon events in the open D_2O data set. Unless it is otherwise noted all of the data selection cuts outlined in the above subsection are applied to both the data and simulated events. Because of the spoiling of the acceptance as a function of zenith angle for the data there are more Monte Carlo events than data events. For this reason when histograms are plotted in this subsection they are normalized to equal area. It should also be noted that while it is not shown, most of the parameters vary considerably as a function of muon impact parameter. In all of the plots simulated events will be shown in black and the open D_2O data in red.

The first and most obvious parameters to compare between data and Monte Carlo are the number of PMTs that fire and the charge that they observe. These two parameters are shown with data and simulated events plotted on top of each other in Figure 5.8. While the distributions look fairly similar by eye they are not identical, and a KS test performed on the two distributions returns a very small likelihood that each is drawn from the same underlying distribution. The high charge tail for both data and Monte Carlo are cut off. The true distributions extend to summed charges as large as 600,000 detected photoelectrons in simulation and 800,000 detected photoelectrons in the data.

The distributions of the number of hit PMTs and number of detected photoelectrons

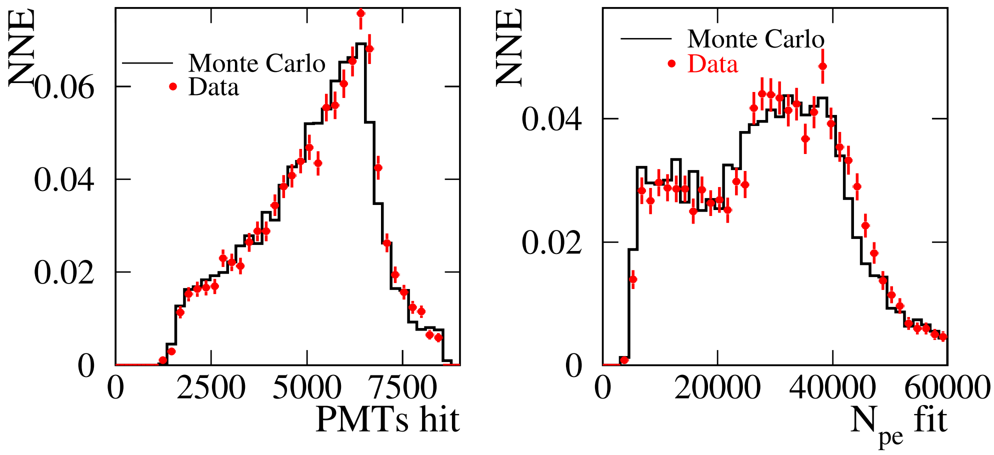

Figure 5.8: These two plots show the distribution of number of calibrated PMTs hit and fit number of photoelectrons. Each plot has been normalized to have unit area, so the ordinate is the normalized number of events (NNE). The structure in the diagram on the right is believed to be due to the acrylic vessel. The number of detected photons per unit path length is smaller inside the AV, and as a consequence the rate of change with respect to impact parameter squared is also smaller. The fit number of photoelectrons distributions are cut off at 60,000 to make the bulk of the data visible.

for data and simulated events are not identical, but agree to a considerable extent given that the detector has not been explicitly calibrated for muon events. The presence of the acrylic vessel makes the job of calibrating the detector considerably harder than it would be otherwise be for several reasons. The acrylic attenuates photons, especially at shorter wavelengths, and it is difficult to separate the acrylic and water attenuation with *in situ* measurements. The acrylic also restricts the movement of calibration sources, and because of these two effects it is difficult to model the H_2O optics. The same effects also make it difficult to measure the angular response of the phototubes at large angles *in situ*, so in producing optical models[10] the simulation has had to use *ex situ* measurements for angles larger than roughly 40 degrees [61].

The differences between the simulation and the data are more pronounced for the anal-

[10]To account for changes in attenuation and concentrator reflectivity over time, several period specific optical models have been developed. The angle at which they must switch from *in situ* to *ex situ* measurements depends upon the amount of data taken during the optics calibration runs for that time period.

ysis parameters dealing with charge and time than they are for the NHit and number of detected photoelectrons. The histograms in Figure 5.9 show the distributions of the four charge and time parameters for data and simulated events. In each of the plots, the simulated events are shown as a black histogram and the data events are shown as red data points. The top two plots show the raw time and charge RMS parameters. The data distribution has narrower raw times than the simulation, and the peak of the charge RMS distribution is larger than that for simulated events. When PMTs are hit by many thousands of photons in the detector the analog to digital converters for both time and charge may return a value of 0. In this case the PMT does not have a good calibrated time and charge, so its information is not used. This effect is not simulated, so larger charges, and thus larger charge RMS values, are possible in the simulation. For this reason the tail of the raw charge RMS distribution is much longer in the simulation than in the data.

The bottom two plots of Figure 5.9 show the Cherenkov cone in-time ratio and the Cherenkov cone PMT fit time RMS. The data events have considerably narrower time residuals than the simulated events. If the intrinsic PMT single pe transit time distribution is narrower than that used in the simulation this would lead to larger fit time residuals for simulated events. A recent bench-top measurement on one SNO PMT has indicated that this may be the case [62]. Another possible explanation for such a discrepancy would be a mismatch in the PMT wavelength response between the data and the Monte Carlo. Photons in the red part of the spectrum travel faster through water than blue photons, so if the PMT response to red photons was too small (large) in the simulation this would result in a slower (faster) average speed of light. Because the speed of light used by the FTI algorithm was not tuned using the Monte Carlo it is possible that the speed of light used in FTI matches that in the data better than the Monte Carlo. The large difference in the Cherenkov cone PMT fit time RMS for data and simulated events will obviously affect the linear discriminant which will be discussed in subsection 5.3.7. A plot comparing the linear discriminant for

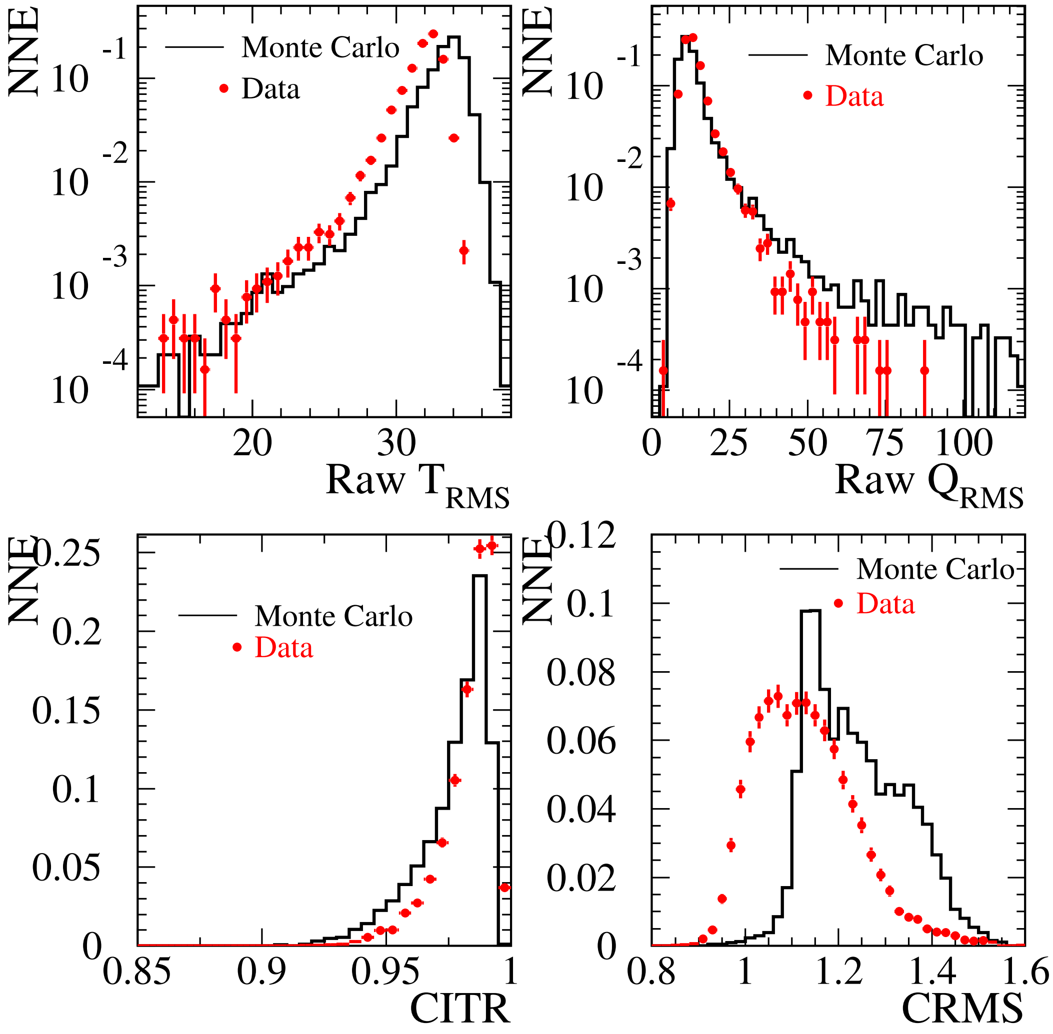

Figure 5.9: These four figures compare the distributions of the analysis parameters for data and simulated events. In each plot the simulated events are shown as a black histogram and the data as red data points. The ordinate in each plot is the normalized number of events (NNE). For each of the parameters the simulation has a larger tail under the cut defined in subsections 5.3.3 and 5.3.4 than does the data.

neutrino induced muon data, cosmic muon simulations, and internal neutrino interaction simulations will be shown in a few pages.

While it is not particularly pleasant that the agreement between the data and Monte Carlo distributions shown in Figure 5.9 is not optimal, this is not necessarily a problem for the analyses in this thesis. In each case the Monte Carlo has a larger tail under the cut value than does the data, so the acceptance measured by Monte Carlo represents a lower bound on the true acceptance of the cuts. In section 5.3.8 we will see that the measured acceptance is very close to one, so the uncertainty on this acceptance will be small compared to other sources of uncertainty.

A comparison of the final two analysis parameters, TRICH and dE/dx, is shown in Figure 5.10. Again, while the data and simulated distributions do not match exactly, they have only small tails below the chosen cut values, and the tail in the simulation is larger than that in the data. It is not shown here, but the largest TRICH values tend to occur for the events with large values of dE/dx. This makes sense, as the possibility of a PMT in the cone not firing decreases with increasing total light. The data peaks at a lower value than the Monte Carlo. This is due to the timing calibration failing for PMTs that observe many hundreds of photoelectrons. In calculating the TRICH fraction, tubes with failed calibrations should have been ignored in both the numerator and the denominator, but they were unfortunately included in the denominator. This means that as the amount of light increases past a certain level, the TRICH parameter actually starts to decrease in data. Although this is a small effect it should be corrected in future analyses.

5.3.7 Background Estimates for the Neutrino Induced Muon Analysis

The three sources of background events to the neutrino induced muon flux and oscillation analysis are instrumental contamination, atmospheric neutrino interactions inside the PSUP, and cosmic muons. In this subsection we will estimate the contribution from each

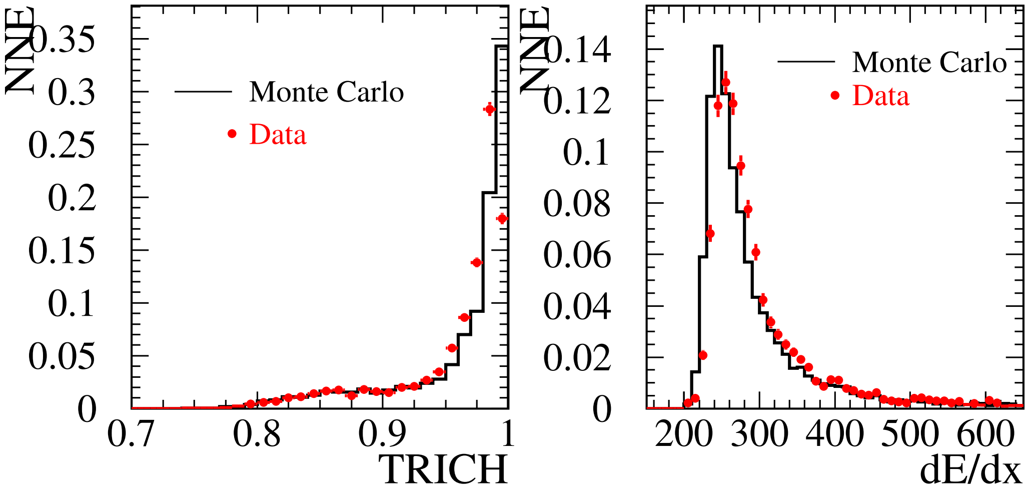

Figure 5.10: The left hand plot shows the distributions of the fraction of PMTs hit inside the trimmed Cherenkov cone for data and simulated events. The right hand plot shows the dE/dx distributions. Note that as for the total charge plots it appears that the total amount of light in the simulation is too small.

of these sources to the signal.

Instrumental Background Estimate

Several plots in subsections 5.3.3 and 5.3.4 showed comparisons of simulated muons to events which failed burst cuts. None of the high or low level analysis cuts has perfect rejection of background events, but in combination they work very well. The plots shown in Figure 5.11 show the events in the range $-1 < \cos\theta_z < 0.4$ which remain in the open data set after applying all of the cuts but CITR and CRMS. The cosmic muon region above $\cos\theta_z = 0.4$ includes multiple muon events which are systematically different from the single through-going muons expected in neutrino induced muon events. The data selection cuts are not 100% efficient at rejecting multiple muons, but they are removed efficiently with a zenith angle cut.

All of the events in Figure 5.11 were scanned by eye to classify them. The eight events with CRMS larger than 2ns in the left hand plot include four PMT WEB events, one bubbler event, one internal atmospheric neutrino interaction event that had a decay electron follower, and two extremely high charge sum events that the scanner (the author) was not able to classify. The dashed line which passes through those events is the equation CITR=3.1-CRMS, and this function was used to create a linear discriminant for instrumental background events, L=3.1-CRMS-CITR. The distribution of linear discriminant values for the open data is shown in the right hand frame of Figure 5.11. A cut on the minimum value of the linear discriminant was set at 0.6 based on simulated cosmic muons.

The data events with linear discriminants in the range (0.2,0.6) included one through-going muon and five events which were classed as events originating inside of SNO. Of the remaining 53 events with linear discriminant larger than 0.6, 50 were classified as through-going muons and 3 were classified as interactions originating inside of the detector. Between 7-29 of the outward looking (OWL) tubes fired in the events classed as through-going muons, while for the events classed as internal interactions two events had 0 OWL hits and one had two.

The instrumental background events in the histogram of Figure 5.11 were fit with a Gaussian in the linear discriminant range (-0.3,0.2). The fit returned a mean of -0.06 ± 0.03, RMS of 0.09 ± 0.03, and constant of 1.7 ± 0.7. If the distribution of linear discriminant values for the backgrounds are hypothesized to be normally distributed, then a limit may be set on the number of events above the cut of 0.6. The cut is $\sim 7\sigma$ away from the mean, so this would imply a negligible background.

It is also possible to constrain the worst-case hypothesis that the background is actually flat rather than distributed near zero. In the range -0.3<L<0.2 there are 0.7 instrumental or unidentified background events per bin, so if the background were flat we would expect to see 15.4 background events in the region with discriminant larger than 0.2. Assuming Poisson statistics the probability of observing zero events is 2×10^{-7}. Similarly, the probability of observing zero events in the range 0.2<L<0.6 is 0.004. From this analysis we conclude that the instrumental background is negligibly small.

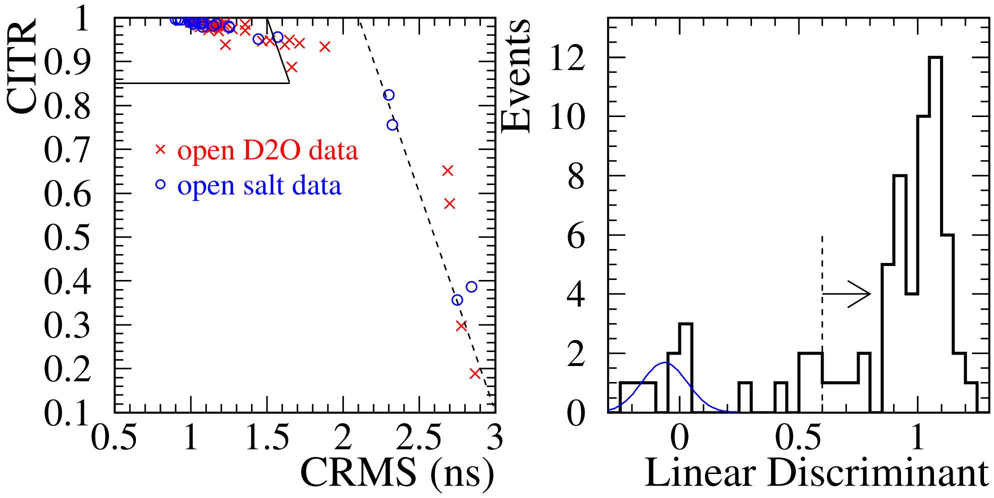

Figure 5.11: The left hand scatterplot shows the Cherenkov cone in-time ratio vs. the cone time residual RMS for events in the open data set. Hand scanning indicates that none of the eight events with CRMS larger than 2ns are through-going muons. The dashed line shows where the linear discriminant (L) mentioned in the text has a value of zero, and the solid lines show the cuts CITR>0.85 and L>0.6. The right hand plot shows the distance of each event from the dashed line in the left hand plot, with a fit to the background events near L=0. The dashed line in the right hand plot shows the chosen cut value of 0.6.

We now compare the value of the linear discriminant in the simulation to that in the data. The value of the linear discriminant for simulated muons in the D_2O phase, muons passing all other selection cuts from the D_2O and salt open data sets, and for simulated internal atmospheric neutrino interactions, is shown in Figure 5.12. In making this plot no cut was placed on the zenith angle, so the data are dominated by downgoing cosmic muons. The fact that the linear discriminant has larger values in the data than in the simulation means that the acceptance of the cut may be underestimated. It will be shown that the cut acceptance is close to one, so this difference will result in only a small uncertainty on the number of detected events.

Internal Neutrino Interaction Background Estimate

Interactions of electron neutrinos and muon neutrinos in the water inside of the PSUP (the "fully contained" and "partially contained" events of Super Kamiokande) are a background

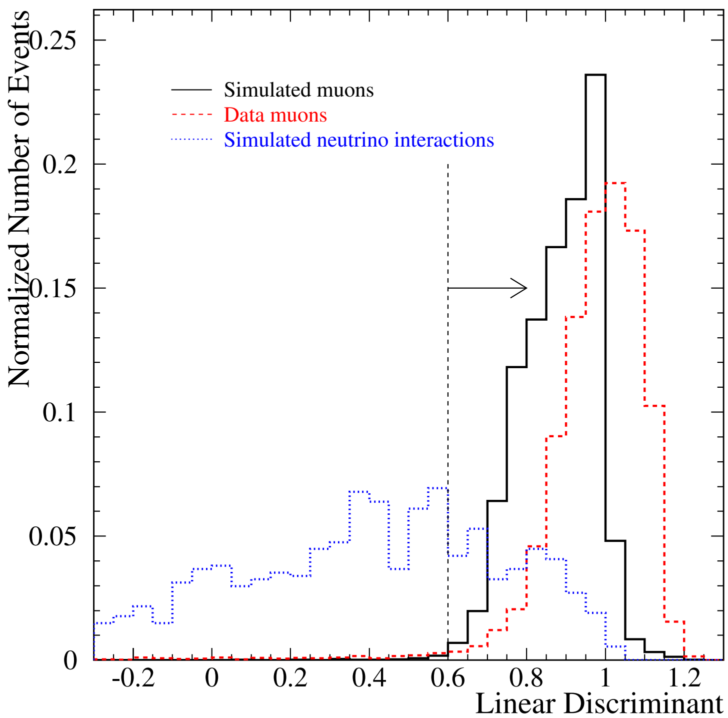

Figure 5.12: This plot shows the value of the linear discriminant for simulated muons (solid black), the D_2O and salt open data (dashed red), and neutrino interactions inside the PSUP (dotted blue). All three histograms are normalized to have area 1. The dashed black line shows the chosen cut value of 0.6. As discussed earlier, most of the difference in the linear discriminant value in data and the simulation is due to the difference in the spread of the fit time residuals.

to the neutrino induced through-going muon analysis. The data selection cuts described above, especially dE/dx and the linear discriminant, do a fairly good job of rejecting most such internal neutrino interaction events, but they were designed to reject instrumental events. None of these cuts were tailored to target the internal neutrino interaction. As a result some of these events pass the data selection cuts.

A simulated data set of 33 times the data expected in the D_2O phase was generated using the neutrino propagation code "nuance [11]" using the unoscillated Bartol04 fluxes for SNO [26]. The data set consisted of 17044 internal neutrino interaction events, of which 5388 were internal electron neutrino interactions and 11706 were internal muon neutrino interactions. After applying all of the data selection cuts and a zenith angle cut of $\cos\theta_z < 0.4$, 326 such events remained. This corresponds to an estimated background mean number of events in the case of no oscillations of 9.9 ± 0.5(stat). The systematic uncertainty on this background is large, due to uncertainties in the neutrino flux prediction and the cross section of neutrinos in the detector. To account for these uncertainties a conservative systematic of 30%, or 3 events, is assigned.

The internal neutrino interaction events that pass the data selection cuts are very bright events that light up a lot of the detector. It is for this reason that these events are able to pass the dE/dx and TRICH cuts. The neutrinos that cause these events typically have energies in the 1-10 GeV range, with a tail out to about 100GeV. The three plots in Figure 5.13 show the reconstruction properties of the events which pass the cuts. The first frame in Figure 5.13 shows that FTI does a good job of fitting the simulated neutrino direction. The second frame shows the fit zenith angle distribution. The third frame shows the distribution of fit impact parameters.

[11] The name of the Monte Carlo program nuance does not have caps. In this thesis a different font will be used in order to distinguish the word nuance from the code nuance.

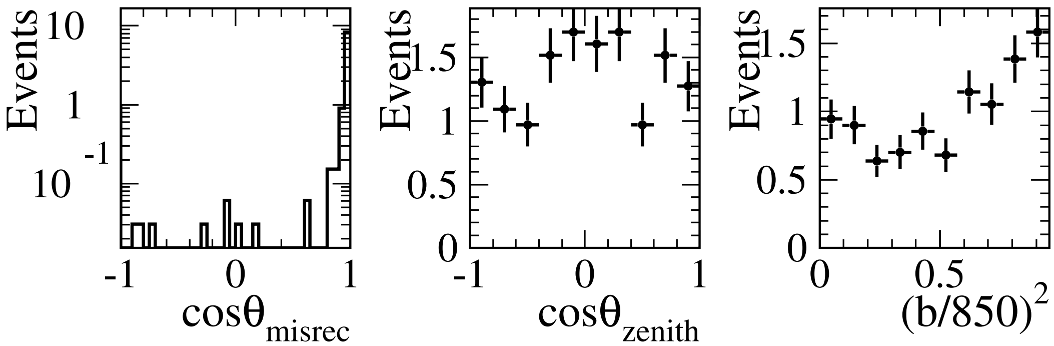

Figure 5.13: These plots show the internal neutrino interaction background to the neutrino induced through-going muon analysis. The left hand plot shows the cosine of the angle between the direction the neutrino was traveling and the fit direction from FTI. Almost all of the events reconstruct in the neutrino direction. The center plot shows the zenith angle distribution of the internal neutrino interaction events. The right hand plot shows the distribution of fit impact parameters. The left and right hand plots have been scaled to have an integral of 9.9 events, the central plot is scaled to 13.6 events

Cosmic Muon Background Estimate

The cosmic muon signal is orders of magnitude larger than the neutrino induced muon signal, so misreconstructed cosmic muons are a potential background to the neutrino induced muon analysis. The size of this background was estimated using simulated cosmic muons. The zenith angle distribution for a $\sim$2.5 times statistics (854.5 day) simulation of cosmic muons in the $\mathrm{D_2O}$ phase is shown as the black histogram in Figure 5.14. Two events are observed to reconstruct below a $\cos\theta_z = 0.4$, and both events had a generated $\cos\theta_z$ direction of less than 0.4. The 90% confidence level interval corresponding to two events is given by [63] as (0.53, 5.91), so the 90% confidence level estimate of the number of cosmic background events per year is (0.22, 2.52). A careful observer may note that the shape of the zenith angle distribution is steeper in the data than in the Monte Carlo. The possible causes of this will be discussed in Chapter 7; for now it is sufficient to note that this means the simulated cosmic background estimate may be overestimated.

In simulating cosmic muons most of the time is spent propagating the event through the SNO detector rather than simulating the muon's passage through the rock. It is pos-

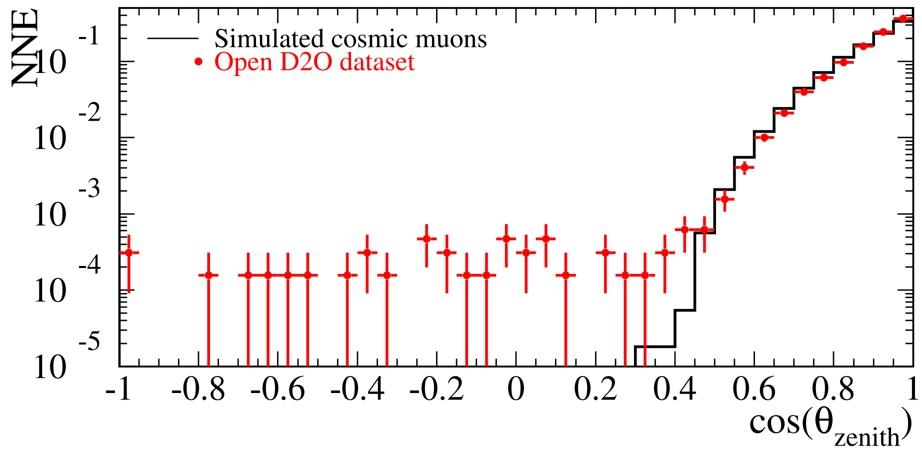

Figure 5.14: This plot shows the zenith angle distribution for 854.5 days of simulated cosmic muons (solid black) and the D_2O open data (red points). The two histograms are normalized to have area 1. The observation of two simulated events below $\cos\theta_z = 0.4$ sets a 90% confidence interval of (0.22, 2.52) cosmic muon background events per year.

sible to run the MUSIC cosmic muon simulation in SNOMAN without propagating the events through the SNO detector, and this makes simulating many times the expected rate possible. Cosmic muons were simulated with a surface zenith angle distribution in the range[12] $0.25 < \cos\theta_z < 0.65$ and propagated to SNO. Muons that had less than 2000MeV were discarded, as these events would not pass the data selection cuts. The resulting zenith angle distribution expected at SNO was normalized to the 458 events in the range $0.4 < \cos\theta_z < 0.6$ in the 2.5 times D_2O phase statistics simulated data set. The normalized number of events predicted with $\cos\theta_z < 0.4$ is 1.51 ± 0.14. The probability of observing two events with an expected mean of 1.51 is about 25%. For the 337.3 day D_2O data set this corresponds to an expected cosmic muon background of 0.59 ± 0.05 events. Due to the disagreement between the simulated zenith angle distribution with that observed in the data a conservative 100% uncertainty is assigned to this background.

[12]The normal range is $0.3 < \cos\theta_z < 1$.

5.3.8 Estimate of Signal Acceptance

The final task of data selection which must be addressed is the acceptance of the cuts to the through-going and stopping muon signals. When an experiment uses calibration sources to measure acceptance the acceptance is just the number of events passing the cuts divided by the number of events produced by the source. The situation is not quite so clear when the acceptance is measured using a Monte Carlo technique. The Super-Kamiokande experiment, for example, defines acceptance [10] as the number of events that reconstruct as through-going, divided by the total simulated number of through-going events. Because stopping muons are sometimes misidentified as through-going, this results in an acceptance larger than one.

In the interest of avoiding the possibility of having an acceptance larger than one this analysis defines acceptance in the following way:

$$\frac{N_{830f}}{N_{830MC}} = a_{TG} \times N_{ATG} \tag{5.10}$$

where N_{830f} is the number of events which pass the cuts and reconstruct with a fiducial radius less than 830cm, N_{830MC} is the number of simulated events which had a true impact parameter less than 830cm, a_{TG} is the acceptance for through-going muons, and N_{ATG} is a normalization term for the area of the detector. More explicitly, a_{TG} is the number of through-going events generated inside of 830cm that reconstruct inside 830cm and pass cuts, divided by the number of through-going events that reconstruct inside 830cm. N_{ATG} is the number of through-going events that reconstruct inside 830cm, divided by the number of through-going events that were generated inside 830cm. Stopping muons are not included when estimating any of these variables.

The normalization term accounts for the fact that if there is any radial bias, events with simulated impact parameter greater than 830cm may reconstruct inside of 830cm. It can be written in the form

$$N_{ATG} = \frac{N_{ii} + N_{oi}}{N_{ii} + N_{io}} \tag{5.11}$$

where N_{ii} is the number of events generated inside of 830cm that reconstruct inside 830cm, N_{oi} is the number of events generated outside of 830cm that reconstruct inside of 830cm, and N_{io} is the number of events generated inside that reconstruct outside, or do not reconstruct at all. Since N_{ii} is much larger than the other terms this can be approximated as

$$N_{ATG} \approx 1 - \frac{N_{oi} - N_{io}}{N_{ii}} \tag{5.12}$$

to simplify the calculation of the uncertainty. Assuming the distributions of the Ns follow Poisson statistics ($\delta N_{ii} = \sqrt{N_{ii}}$) then the uncertainty on Equation 5.11 is:

$$\delta N_{ATG} = \frac{\sqrt{N_{io} + N_{oi}}}{N_{ii}} \tag{5.13}$$

The acceptance and normalization factors were estimated by using simulated cosmic muons. In estimating the acceptance one should use a different dataset from the one used to set the cuts, and two equal sized data sets are available for this purpose. The first data set was created using SNOMAN 5_0184, and the second with SNOMAN 5_0185. There were three differences between the two versions. The first was a fix of a simulation bug which had been responsible for adding some jitter to the Cherenkov photon creation times. The second was that the value of the constant hc was changed by ~0.5%, as it had been discovered that that value had been set incorrectly in the simulation. The third change was an improvement in the numerical precision of the selection of radiative processes. The final change (which was specific to the data set used here rather than a change in the default settings) was to shift the position of the acrylic vessel in the simulation downward by 5.32cm, as the original version of the simulation had incorrectly centered the acrylic vessel

Data set	a_{TG}	N_{ATG}
Used to set cuts (5_0184)	0.9948 ± 0.0005	0.9962 ± 0.0007
5_0184	0.9959 ± 0.0005	0.9956 ± 0.0007
5_0185	0.9943 ± 0.0006	0.9978 ± 0.0006
Combination of 84 and 85	0.9951 ± 0.0004	0.9967 ± 0.0004

Table 5.3: This table shows the estimated values for the acceptance and fiducial area for three independent simulated data sets. The first data set is the one that was used to set the cuts. The second was generated using SNOMAN version 5_0184 and the third with SNOMAN 5_0185. It is possible that the changes between the SNOMAN versions changed N_{ATG}, but this difference is very small compared to the expected statistical variation in the number of events.

at the center of the PSUP. The different estimated values for a_{TG} and N_{ATG} are shown for each of the three data sets in Table 5.3. This analysis will use the combined value from the two data sets not used to set the cuts, $a_{TG} = 0.9951 \pm 0.0004$ and $N_{ATG} = 0.9967 \pm 0.0004$.

The through-going muon acceptance as a function of impact parameter for simulated muons is shown in Figure 5.15. The black points show the acceptance if all cuts including the fiducial area cut are applied, the red points show the acceptance of the cuts regardless of fit impact parameter. The top plot shows the acceptance over the entire area of the SNO detector, and the bottom plot shows the acceptance near the fiducial cut of 830cm. The black points drop off earlier than 830cm because of the fitter's slight outward radial bias. This bias is the reason that N_{ATG} is slightly smaller than one.

The acceptance for stopping muons is complicated by the fact that the acceptance depends upon the energy of the muon when it enters SNO. This means that the acceptance for stopping cosmic muons may be different from the acceptance for stopping neutrino induced muons. To avoid having to understand the stopping acceptance as a function of energy and impact parameter we instead study the stopping acceptance as a function of the fraction of the distance that the muon propagated through SNO. The left hand frame of Figure 5.16 is a scatterplot showing the fraction of the distance through SNO that the muon traveled vs. the square of the simulated impact parameter. The distance the muon traveled is defined as

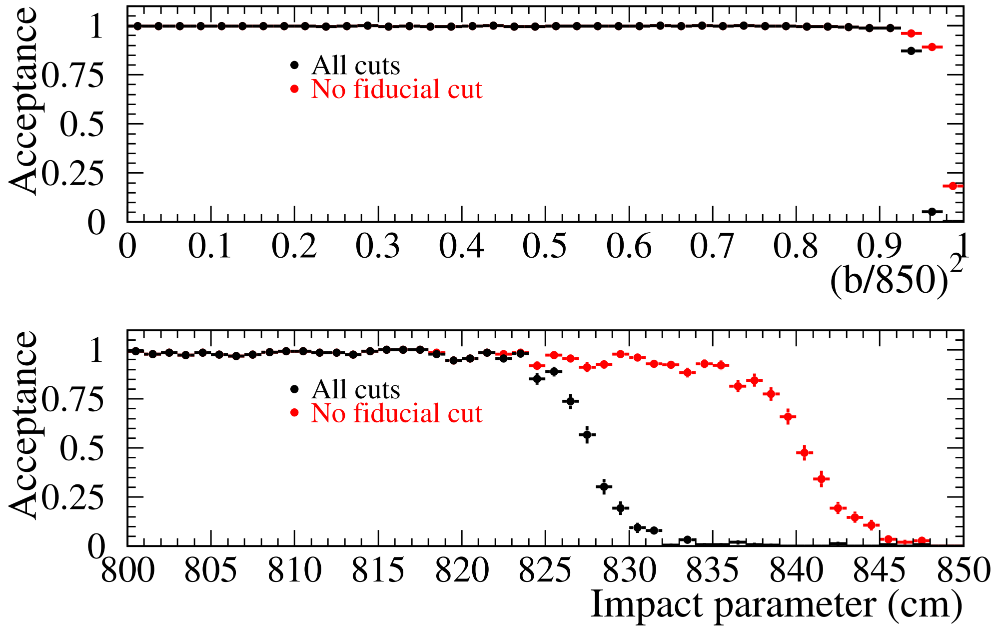

Figure 5.15: The acceptance of the cuts for simulated through-going muons is shown against simulated impact parameter. The top plot shows the entire detector area, the bottom plot is a blow up near the fiducial radius cut of 830cm. The black points show the acceptance with all cuts applied, while the red points include all cuts but the fiducial radius. The black curve drops off before 830cm because of the slight outward radial bias of FTI. Due to this bias the fit impact parameters are slightly larger than the true values near 830cm.

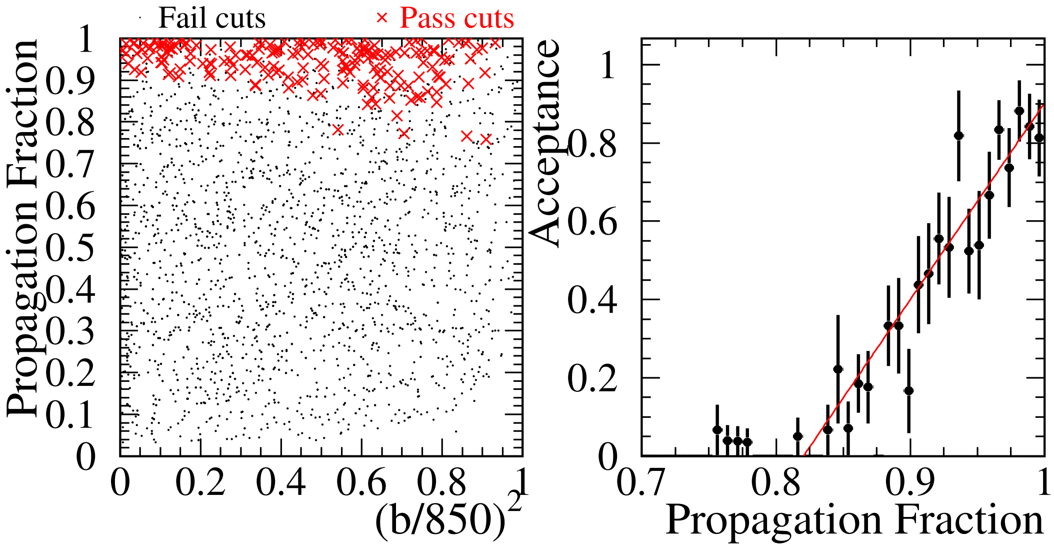

Figure 5.16: The left hand scatter plot shows the fraction of the distance that the muon propagated through SNO vs. the square of the simulated impact parameter. The right hand plot shows the integrated acceptance as a function of the muon propagation fraction, along with a first order polynomial fit between 0.8 and 1.

the distance from the point at which it entered SNO to the point at which it stopped, and because of multiple scattering this distance is actually slightly smaller than the distance that the muon actually traveled. The events that fail the cuts are shown as black dots, while the events that pass the cuts are shown as red x'es. If we ignore the variation with impact parameter we can get an acceptance as a function of the fraction of the distance that the muon propagated through SNO, and this is shown in the right hand frame. A linear fit to the data above a propagation fraction of 0.8 is shown. The fit has a slope of 5.0 ± 0.3, an intercept of -4.1 ± 0.3, and the χ^2/dof is 18.9/20. This equation for the stopping muon acceptance will be used in Chapter 7 to estimate the acceptance of stopping neutrino induced muons. The integrated acceptance for stopping cosmic muons was estimated as 0.065 ± 0.012. using simulated cosmic muon events.

Acceptance of the Neck Tube Cut

The veto tubes installed in the detector neck and glovebox are not simulated in SNOMAN, so the acceptance of the neck tube cut cannot be evaluated using the simulation. It was found that after all cuts but the neck cut were applied to the open $\mathrm{D_2O}$ data set 6428 events remained, and the neck cut was responsible for cutting two events. Those two events were scanned by eye and identified as being through-going downward going muons. This indicates that the acceptance loss due to the use of the neck cut is on the order of the uncertainty of the acceptance estimate, so the acceptance of the neck tube cut will be assumed to be exactly 1.

5.4 Detector Live Time

The uncorrected live time for muon analyses in the $\mathrm{D_2O}$ phase is $2.91459 \pm 0.00016 \times 10^7$s. The retrigger and burst cuts removed a total of 7431s, so the total muon live time in the $\mathrm{D_2O}$ phase is $2.9138 \pm 0.00016 \times 10^7$s, or 337.25 ± 0.02days.

5.5 Summary of Data Selection

This chapter has discussed the problem of separating through-going muon events from instrumental and physics backgrounds. The first step of data selection was to select the set of "runs", intervals of time referenced by an integer number, which make up the data set. Next, we discussed two sets of simple and complex analysis parameters which are used to discriminate muon events from the backgrounds. We then compared the distributions of the analysis parameters from simulated events to those for the data. The distributions were found to not match perfectly, but they do not differ in a way that is expected to affect the extraction of physics parameters from the data. Next, we estimated the amount of physics and instrumental background events which remain after performing all of the cuts. Finally, we estimated the acceptance of the cuts to through-going and stopping muon events.

In the next chapter, we will discuss how systematic uncertainties in the muon simulation affect the analysis parameters, such as the acceptance, which we have just measured.

Chapter 6

Systematic Effects and Uncertainties

In Chapter 5 a set of analysis parameters which allow the separation of through-going muon events from backgrounds were presented. The acceptance of through-going muons to these cuts, and the effective fiducial area were estimated by using simulated events. These estimates may therefore be sensitive to parameters in the SNOMAN detector simulation, such as the model of PMT prepulsing or the probability of Rayleigh scattering. In the first part of this chapter we will study the effect that varying input SNOMAN parameters has on the estimated acceptance and fiducial area. In the second part of this chapter we will compare the simulated data sets to the open D_2O data set.

6.1 SNOMAN Model Parameters

Many of the input parameters of the SNOMAN simulation are measured and tuned in order to make the simulation match the calibration data. At low energies, the primary background to the solar neutrino signals that SNO was built to study are betas and gammas, produced by the radioactive decay of uranium and thorium daughters. The number of low energy background events that are estimated to have an energy above a given threshold is highly

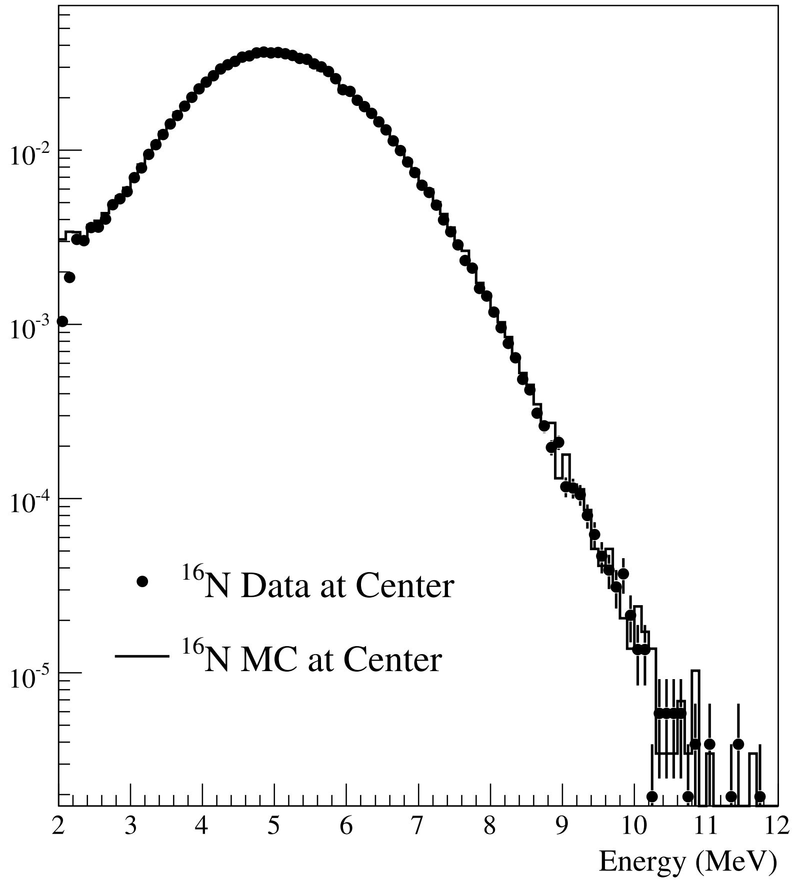

Figure 6.1: The estimated energy of ^{16}N events with the source near the center of the detector is shown for simulated events (histogram) and data (points). The total number of events has been normalized to 1 for both the data and the simulation. Figure courtesy of Monica Dunford.

dependent on the resolution of the energy estimator. For this reason a great deal of the running time of the experiment has been devoted to taking calibration data[1], and the majority of analysis effort has been devoted to modeling the detector. The result of this effort is that the data and simulation agree very well for source calibration events in the D_2O. Figure 6.1 compares the estimated energy distribution of events from the ^{16}N source at the center of the detector for both calibration data and simulated ^{16}N events.

There is more disagreement between data and simulation in muon events than is ob-

[1]The ratio of time spent taking calibration data to "neutrino" data during the ~seven years of the experiment is roughly 2:7.

served in ^{16}N events, as was shown in Figures 5.8, 5.9, 5.10, and 5.12 of the last chapter. The two important differences between muon and solar neutrino events from a calibration standpoint are the total light level and the fact that the muons produce light in the H_2O. The differences due to the light level were discussed and accounted for in Chapter 3. Differences due to the different event geometries will be discussed below.

6.1.1 Potential Sources of Differences Between the Data and Simulation

Attenuation and PMT Response

Most of the optical and energy calibrations are performed using sources in the D_2O, which makes it difficult to separate the twin effects of Cherenkov photon light attenuation in the acrylic and in the light water. This ambiguity has a larger impact on events taking place in the H_2O than it does on events in the D_2O. The reason for this is that the total Cherenkov photon path length through acrylic and H_2O does not vary greatly if the photon originated in the D_2O. In muon events, in contrast, the photon path length through H_2O is very different for different PMTs, as photons may pass through the acrylic zero, one, or two times. Because of the velocity dispersion of optical photons in water, uncertainties in the attenuations can affect both the total amount of light detected from the muon and the arrival times of those photons. Photons at the red end of the Cherenkov spectrum travel faster in water than those at the UV end, so a systematic shift in the detected Cherenkov photon wavelength distribution would result in a systematic shift in the average observed photon speed of light.

In addition to the attenuation effects discussed earlier, Cherenkov photons produced in the light water are likely to be more sensitive to geometrical effects related to the PMTs.

The angular response of the PMTs decreases rapidly beyond an incident angle[2] of 40 degrees, and because the AV restricts the possible positions at which calibration sources can be deployed, the angular response cannot be measured *in situ* beyond roughly 40 degrees. The uncertainty on the PMT angular response is one of the dominant systematics for the solar neutrino energy estimators; the light produced by muons tends to have much larger PMT incident angles than light produced inside of the D_2O.

The six plots shown in Figure 6.2 will be used in the next few paragraphs to illustrate the importance of each of the following effects: acrylic attenuation, PMT angular response, and PMT wavelength response. The three pairs of plots show the wavelength distribution and photon incident angle distributions for 5MeV electrons inside a radius of 550cm, a simulated muon at an impact parameter of 380cm, and a simulated muon at an impact parameter of 800cm. In each plot the black histogram shows the distribution for all photons which hit the PMT, and the red points show the distribution for photons which liberated an electron from the photocathode. The phototubes have a quantum efficiency of roughly 20%, and other losses reduce the number of detected photons further. If the distributions were plotted with their normal scale the structure of the photoelectron distributions would not be visible. For this reason the scale of the two distributions has been modified. The true ratios between total number of photons and number of photoelectrons are (9.3:1), (10.5:1), and (29.9:1) respectively.

The acrylic attenuates most of the light below 300nm, and its effect can be easily seen by studying the black histograms in the left hand plots of Figure 6.2. All of the light produced by electrons in the D_2O must pass through the acrylic, and the resultant feature in the range 200nm-280nm is clearly visible when comparing the plots. The muon in the middle left plot produces some light in the H_2O which does not have to pass through the

[2]The PMT incident angle is defined with respect to the PMT symmetry axis. A photon normal to the PMT face (parallel to the symmetry axis) thus has an incident angle of 0 degrees.

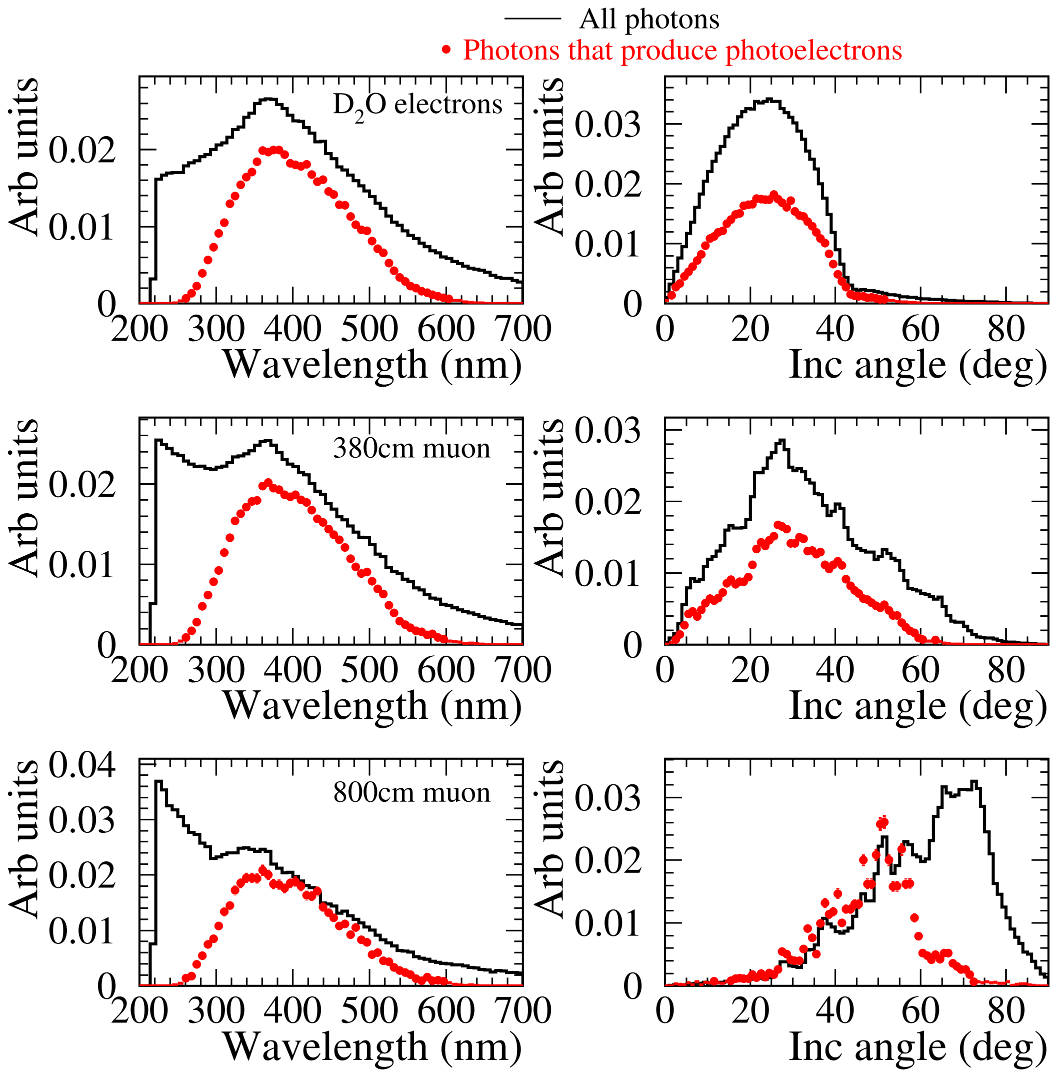

Figure 6.2: This series of plots shows the simulated distributions of Cherenkov photon wavelengths and PMT incident angles for three different event classes. The upper set is for 5MeV electrons in the D_2O with radii smaller than 550cm, the middle and lower sets are for individual muons with impact parameters of 380cm and 800cm respectively. The black histograms show the distributions for all photons that hit PMTs, and the red points show the distributions for photons which caused the PMT to fire. The red data points have been scaled up to allow the structure to be seen. The true scalings (all:hit) are: For D_2O electrons (9.3:1), the 380cm muon (10.5:1), the 800cm muon (29.9:1). The largest component of the photon loss is the PMT quantum efficiency, which is on the order of 20%.

acrylic, so considerably more low wavelength light reaches some PMTs. The situation is much more drastic in the bottom left plot, where all of the light was generated in the H_2O, and much of it is directly incident on the phototubes without passing through the AV. Because the acrylic attenuates the slowest photons, a change in acrylic attenuation is likely to effect the overall charge of muon events but not the PMT time residual distributions.

The black histograms in the right hand plots illustrate the earlier statement that the incident angle of photons produced in the D_2O is generally smaller than those produced in the H_2O. Only 6% of the light produced by electrons in the D_2O hits the PMTs with an angle greater than 40 degrees. This increases to 32% for the muon with an impact parameter of 380cm, and to almost 91% for the muon at 800cm. The corresponding fractions for the total detected light are 5%, 23%, and 75.8%. Because the angular response is not measured *in situ* past roughly 40 degrees, changes in the response from the original *ex situ* measurements are neither probed nor modeled. Changes in the PMT angular response will affect all wavelengths, so such a change could affect both the total charge and the timing of muon events.

The simulation makes use of an *ex situ* measurement of the relative response of the SNO phototubes as a function of wavelength. As the reader has probably already noticed, the response of the PMTs increases from zero quickly after about 250nm, and drops away more gradually after the peak. This structure is mainly due to the design of the PMTs. At short wavelengths the structure is determined by the choice of phototube glass, as glass attenuates ultraviolet light. The large wavelength behavior comes from two factors. The first is that the quantum efficiency of the photocathode is inversely proportional to the photon wavelength. The second is that the probability of an electron escaping into the vacuum is dependent upon the wavelength and the photocathode material used [64]. A systematic difference between the wavelength response of the SNO tubes and the *ex situ* measurement could affect both the total charge and the timing of muon events.

Late Light

A time residual plot for tubes illuminated by a source at the center of the detector was shown in Figure 2.3 on page 34. Roughly 20% of the hits come after the “prompt peak”. These hits are due to PMT late pulsing, Rayleigh scattered photons, and photons reflected from the acrylic vessel, the photomultipliers, and the PMT light concentrators. If these effects are not simulated properly they could have some effect on the overall light level of muon events. They are not, however, likely to have any impact on the timing, because prompt photons will almost always have already fired the channel discriminator.

Electronics Effects

A number of electronics effects are not simulated in the Monte Carlo. The measured times and charges, for example, have discrete rather than continuous values in the data, due to the use of analog to digital converters. The difference in ADC values is small enough (~ 0.1ns and 0.02pe) that this is unlikely to result in much of a difference in muon events. Another example is that when a phototube observes a very large charge the ADCs may return a railed value (0 or 4095), and it is suspected that this is the source of the difference between the data and simulation values of raw charge RMS seen in Figure 5.9. In normal through-going muon events a very small fraction of tubes observe such large charges.

Three electronics effects are known to have an important impact on the measured charges and times. The first is a nonlineararity problem in the commercial ADCs used in the digitization of the times and charges. This problem was corrected using a hardware fix on August 18, 2000. The second and third effects are known as the “time since last hit” (TSLH) and “charge since last hit” (QSLH) effects. In the early stages of SNO it was observed that the time residuals for calibration runs taken at low rate were considerably larger than the time residuals in the normal high rate calibration runs. Special calibrations

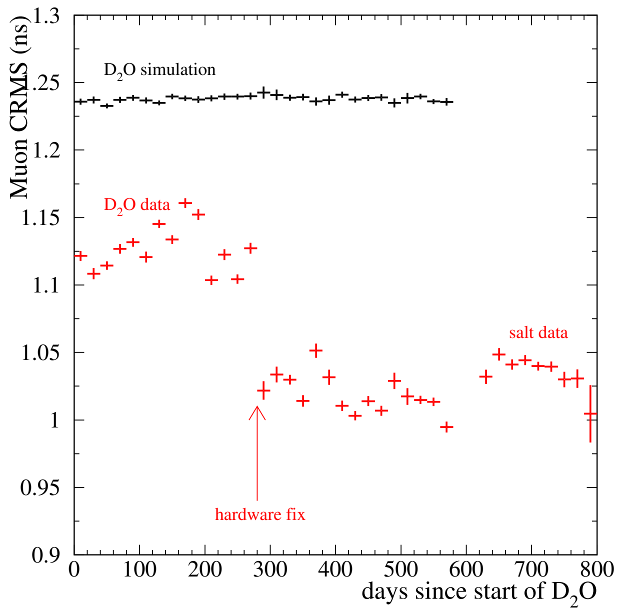

Figure 6.3: The mean time residual RMS for muon events as a function of time, with the start of production data taking as day zero. Two hardware fixes were implemented in August of 2000, one to fix the ADC nonlinearity problem, the other to mitigate the TSLH effect described in the text. Note that the plot has a suppressed zero.

were performed and code was written to remove this problem. The software fix has been applied to the data used in this thesis. The TSLH and QSLH problems were also found to be correlated with the temperature of the QUSN7 analog storage chip. The magnitude of the effects was reduced by with the installation of extra cooling fans in the racks, and a change in the operation of the chips. The effects of these changes can be seen in Figure 6.3 which shows the mean value of the muon time residual RMS as a function of time. The striking feature of the plot is that after the hardware fixes in August of 2000, the RMS of the time residuals decreased by an amount equivalent to a smearing of roughly one half of a nanosecond. By coincidence, the data shown in Figure 5.9 were all taken before the hardware fix, so the disagreement between CRMS in the data and simulation is actually greater than was shown in that plot.

Effects associated with electronics tend to smear out timing distributions rather than to narrow them. For this reason it is unlikely that the observed differences between the data and the Monte Carlo are due to an effect associated with the electronics.

Intrinsic PMT Single Photoelectron Transit Time Distribution

It is possible that the discrepancy between the muon time residuals in the simulation and the data is largely due to the simulated single pe transit time distribution being too wide. As was mentioned in Chapter 5, a new bench-top measurement on a single SNO PMT has indicated that the single photon transit time distribution may have an RMS close to 1.25ns rather than the $\sim$1.6ns in the simulation. The effects of the transit time distribution were studied in simulated events by removing it, ie. by not smearing the photon arrival times and not applying the multiple photon timing correction. The muon events generated by this simulation had a mean FTI fit CRMS of $\sim$0.78ns. This implies that the 1.6ns smearing in the single photoelectron transit time distribution corresponds to roughly a 0.96ns smearing of the first photon arrival times for muon events. If we assume that the smearing observed in muon events scales linearly with the input single pe transit time RMS then we find that tubes with a 1.25ns single pe transit time spread would have a CRMS of $\sim$1.08ns. If this is truly the source of the discrepancy then it is a fairly benign problem, as this degree of smearing is unlikely to affect reconstruction. One should keep in mind, however, that a measurement performed on a single tube has very little statistical significance.

PMT Prepulsing

The PMT single photon transit time distribution used in the simulation is not a function of the incident angle of the photon on the PMT. Recall, however, that the prepulses are due to photons that eject electrons directly from the first dynode. The dynode stack is set back roughly 25cm from the front of the light concentrator, so light which does not illuminate the PMT at close to normal incidence is not expected to cause a prepulse. The fraction of PMTs in the cone which have a prepulse is about 2.5 times larger in muon simulations than in the data, but as this fraction is less than 1% it is unlikely to have much of an effect upon reconstruction.

Detector Geometry

The model of the detector in the simulation is systematically different from the true detector in some respects out of computational necessity. For example, the shape of the light collecting reflectors is approximately a Winston cone, but because the Winston cone does not have a tractable analytic form they are instead modeled as an elliptical torus [65]. It is known that the reflectivity of the concentrators is not uniform, but this non-uniformity is not modeled[3] in the simulation of photon reflection from the concentrators. The positions of individual elements in the detector are obviously not known to the floating point precision used in the Monte Carlo.

It is not possible to use the simulation to test what effects these sorts of geometrical simplifications have, but they are unlikely to have an effect on muon reconstruction. The multiple scattering of the muon and of the knock-on electrons the muon produces as it passes through SNO should act to smear out any effects associated with short distance scale position uncertainties.

Upward vs. Downward Asymmetries

The largest deviation in the full rotational symmetry of the SNO detector is the presence of the neck and the AV support ropes. The acrylic in the neck is not transparent, so muon light that hits it will be absorbed. The ropes which support the AV hang vertically in the top half of the light water volume. The effect of these known asymmetries can be easily checked using the simulation.

[3]The efficiency of each PMT is measured and simulated, so the most direct effect of the reflectors is taken into account.

Low Energy Electrons

As muons pass through SNO they produce many knock-on electrons, and these electrons produce their own Cherenkov light. Prior to SNOMAN release 5_0184 a cut of 2MeV was placed on the minimum energy at which knock-on electrons were simulated. It is unclear why this unphysical cut was part of the simulation. In the course of performing systematic studies it was realized that the cut had an effect on the muon time residuals. Without these low energy electrons the time residuals in simulation were even larger than was shown in the earlier plots.

Electrons and their production of Cherenkov photons are simulated in SNOMAN using the EGS4 Monte Carlo. Simplifications used in the EGS4 Monte Carlo may therefore result in biased results. In particular, some resultant EGS variables are dependent upon the choice of the electron step size used in the EGS simulation. Lay discusses the choices made in integrating EGS4 into SNOMAN in [65]. If there is a small bias in, say, the total production of Cherenkov light in EGS4, the size of its effect on the SNO simulation will be smaller than the bias itself. This is because the SNO simulation is calibrated to reproduce the ^{16}N calibration data. Therefore, a bias in the total number of electrons would be removed by modifying the phototube response.

In contrast to electrons, muons are propagated in the SNO detector using the GEANT3 Monte Carlo. The production of Cherenkov light is handled with a set of routines that are analogous to those used by EGS. It is possible, however, that a bias in the production of Cherenkov photons from EGS which was artificially removed from single electron events will result in a bias in the production of light from the muon.

Monte Carlo Step Size for Muons

In the GEANT3 Monte Carlo simulation muons are propagated by choosing an energy dependent step size, and then moving the muon forward in discrete steps smaller than this maximum size. If the step size used by the simulation is too large this can result in biases. The reason for this is the fact that the cross sections for many processes are energy dependent. The effects of step size are easy to test, as the size can be arbitrarily reduced by the user (at the expense of a slower simulation).

Multiple Photoelectron Timing Correction

The multiple photon timing correction was estimated for the SNOMAN Monte Carlo in subsection 3.1.4, and for the data in 3.1.5. The Monte Carlo correction accounts only for the multiple photoelectron timing effect, while the data correction has to take care of both the multiple pe effect and the discriminator walk. The discriminator walk for individual tubes is known not to be identical [40], but the data correction does not take this into account. Given that the time residuals are narrower in the data than the Monte Carlo this does not appear to be an important difference.

6.1.2 Tests of Input SNOMAN Model Parameters

The previous subsection listed a number of assumptions that went into building the SNOMAN simulation, and how systematic errors on those assumptions could potentially affect muon analyses. For many of these assumptions it is possible to test whether or not muons are sensitive to their associated parameter settings. This has been done by running simulations of the D_2O data set with individual tunable parameters changed from their default values. Because the simulation of the entire data set takes a significant amount of computer time these systematic tests were performed with scaled down statistics, usually roughly 1/5 the amount of cosmic ray muons expected in the D_2O data set.

The studies of these systematic effects took place over a period of several months, during which the SNOMAN muon simulation was changed. The bug that affected the outgoing angle of bremsstrahlung photons was fixed in SNOMAN release 5_0183, the unphysical cut on knock-on electron production was removed in SNOMAN release 5_0184, and a number of minor bugs were fixed in SNOMAN release 5_0185. Because some of these changes affected the CRMS parameter, the acceptance of the linear discriminant cut changed considerably between release 5_0182 and 5_0185. For this reason the linear discriminant cut is not applied when making comparisons between the versions.

The effect of changing input parameters of the SNOMAN model upon the acceptance (a_{TG}) and area normalization factor (N_{ATG}) of the previous chapter is shown in Table 6.1. The table lists each input parameter which was tested, the SNOMAN release used to test it, and the estimated values of a_{TG} and N_{ATG} with their errors. The size of the estimated uncertainties varies due to varying sample sizes. The following list explains each input parameter that appears in the table:

Chapter 5: The values estimated in Chapter 5 are listed at the top of the table, for comparison with the modified versions of the code. This is the only entry in which the linear discriminant cut is applied.

Nominal: The table lists the estimated values for three different SNOMAN versions. The FTI tables were produced using SNOMAN release 5_0086.

No Rayleigh: The effect of Rayleigh scattering was tested by turning the process off entirely.

Knock-on electrons: The 2MeV cut on production of knock-on electrons and bremsstrahlung was reduced to the minimum values allowed by EGS and GEANT.

Transit Time Spread (tts): The single photoelectron transit time spread was reduced to zero. The multiple photoelectron timing correction was also disabled, as the multiple photoelectron effect is caused by the single photoelectron transit time spread.

No Prepulse: The simulation of the prepulse was disabled.

Step Size: The maximum allowed muon step size was reduced from 100cm to 0.1cm.

PMT efficiencies: The PMT collection efficiencies were increased by 6.4%. This is roughly equivalent to the difference in the total light between the simulated events and the data in SNOMAN release 5_0183.

PMT angular response: The PMT response at angles larger than 40 degrees was increased by 6.7%. This was done with code graciously provided by Gann [61].

AV attenuation: The attenuation of photons by the acrylic vessel was disabled.

EGS step size: The maximum allowed electron energy loss per step in EGS was reduced from the SNOMAN nominal 3% to 1%.

Wavelength response: The wavelength response was changed so that between 500-520nm the response for 500nm was used, and for wavelengths greater than 520nm the response was: $R_{\text{new}}(\lambda) = R_{\text{nominal}}(\lambda - 20\text{nm})$

Zenith dependence: The up/down asymmetry of the detector was checked by simulating 40GeV muons with isotropic directions. While the distribution of starting positions and directions is isotropic it is not random, the events were simulated in such a way that their distribution of squared impact parameters would be flat. This was necessary in order to prevent an artificial bias in the normalization coefficient. The detector was divided into four zenith angle bins as follows: Zenith bin A is ($0.5 < \cos\theta_z < 1$)

Systematic	SNOMAN version	a_{TG}	N_{ATG}
Chapter 5	5_0184–5	0.9951 ± 0.0004	0.9967 ± 0.0004
Nominal	5_0184–5	0.9976 ± 0.0003	0.9967 ± 0.0004
Nominal	5_0181	0.9962 ± 0.0003	0.9972 ± 0.0004
Nominal	5_0086	0.9956 ± 0.0010	0.9963 ± 0.0013
No Rayleigh	5_0182	0.9943 ± 0.0011	0.9978 ± 0.0016
Knock-on e	5_0182	0.9978 ± 0.0007	0.9980 ± 0.0013
tts=0ns	5_0182	0.9973 ± 0.0007	0.9971 ± 0.0010
no prepulse	5_0182	0.9948 ± 0.0011	0.9994 ± 0.0012
step size	5_0183	0.995 ± 0.003	0.993 ± 0.005
PMT eff	5_0183	0.9979 ± 0.0010	0.997 ± 0.002
angular response	5_0185	0.9989 ± 0.0008	0.996 ± 0.002
AV attenuation	5_0185	0.9966 ± 0.0011	0.9978 ± 0.0018
EGS step	5_0185	0.9979 ± 0.0009	0.9968 ± 0.0015
wavelength response	5_0185	0.9989 ± 0.0008	0.9973 ± 0.0018
Zenith bin A	5_0185	0.9965 ± 0.0008	0.9967 ± 0.0013
Zenith bin B	5_0185	0.9965 ± 0.0017	0.9951 ± 0.0011
Zenith bin C	5_0185	0.9725 ± 0.0024	0.9941 ± 0.0014
Zenith bin D	5_0185	0.9827 ± 0.0019	0.9956 ± 0.0013

Table 6.1: This table shows the estimated acceptance and area normalization factor and their uncertainties for several different modified versions of SNOMAN. The variation of the size of the N_{ATG} uncertainty reflects variation in the sample sizes. The linear discriminant cut was not applied when evaluating the acceptances, as the size of CRMS changed between SNOMAN releases 5_0183 and 5_0184, when a poorly motivated cut on knock-on electrons was removed.

(downgoing), Zenith bin B is (0<$\cos\theta_z < 0.5$), Zenith bin C is (-0.5<$\cos\theta_z < 0$), Zenith bin D is (-1<$\cos\theta_z < -0.5$).

The data in Table 6.1 indicate that the muon analysis is insensitive to the simulation inputs that have been tested. Most of the simulation input entries agree within one or two standard deviations with the nominal values for the nearest SNOMAN release. The statistical uncertainty on the measured neutrino induced muon flux for the entire SNO dataset (D_2O, salt, and NCD phases) is expected to be on the order of 5%. In comparison to that statistical uncertainty the systematic variation due to the changes in the simulation inputs studied are all negligible.

On the other hand, the measurements of the acceptance in the detector symmetry (zenith angle) test show far greater variability than any of the simulation inputs. This change in acceptance is due changes in the fit dE/dx value as a function of zenith angle. The shift in the mean dE/dx value as a function of $\cos\theta_z$ fits well to a first order polynomial with slope -6.7MeV/m. To account for the changing acceptance a 2.4% systematic uncertainty will be applied to the measurement of the muon flux. This shift in fit dE/dx can be accounted for in future SNO analyses.

Comment Regarding EGS4 Step Size

When the EGS4 maximum electron step size was changed from 3% to 1% a slight but noticeable reduction in the CRMS parameter was observed. The reduction was, however, no where near the size required to bring the simulated events in line with the data. A second consequence of the change was that the total number of observed photons was increased by roughly 15%, bringing the data and simulation into far worse agreement in that respect.

6.2 Testing the Data and Simulation Agreement

In Section 6.1 we studied the effect that changing input parameters of the SNOMAN simulation had upon the estimated muon acceptance and fiducial area. In this section we will compare the distributions of fit impact parameters of data and simulated events in order to attempt to constrain the disagreement.

6.2.1 Reconstructed Impact Parameter Comparison at Large Impact Parameters

In the previous chapter we saw that the simulation does not perfectly model the detector. It is conceivable that these imperfections could result in a fitter bias between the data and

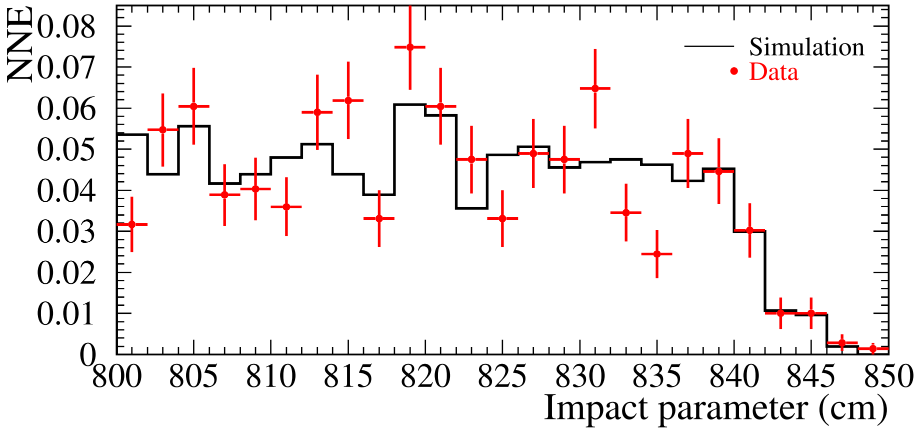

Figure 6.4: The impact parameter distribution for large impact parameter muons is shown for simulated cosmic muon events (black histogram) and the open D_2O data (red points). The area of each histogram has been normalized to one.

simulation. The most important potential fitter bias is radial impact parameter bias, as a differential bias between the data and the simulation would result in different fiducial areas, and therefore different measured neutrino fluxes. With this in mind, one question that should be asked is "how do the fit impact parameter distributions for data events and simulated events compare near the fiducial radius?".

The fit impact parameter distributions for data and simulated events near the edge of the detector are shown in Figure 6.4. For both the data and simulation all of the data selection cuts other than the fiducial radius cut were applied. The simulated events are shown as a black histogram, the data events as red points. The data and simulation are in agreement inside of the fiducial radius cut, and beyond the cut the data and simulation track each other very well. The acceptance of background events was only measured inside of 830cm, so it is possible that some of the data events beyond 830cm are due to instrumental backgrounds. The data and simulated distributions in Figure 6.4 agree very well, with no indication of a significant differential bias.

6.2.2 Testing Radial Bias Using the KS statistic

The Kolmogorov-Smirnov (KS) test is a technique which tests the degree to which two event distributions are likely to have been drawn from the same underlying distribution. The KS test accomplishes this by calculating the cumulative distribution for each of the two distributions and measuring the maximum difference between the two cumulative distributions. After measuring the maximum value of the difference between the cumulative distributions it is possible to assign a significance to that difference. This significance, the KS probability, represents the probability that a random sample of events drawn from the same underlying distribution would return a larger maximum difference.

The *Numerical Recipes* routine "kstwo" [55] was used to calculate the KS probability for the distribution of impact parameters in the data, and for the $\sim$2.5 times expected statistics simulated $\mathrm{D_2O}$ dataset. For the selected events with impact parameters of less than 830cm the KS probability was 0.57, which indicates that the two distributions are consistent with the hypothesis that they are drawn from the same underlying distribution. The KS probability was also measured for the entire range of fit impact parameters, and for the range 800cm$<$b$<$850cm which was shown in Figure 6.4; these tests returned KS probabilities of 0.63 and 0.57 respectively.

The KS probabilities measured above indicate that the two distributions are consistent with the hypothesis that they are drawn from the same underlying distribution, but they do not tell us whether the distributions are also consistent with some other hypothesis. An example of such a hypothesis would be that the data has a small outward radial bias compared to the simulation. One way to constrain these alternate hypotheses is to create fake data sets based upon the simulated events, and compare the KS probability distributions for these fake data sets with the simulated events to the KS probability measured above.

As a first test we will study the effect that a larger impact parameter uncertainty would have. The procedure is to draw an event from the simulated data set, apply a Gaussian

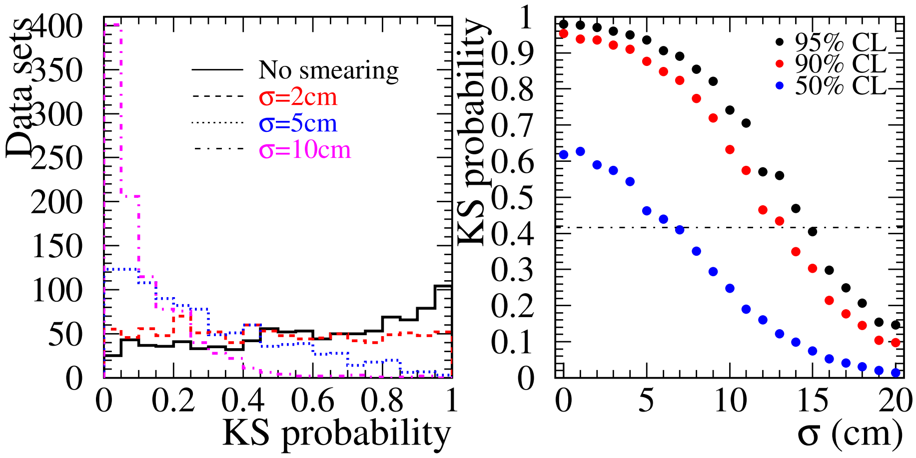

Figure 6.5: These figures show the effects that smearing the fit impact parameter of simulated events has upon the measured KS probabilities. The left hand plot shows the distribution of measured KS probabilities, while the right hand plot shows confidence bands for 50%, 90%, and 95%. The horizontal line shows the KS probability measured in the data in the impact parameter range 0cm<b<760cm.

smearing to the impact parameter, and keep the event if its resultant impact parameter is smaller than the fiducial radius. In order to avoid effects due to sampling near the edge of the data, the fiducial radius is reduced to 760cm. For this impact parameter cut the KS probability that the open D_2O data and the simulation agree is 0.42. The distributions of measured KS probabilities for 1000 fake data sets which had a smearing of 0cm, 5cm, 10cm, or 15cm applied are shown in the left hand plot of Figure 6.5. In the case of no smearing the KS probability distribution should be flat. The observed slope is not surprising, given that rather than sampling from an underlying distribution we sampled from the the set of simulated events that is being tested against. The lowest set of points in the right hand plot shows the median KS probability obtained in the 1000 trials as a function of the size of the smearing. The two higher curves show the 90th and 95th percentiles respectively. Based on these curves we see that a KS probability greater or equal to that measured in the data would occur only 10% of the time if the smearing were as large as 13cm.

The effect of smearing the fit impact parameter is to remove the structure observed in

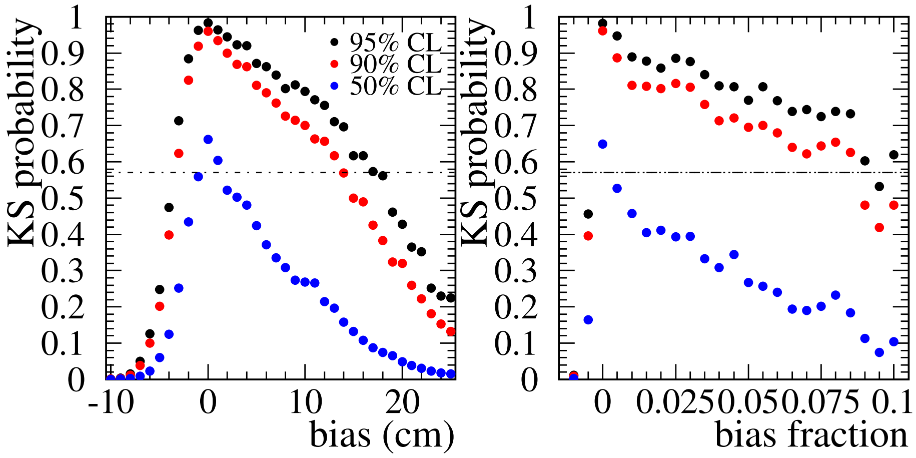

Figure 6.6: These figures show the effects that biasing the fit impact parameter of simulated events has upon the measured KS probabilities. Both plots show the confidence bands for which 50%, 90%, and 95% of the simulated data sets have a smaller KS probability. The left hand plot shows the bands for a constant bias, $b \to b+x_{const}$. The right hand plot shows the bands for a linear bias, $b \to b(1 + x_{frac})$. The horizontal lines show the KS probability measured in the data in the impact parameter range 0cm<b<830cm.

Figure 4.14. With the exception of possible effects near the edge of the detector, the impact parameter resolution is unlikely to have an effect on the measured muon flux. A much more important potential effect is impact parameter bias. In a manner similar to that described above it is possible to study what effect introducing an artificial bias would have upon the measured KS probabilities for simulated data sets. In this case, rather than smearing the data two biases are applied, first a constant bias of x_{const} cm ($b \to b + x_{const}$), and second a fractional bias x_{frac} that is linear in impact parameter ($b \to b(1 + x_{frac})$). The effect of these biases is shown in Figure 6.6.

The left hand plot of Figure 6.6 shows the effect that a constant bias is expected to have on the KS probabilities. One can conclude from this plot that if the data has a constant outward bias there is less than a 5% probability that it is large as 15cm. A 15cm bias would correspond to a bias in the fiducial area of 3.7%. The second plot indicates that the KS test has very little ability to constrain the possibility of a large bias in the impact parameters in the data if that bias is proportional to impact parameter. While the KS test is not able to place a stringent quantitative limit on the maximum potential radial bias, Figure 6.4 makes it difficult to imagine that such a bias could be large.

6.2.3 Comparison of Charge-Only and Time-Only FTI Fits

In 4.3.1 the possibility of running the FTI reconstruction routine in "charge-only" and "time-only" modes was discussed. These different reconstruction modes may be used to test the similarity of simulated and real muon events. Imagine, for example, that for some reason the average speed of light was not the same in the simulation and the data. It would be reasonable to assume that this could have some effect on the time part of the fit, but it should have no impact on the measured charges. In this scenario we would expect that the difference between the time-only and charge-only fit for each event would be different for data and simulated events. Of course, it is possible that some unknown difference between the data and the simulation has a perfectly correlated effect on the charge-only and time-only fits, in which case this test would be insensitive to it. So as was the case for the KS tests, the charge-only and time-only comparisons which will be performed are a consistency test; it cannot prove that a differential bias does not exist.

An additional problem in making this comparison is that the time-only FTI fit begins to occasionally fail at impact parameters larger than 780cm (84.2% of the detector area). The part of the FTI likelihood related to charge provides a very strong constraint on the maximum possible impact parameter, because the total number of detected photons is strongly correlated with impact parameter. With the loss of the charge constraint the fitter is free to greatly reduce the number of tubes in the Cherenkov cone, by pushing the fit to the edge of the detector. The same behavior was observed in the original SNO time residual minimization fitter, FTM, and this restricted the maximum impact parameter used in earlier analyses to 750cm. For the purposes of comparing the charge and time-only fits events will be rejected if the time-only fit returned an impact parameter of greater than 845cm, so the

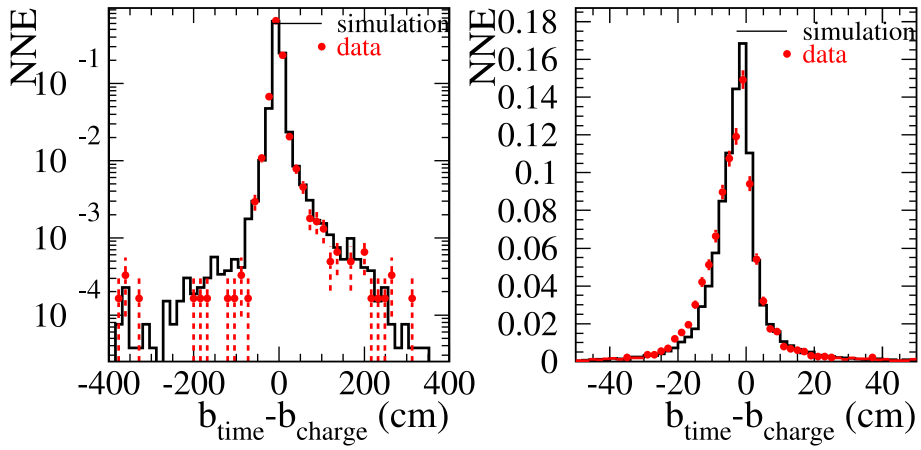

Figure 6.7: The difference between the time-only and charge-only fit impact parameters is shown on a linear and log scale. In each plot the simulated distribution is shown as a black histogram and the data is shown as red points.

results at FTI fit impact parameters greater than 780cm should be taken with a grain of salt.

With those caveats, let us now compare the data and simulation. The distribution of the difference[4] of the time-only and charge-only fits is shown in the plots of Figure 6.7. The simulated distribution is shown as a solid black histogram, the data as red points. The number of events in each plot has been normalized to 1 to account for the different statistics. The difference between the two fits is larger than 50cm for 1.6% of the data events and 2.4% of the simulated events. A difference between the data and simulation is visible in the right hand plot; the distribution of charge-only and time-only differences is slightly narrower in the simulation than in the data, and the means of the two distributions are different by $\sim$1.1cm.

It is reasonable to study further whether the deviation in charge-only and time-only fit impact parameters varies with impact parameter. The charge-only and time-only fit difference is shown as a function of the combined FTI fit impact parameter in Figure 6.8. The difference between simulated and data events seems to be worse for muons that do not pass through the acrylic vessel ($(b_{FTI}/850)^2 > 0.5$ in the plot). This should not be

[4]Time-only minus charge-only.

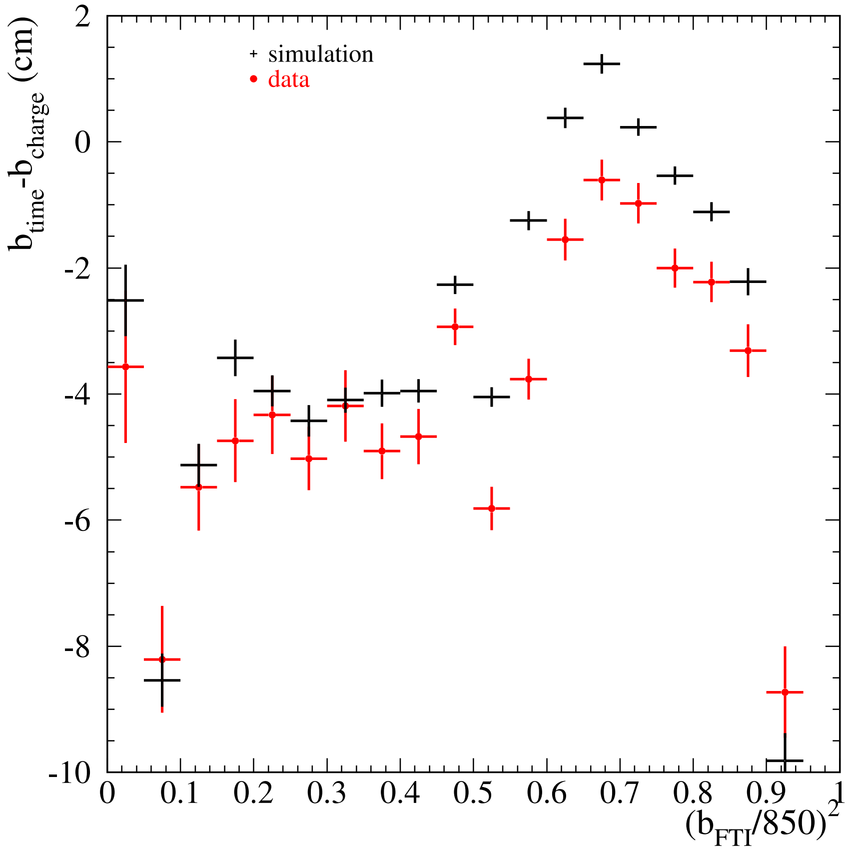

Figure 6.8: The difference between the time-only and charge-only fit impact parameters is shown as a function of the combined FTI fit impact parameter b_{FTI}. The fit difference in the simulation is shown in black, and the difference in the data are shown as red points. The "hump" between 0.55 and 0.85 comes from a changing radial bias in the time-only fitter.

surprising, as the production and transport of light is best understood at the center of the detector.

The source of the 1.1cm difference between the time-only and charge-only fits is not known, so we cannot claim to know the fiducial radius to within less than this difference. At 830cm a 1.1cm difference in impact parameter corresponds to a 0.3% difference in area. Due to this observation the fiducial area normalization term will be taken to be $N_{ATG} = 0.997 \pm 0.003$. This value is consistent with all of the systematic effects tested in Table 6.1. The product of a_{TG} and N_{ATG} is therefore 0.992 ± 0.003, and thus the uncertainty on the flux due to the uncertainties on the cut acceptance and fiducial area is 0.3%.

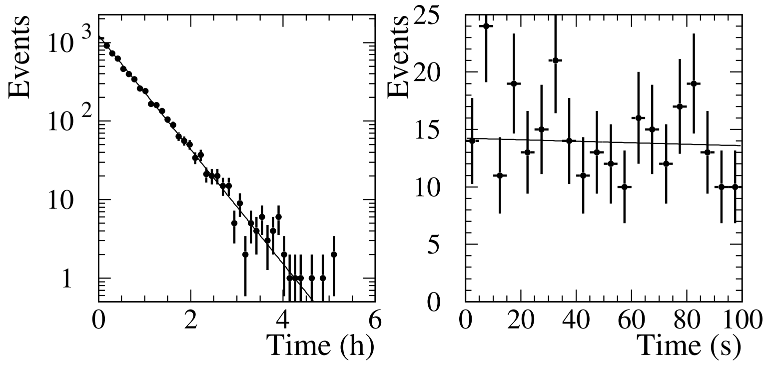

Figure 6.9: The time intervals between subsequent events in the open D_2O data set which pass all of the muon cuts are shown at long time scales (left), and short time scales (right). The histogram on the left was fit with an exponential, which is shown. The same exponential was scaled to account for the different bin size, and is plotted again in the right hand plot.

6.3 Time Intervals Between Muons in the Data

As a final check for instrumental background events in the open D_2O data set, the time intervals between muon events is studied. If the arrival times of muons are truly random then the time intervals between muon events should be exponentially distributed. If any instrumental backgrounds which produce bursts of light at short time intervals (such as PMT wet end breakdown events) are present, they could be expected to produce an excess of events at short time intervals. Alternatively, if some instrumental background is present during only one period of the data taking this could be expected to make the distribution of time intervals non-exponential. The distribution of time intervals between events in the open D_2O data set is shown in Figure 6.9 at long and short time intervals, along with an exponential fit to the long time interval distribution. The exponential fit had a χ^2/dof of 34.7/38, so the data are consistent with the random arrival time hypothesis.

6.4 Summary of Systematic Effects and Uncertainties

In Chapter 5 we observed differences between the simulated and real data. In this chapter we have attempted to understand the possible sources of those differences and to study how much of an effect those differences could have on the through-going muon acceptance and fiducial area. We first studied the effects that changing SNOMAN model parameters had upon the simulation, and we then compared the impact parameters in the data to the simulated events. This comparison was done using the Kolmogorov-Smirnov test, and also by using the time-only and charge-only versions of the FTI reconstruction routine.

The first result of this investigation was the observation that even fairly gross changes to input parameters of the SNOMAN simulation do not appear to have much of an effect on the estimated acceptance and area normalization coefficient. The up-down asymmetry was seen to result in an acceptance uncertainty of $\sim 2.4\%$, considerably larger than the change seen in any of the simulation input tests. The second result of this investigation was that the impact parameter distributions of the data and simulated events are consistent with the hypothesis that they are both drawn from the same underlying distribution. The third result was that there is a difference between the charge-only and time-only impact parameter fits for data events and simulated events of 1.1cm. As a result of the observance of this discrepancy, the uncertainty on the fiducial area was increased by a corresponding factor. The uncertainty on the flux due to the data selection cut acceptance and fiducial area uncertainty is 0.3%. The uncertainties on the flux which are not directly related to the detector will be discussed in the next chapter.

In the next chapter we will use the results of this and all earlier work to study the cosmic muon signals and the neutrino induced muon signal. The depth-intensity relation for SNO will be measured, the production of neutrons from muons will be studied, and finally we will measure the neutrino induced muon flux.

Chapter 7

Neutrino Induced Muon Flux and Other Analyses

In chapters 3 through 5, techniques for calibrating, reconstructing, and identifying muons in the Sudbury Neutrino Observatory detector were described. The previous chapter tested and quantified the extent to which the detector is understood. In this chapter, the results of these techniques will be used to make several measurements. First, we will measure the neutrino induced muon flux at the SNO detector for data from the $\mathrm{D_2O}$ phase of the experiment. We will use this measurement to estimate the atmospheric neutrino flux and the oscillation parameter Δm^2, assuming maximal mixing. Second, the cosmic muon flux intensity will be measured as a function of depth for the $\mathrm{D_2O}$ phase data. Third, the muon flux as a function of time will be checked for the predicted seasonal variation in both the $\mathrm{D_2O}$ and salt data sets. Finally, the detection of free neutrons after a muon has passed through SNO will be discussed, again for both the $\mathrm{D_2O}$ and salt phases.

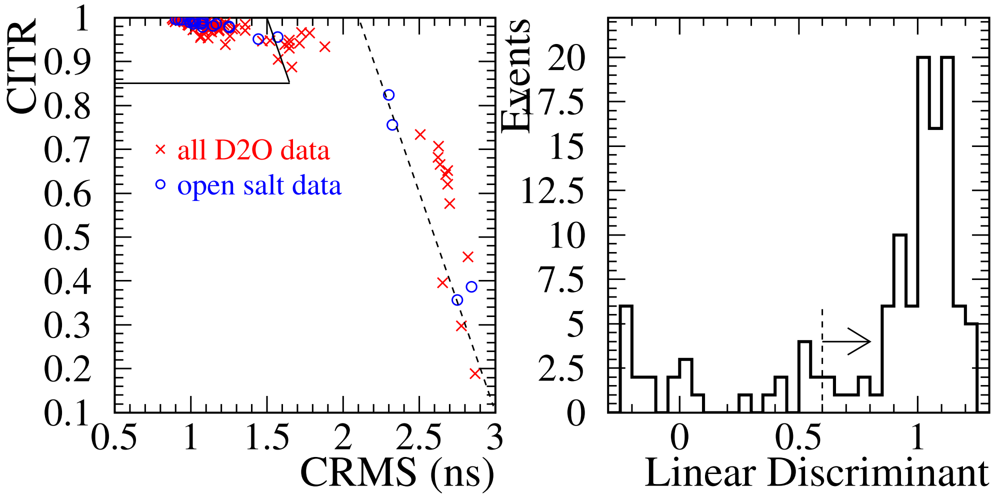

Figure 7.1: The left hand scatterplot shows the Cherenkov cone in-time ratio vs. the cone time residual RMS for events in the open data set. The dashed line shows where the linear discriminant (L) has a value of zero, and the solid lines show the cuts CITR>0.85 and L>0.6. The right hand plot shows the distance of each event from the dashed line in the left hand plot. The dashed line in the right hand plot shows the chosen cut value of 0.6. This figure is consistent with Figure 5.11, which was produced with a smaller sample of events.

7.1 Consistency Check of the Newly Included Data

Before proceeding to the measurements let us quickly perform a consistency check between the open D_2O data studied in Chapter 5 and the entire D_2O data set. In Figure 5.11 the CITR and CRMS through-going muon event selection parameters were compared, and these parameters were combined to form a linear discriminant. The linear discriminant was used to estimate an upper limit on the number of instrumental backgrounds in the signal region. Figure 7.1 shows the same distributions as Figure 5.11, but with all of the D_2O data included. The two figures are consistent with each other, and there is no evidence that a previously unknown background has appeared.

7.2 Neutrino Induced Muon Flux and Oscillation Analysis

In this section, the neutrino induced through-going muon flux will be measured as a function of the cosine of the incoming muon zenith angle. Only data from the D_2O phase of the experiment will be considered. Below a zenith angle of 90 degrees it is expected that $\nu_\mu \rightarrow \nu_\tau$ oscillation will result in a distortion of the $\cos\theta_z$ distribution. Due to the relative sizes of the SNO and Super-Kamiokande detectors, SNO will not have enough statistics to make a competitive measurement of the neutrino oscillation parameters. The SNO experiment is particularly insensitive to the mixing angle, and this is the reason that only maximal mixing will be considered. However, we will see that due to the fact that SNO can observe atmospheric neutrino induced muons at very short baselines the inclusion of the SNO data in a combined SNO+Super-K analysis helps to constrain the oscillation parameter space.

The neutrino induced muon flux will be estimated using a maximum likelihood parameter estimation. The backgrounds that are present are cosmic ray muons and internal atmospheric ν_e and ν_μ interactions. The predicted neutrino induced muon flux and zenith angle distribution will be calculated using the nuance neutrino generator [66], with the flux predicted by the Bartol04 [26] neutrino flux tables specific to the geomagnetic latitude of SNO. The nuance Monte Carlo program is written in lower case, so in this thesis nuance will be distinguished from the word nuance by using a different font. Version three of nuance was used in all simulations. In this version of nuance the deep inelastic cross section is modeled using the prescription of Bodek and Yang [67], the model of Smith and Moniz [68] is the framework for calculating the quasi-elastic scattering cross sections, and resonant processes are modeled using the Rein-Sehgal [69] model.

The production and detection of atmospheric neutrinos was discussed in Section 1.2.5. As a reminder, atmospheric muon neutrinos can undergo a charged current interaction in the rock surrounding SNO, producing a muon (or an anti-muon). These muons are then able to propagate many meters through the rock, and may reach the SNO detector before being stopped in the rock. Muons that enter SNO from outside can either continue propagating all the way through the detector and out the other side, or they can stop inside the detector. The SNO collaboration has not yet developed a stopping muon fitter, so this analysis will only consider the through-going muons.

7.2.1 Generation of Probability Distribution Functions

The nuance neutrino generator was used with the Bartol04 flux predictions for SNO to generate predictions for the neutrino induced through-going muon zenith angle distributions for both SNO and Super-Kamiokande. This subsection explains the generation of zenith angle probability distribution functions (PDFs). These PDFs are later used in a maximum likelihood flux extraction.

Nuance was used to generate two hundred years of simulated data for SNO at neutrino energies between 600MeV and 8.9TeV. A total of 135.18[1] years of simulated data was generated for Super-Kamiokande. At energies below 10GeV the neutrino flux is expected to change as a function of the $\sim$11 year solar magnetic activity cycle. During the "solar maximum" the cosmic ray flux at Earth is reduced, and at low energies the neutrino flux should track the cosmic rays. Predictions for SNO were generated using the flux tables for solar maximum. The Super-K predictions were generated using a ratio of 1.3518 years of solar maximum to 3.1542 years of solar minimum. The Super-K cavity was simulated as described in [10], the SNO cavity was taken to be a cylinder with radius 11m and height 34m. Events were generated both in the rock surrounding the experiments and in the water gap between the cavity wall and the detectors. In both detectors only a few percent of the through-going muon events are due to interactions in the water between the detector and the rock.

[1] This is 30 times the sample reported in [10]

Nuance propagates muons through rock using the PROPMU [70] Monte Carlo simulation package. Only muons with an energy greater than 600MeV (1.5GeV) were considered for SNO (Super-K), as muons with energies smaller than this will not cause through-going muons. Super-K rock was taken to be "standard rock", ρ=2.65g/cm^3, Z=11, A=22. The rock around SNO was simulated with ρ=2.85g/cm^3, Z=12, A=25. Higher density rock attenuates muons more quickly, but due to the increased number of targets there are also proportionally more neutrino interactions. These effects offset each other well; Super-K does not assign a flux uncertainty due to the difference between Super-K rock and standard rock.

Muon Propagation in the Detectors, and Stopping Muons

A quadratic range-energy formula was used to model the propagation of muons through the detectors. This energy loss approximation was parameterized with a fit to SNOMAN simulated stopping muons . Muons in Super-Kamiokande may travel longer distances than can be easily simulated using SNOMAN, so the quadratic fit was constrained to match the data given by Groom [71] at 10GeV kinetic energy. The muon range function fit using this procedure is:

$$\text{Range in cm} = -47.59 + 0.46094 E_\mu - 3.4297 \times 10^{-6} E_\mu^2 \tag{7.1}$$

where E_μ is the total muon starting energy in MeV. The same parameterization is used in this parameter estimation for muons at both Super-Kamiokande and SNO.

Muons were required to have had enough energy to pass completely through the detector in order to be considered a through-going muon event at Super-K. A seven meter minimum track length through the Super-K inner detector was also required, as this cut is used to define the fiducial area in Super-K analyses. Super-K's reported [10] zenith angle dependent detector acceptance was also applied. Muons in SNO were required to have an impact parameter of less than 830cm (corresponding to a minimum track length of 367cm), and were accepted as through-going muons if they had enough energy to propagate entirely through the detector. The acceptance of stopping muons in SNO was discussed on page 153. The acceptance curve fit to the data in Figure 5.16 was used to weight the stopped muon data in generating the PDFs.

The energy distribution of the simulated neutrinos that produced through-going and stopping muons in Super-Kamiokande and SNO is shown in Figure 7.2. In order to be considered for this plot stopping muon events in Super-K were required to have enough energy to propagate at least 7m into the inner detector. All SNO events were required to have an impact parameter of less than 830 cm. Since a stopping muon fitter has not yet been written for SNO no track length cut was imposed in making this plot.

The reason that the curves in Figure 7.2 are not the same shape for the two experiments is the relative sizes of the two detectors. The larger the detector, the larger the energy required for a muon to pass all the way through it. One important observation is that because there are more low energy events the effect of oscillations will be more pronounced in the SNO through-going signal than it is in the Super-Kamiokande through-going signal.

Comparison With Super-K Flux Predictions

The Super-Kamiokande experiment has not reported muon flux predictions generated using nuance and PROPMU, so the predictions in this thesis cannot be compared directly to a published Super-K result. They have, however, reported their predicted number of events with the NEUT neutrino simulation package, with GEANT used to propagate muons from the rock into the detector. Table 7.1 shows Super-K's reported predictions for the number of neutrino induced through-going muon events for three different neutrino flux predictions [72, 73, 26], which they refer to as "Flux A", "Flux B", and "Flux C", respectively. The

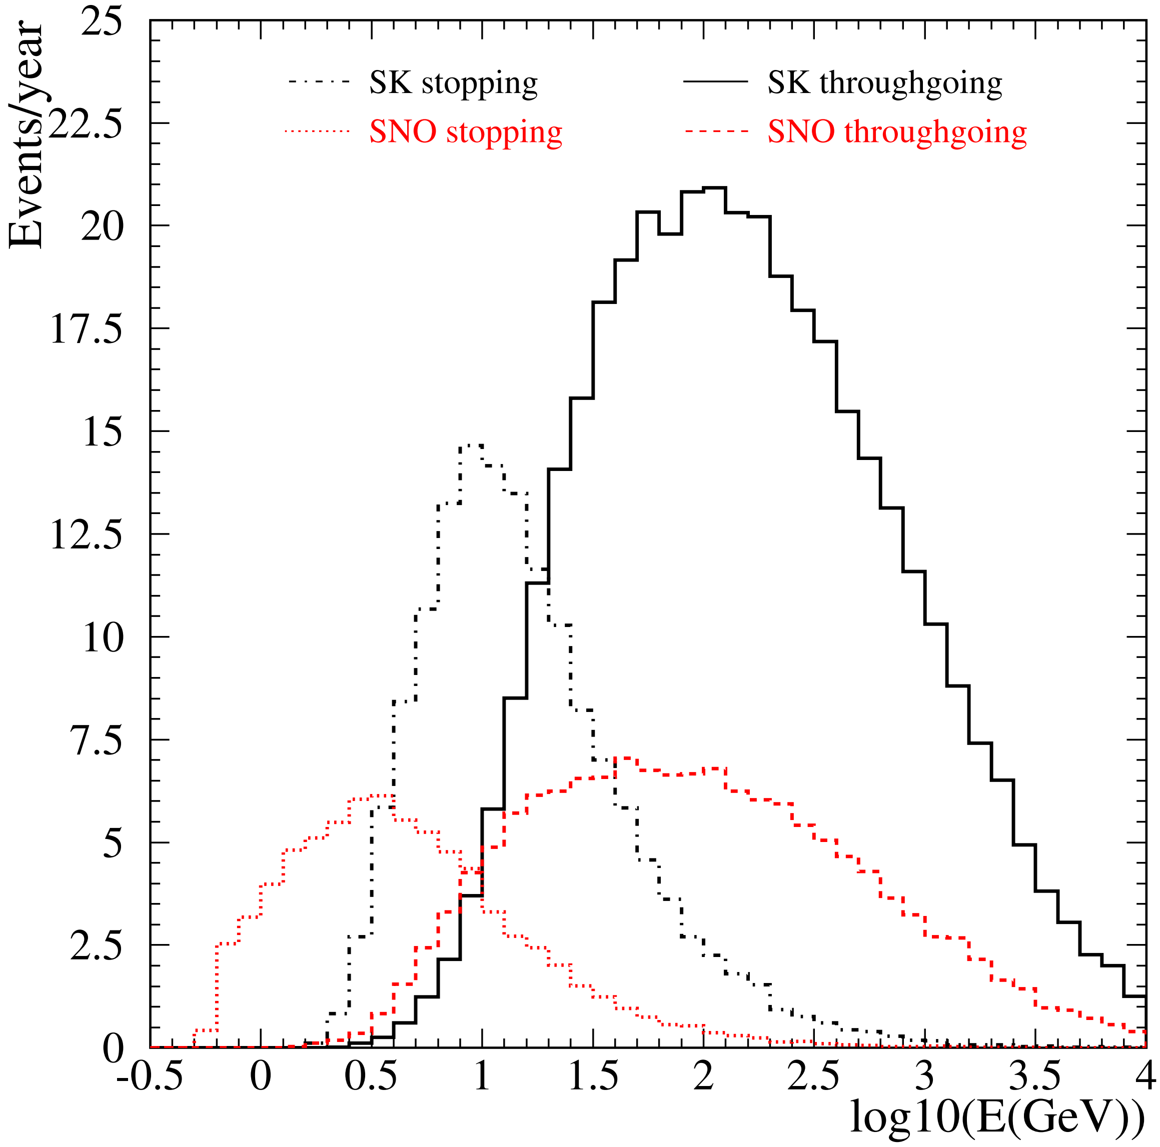

Figure 7.2: This figure compares the energy of nuance simulated neutrino induced through-going and stopping muons at Super-K and SNO. The minimum muon energy at the edge of the detectors is 1.5GeV and 600MeV respectively, so the low side of the stopping muon curves are underestimated. The Super-K events have a 7m track length cut applied, the SNO events have the 830cm fiducial radius cut applied. With the exception of these cuts the acceptance of the detectors has not been taken into account. Oscillations are not simulated.

Flux	Events
Super-K Measurement	1841.6
Super-K Flux A	1669.5
Super-K Flux B	1644.3
Super-K Flux C	1974.9
This work (∼Flux C)	1829 ± 8

Table 7.1: The measured number of neutrino induced through-going muons observed at Super-Kamiokande in 4.506 years is compared to Super-K's three predictions based on different flux models, as well as to the prediction in this work. None of the flux predictions include neutrino oscillations. The fractional Super-K measurement is due to the subtraction of a non-integer number of background cosmic ray muon events. The Super-K predictions come from [10].

Bartol prediction (Flux C) is used in this work. It should be noted that the Bartol table for Sudbury was used in this work for both the SNO and the Super-Kamiokande predictions. Because of the large energy of through-going muon events, the use of the Sudbury specific table should have little effect on the prediction.

The observed flux difference is consistent with the $\sim 10\%$ NEUT/nuance difference that has been found in past comparisons using the Honda flux prediction [74]. At lower energies, Super-Kamiokande reports differences in the predicted flux using the two generators for the fully contained and partially contained muon signals shown in Figure 1.5. These differences range from a 5.1% nuance deficit in the case of "partially contained" events to a 17% deficit for "sub-GeV multi-ring" events. This work used the zenith angle dependent efficiencies reported by Super-K in [10], but it is possible that those efficiencies have energy or position dependencies that are not modeled here, and that these contribute to the difference. This work also used the simple energy loss approximation described above, which neglects both muon range straggling and multiple scattering.

Internal Neutrino Interactions

Neutrino interactions inside the PSUP are a background to a through-going muon analysis in SNO. As was discussed in Section 5.3.7, most such events are fairly low in energy and are easily cut using the dE/dx estimate, which is based on the through-going muon hypothesis. The nuance generator was used to simulate 33 years worth of internal neutrino interactions, and these events were passed through SNOMAN to determine the detector acceptance for such events. The sample that remained after applying the through-going muon selection cuts were used to generate PDFs for internal electron and muon neutrino interactions.

Atmospheric Neutrino Oscillations

With the event samples described above it is possible to generate zenith angle probability distribution functions for the through-going muon signals at SNO and Super-Kamiokande. The data for each event type for both experiments were binned in $\cos\theta_z$ bins of width 0.1. This resulted in 10 bins for Super-K and 14 bins for SNO. The only signal considered in Super-K were the through-going muons, so the event zenith angle distribution was converted to a PDF by simply scaling by the number of events. The different samples in SNO were weighted according to the amount of simulated data used in their generation. Muons that were predicted to stop in SNO were also weighted according to their probability of being accepted as through-going muons.

The neutrino direction and energy were saved along with the muon direction[2], so it is possible to use the same event sample to generate PDFs for any given value of Δm^2. This can be accomplished by simply weighting each event by its survival probability,

$$P_{\nu_\mu \to \nu_\tau} = \sin^2(2\theta_{23})\ \sin^2\left(1.27\Delta m^2_{23}\frac{L}{E}\right) \tag{1.10}$$

[2]FTI fit direction in the case of the internal neutrino interactions in SNO

where E is the neutrino energy in GeV, and L is the distance from the atmosphere to the SNO detector. The neutrino direction is converted to a distance L using the equation:

$$L = \sqrt{((\cos\theta_z^2) \times (R_E - d)^2) + (d + h_a) \times ((2R_E) - d + h_a)} - (R_E - d)\cos\theta_z \quad (7.2)$$

where $\cos\theta_z$ is for the neutrino (not the muon!), R_E is the radius of the Earth, d is the depth of the detector below sea level, and $h_a = 25$km represents the scale height of the atmosphere. For SNO, d was defined as 1.6km, while for Super-K d was set to 0.

Equation 7.2 is an appropriate approximation below $\cos\theta_z = 0$, but above the horizon it is an oversimplification. It is not entirely correct to use the single scale height, h_a, as the interaction height. In principle a Monte Carlo simulation of atmospheric neutrino production above the detector should be used in order to account for the fact that neutrinos are produced over a continuous range of depths in the atmosphere. The distribution of neutrino production heights is especially important for the low energy (sub-GeV) samples at Super-K. This work, however, considers only the neutrinos that produce through-going muons, and at a distance of ≤ 500km these neutrinos have survival probabilities extremely close to one. The effect of changing the depths and atmospheric height used in this equation has less than a 0.1% effect on the flux.

The incoming muon $\cos\theta_z$ probability distribution functions obtained with the method described above are shown in Figure 7.3. Three PDFs are shown for three values of Δm^2 at both SNO (top) and Super-K (bottom). In the context of a combined flux and oscillation measurement it is appropriate to plot the PDFs rather than the number of events, as in such an analysis the normalization is unknown. Note that in Super-Kamiokande most of the information about Δm^2 in the through-going signal comes from the bin nearest the horizon. Although it is not shown, at a large enough value of Δm^2 the oscillated PDFs have the same shape as the PDF for no oscillations. This is because at very large values

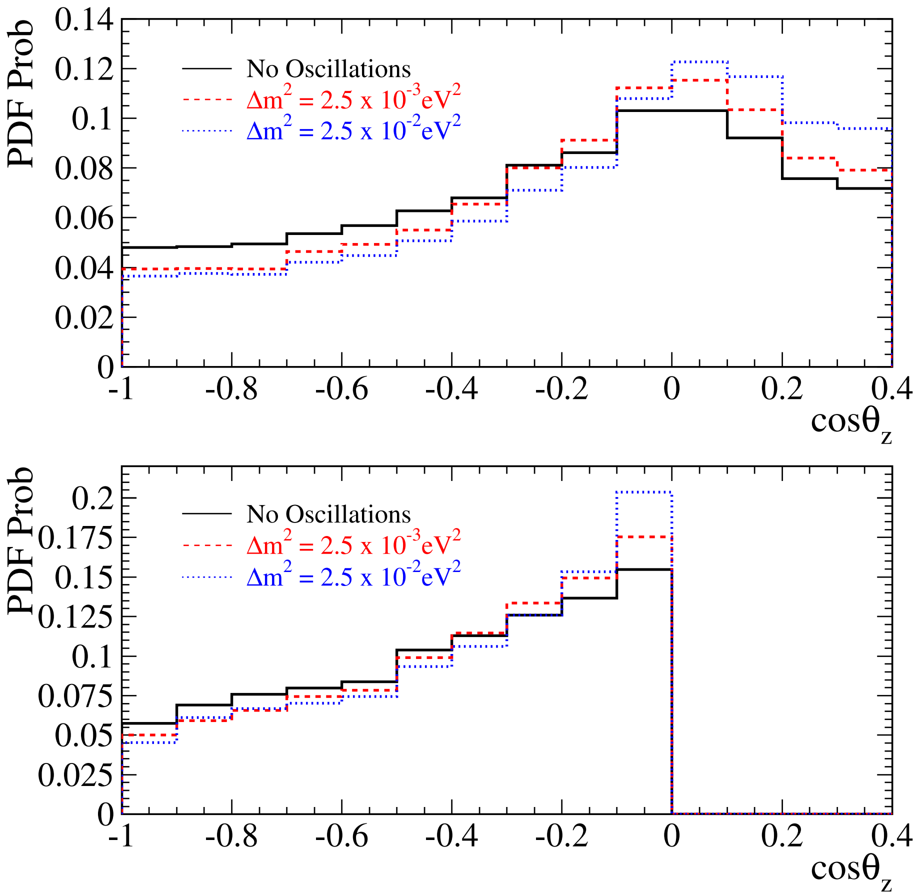

Figure 7.3: The probability distribution functions for incoming muon zenith angle are shown for SNO (top) and Super-Kamiokande (bottom). Three different values of Δm^2 were considered, with $\sin^2 2\theta$ fixed at 1. The small cosmic muon background was not included in the SNO PDF.

of L/E the oscillations are washed out, leaving simply a rate suppression. The value of Δm^2 at which this occurs is larger for SNO than for Super-K, but it is well outside of the Super-K allowed oscillation region for both experiments. The effect of the washing out of the oscillation signal on the "horizontal to vertical" Super-K through-going muon analysis can be seen in Figure 7.4.

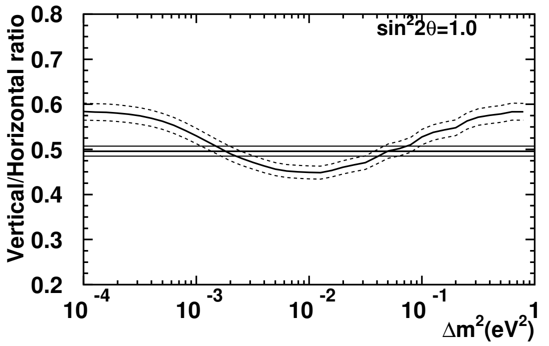

Figure 7.4: In the past Super-Kamiokande has used the ratio of vertical ($-1 < \cos\theta_z < -0.5$) to horizontal ($-0.5 < \cos\theta_z < 0$) muons to estimate Δm^2. In this figure from [10] the measured vertical to horizontal ratio (horizontal line) is compared to the NEUT/GEANT prediction. The bands show the $\pm 1\sigma$ statistical and systematic uncertainties. Note that the ratio is the same for no oscillations and for oscillations with a large Δm^2.

7.2.2 Maximum Likelihood Parameter Estimation

An extended maximum likelihood parameter estimation algorithm was created in order to estimate a set of parameters given a zenith angle distribution at SNO or Super-K. The parameters that are fitted for are the flux (as a multiple of the nuance prediction), $\Delta\mathrm{m}^2$, the number of background cosmic muon muon events in SNO's largest $\cos\theta_z$ bin, the number of internal neutrino events in SNO, and a scale factor Λ for the internal neutrino flux relative to the amount predicted in 5.3.7. The scale factor accounts for the uncertainty on the flux of such events. The fits for background cosmic muons and internal neutrino events are effectively background subtractions.

Given a particular array of the five parameters listed above the parameter estimation routine generates the $\cos\theta_z$ PDF expected for the experiment under consideration. The likelihood for Super-K is defined as

$$\mathcal{L} = \frac{e^{-\phi_K}\phi_K^{N_K}}{N_K!}\prod P_{Ki}(\cos\theta_z|\Delta m^2) \tag{7.3}$$

where $\mathcal{L}$ is the Likelihood, ϕ_K is the expected total number of events at Super-K, N_K is the observed total number of events, and $P_{Ki}(\cos\theta_z|\Delta m^2)$ is the PDF probability for Super-K in $\cos\theta_z$ bin i, given the oscillation parameter Δm^2. The likelihood for SNO is slightly more complicated due to the constraints on the cosmic muon background and the internal neutrino interaction background. The SNO likelihood is given by

$$\begin{aligned}\mathcal{L} &= \left(\frac{e^{-\phi_S}\phi_S^{N_S}}{N_S!}\right)\left(\frac{1}{\sqrt{2\pi}\sigma_I}e^{\frac{-(1-\Lambda)^2}{2\sigma_I^2}}\right)\left(\frac{e^{-\Lambda\phi_I}(\Lambda\phi_I)^{N_I}}{N_I!}\right)\left(\frac{e^{-\phi_\mu}\phi_\mu^{N_\mu}}{N_\mu!}\right)\\ &\quad\times\prod P_{Si}(\cos\theta_z|\Delta m^2)\end{aligned}\tag{7.4}$$

where as before, ϕ_S is the expected total number of events at SNO, N_S is the observed total number of events at SNO, and $P_{Si}(\cos\theta_z|\Delta m^2)$ is the PDF probability for SNO in $\cos\theta_z$ bin i, given oscillations. The first added constraint is on Λ, the systematic difference in flux prediction for internal neutrino interactions and neutrino induced muons. It is constrained with a Gaussian distribution of mean 1 and $\sigma_I = 30\%$. The expected number of internal neutrino interactions from section 5.3.7 is given as ϕ_I, and is modified by Λ. The fit number of internal neutrino interaction events is N_I. The final constraint is on N_μ, the number of cosmic muons fit in the uppermost bin in $\cos\theta_z$. It is constrained to be near the mean expected number from section 5.3.7, ϕ_μ.

While the true number of internal neutrino interaction events must have an integer value, it was found that constraining this value to be an integer considerably reduced the smoothness of the likelihood function. The parameter N_I was therefore treated as a real number rather than an integer, and the factorial term was replaced with a gamma function in the relevant Poisson distribution. A number of constraints were applied in order to avoid problems in evaluating the likelihood at points that could cause numerical problems. The flux was constrained to be between 10^{-5} and 20 times the nuance prediction. The neutrino mass squared difference was constrained to lie in the range $10^{-8}\mathrm{eV}^2 < \Delta m^2 < 100\mathrm{eV}^2$. For the neutrinos under consideration the value $10^{-8}\mathrm{eV}^2$ is indistinguishable from no oscillations. The number of background muons and atmospheric neutrino interactions, and the value Λ were all constrained to be non-negative.

7.2.3 Performance on Simulated Data Sets

In addition to finding the set of parameters that maximizes the likelihood given by Equation 7.4, we would also like to be able to say something about the probability that some other flux and Δm^2 could lead to the observed distribution. The difference of the logarithm of the likelihood at the best fit point from that for some given value of the parameters can be used for such a test. In [75] Cowan gives the log likelihood differences that are expected when performing a fit with n degrees of freedom. Before defining a confidence region for the data using this technique one should test that the data are distributed as expected.

A set of 1000 simulated data sets were generated using the PDFs for neutrino oscillations at SNO with the nominal nuance flux, with $\Delta m^2 = 2.5 \times 10^{-3}\mathrm{eV}^2$, and with $\sin^2 2\theta = 1$. Note that because the PDFs used in the fit were used in creating these data sets, the data sets cannot be considered independent. The simulated data sets were fit to find the parameters that maximize the likelihood given in Equation 7.4, and the logarithm of the likelihood at the best fit point was compared to the input parameters. It was found that the distribution of log likelihood differences of the best fit position and simulated position did not follow any of the distributions predicted by Cowan for n free parameters.

While the likelihood is a function of five parameters, it is not clear that each of these parameters is truly free. The number of cosmic background events fit by the maximum likelihood estimator for 1 year of simulated data is almost always zero, so that particular fit parameter is most certainly not free. The effect on the $\cos\theta_z$ distribution of changing the fit number of internal neutrino interactions or the Λ parameter is similar, although the effect that such a change has on the likelihood is not the same. In addition, changing those parameters can only change the number of events by order 1/10 the total number of events,

CL cut value	In 68%CL	In 90%CL	In 95%CL	In 99%CL
Dependent data set parameters	46%	74%	92%	98%
Cowan 3 free parameters	54%	82%	98%	98%
Cowan 4 free parameters	68%	98%	98%	100%

Table 7.2: The difference of the logarithm of the likelihood at the best fit position and for the true (input) parameters was compared for 50 simulated data sets, using three different log likelihood cuts to define the confidence interval.

so the impact of changing those parameters on the $\cos\theta_z$ distribution itself is quite small.

The 1000 simulated data sets were used to estimate the difference in logarithmic likelihood that corresponds to specified confidence intervals. The differences in log likelihood corresponding to 68%, 90%, 95%, and 99% confidence were found to be 1.4, 2.68, 3.5, and 5.07. Curiously, these were almost exactly halfway between the values given by Cowan for $n = 2$ and $n = 3$.

The confidence intervals defined above were tested using a set of 50 independent simulated data sets. These data sets were created using the PDF for neutrino oscillations at SNO with the nominal nuance flux, $\Delta m^2 = 2.5 \times 10^{-3}\text{eV}^2$, and $\sin^2 2\theta = 1$. The fractions of simulated data sets which fell within the confidence intervals for 68%, 90%, 95%, and 99% using Cowan's prediction for three and four free parameters are given in Table 7.2, along with the fractions of simulated data sets which fell within the confidence intervals defined in the paragraph above. With only 50 simulated data sets it is difficult to draw any conclusions from Table 7.2. Given that the contour plots shown in this thesis are only rough approximations in any case, the confidence level definitions based on the 1000 simulated data sets will be used.

7.2.4 Confidence Intervals for the Super-Kamiokande Data, and for the Combination of the Super-Kamiokande and Simulated SNO Data

The values of the flux and Δm^2 for the Super-Kamiokande through-going muon data that maximize the likelihood given by Equation 7.3 were found to be 1.15 times the nuance default flux and $\Delta m^2 = 2.7 \times 10^{-3}\text{eV}^2$. This can be compared to the Super-Kamiokande global oscillation analysis fit of 1.19 times the NEUT and Honda [72] default flux prediction, and $1.5{\times}10^{-3} < \Delta m^2 < 3.4{\times}10^{-3}\text{eV}^2$ reported in [10]. The approximate 68%, 90%, and 99% confidence intervals for two free fit parameters are shown in the top frame of Figure 7.5. A simulated 1000 day SNO data sample was generated using the nuance default flux and $\Delta \mathrm{m}^2 = 2.5 \times 10^{-3}\text{eV}^2$. A fit to the simulated SNO data returned a best fit flux of 1.02 times the nuance default, a best fit mass squared difference of $\Delta m^2 = 2 \times 10^{-3}\text{eV}^2$, one cosmic muon background event and eight internal neutrino interaction events. The approximate confidence intervals for the simulated SNO data are shown in the bottom frame of Figure 7.5. The maximum likelihood estimation is based upon five parameters, and in making a real contour plot one would fix the flux and Δm^2 and find the best fit likelihood for the remaining three parameters. This would be computationally intensive, and as the goal of this thesis is not to generate the true contour plots a simpler approximate technique was used. During the minimization the maximum likelihood extraction code kept track of the best fit likelihood observed at each point. The likelihood was minimized using simulated annealing, and for each data set the likelihood function was evaluated more than 100,000 times. The approximate contours were generated using the best likelihood observed at a given grid point in Δm^2 and flux during the minimization. These contours are therefore lower limits rather than true confidence intervals. The generation of real contour plots of combined Super-K and SNO data is beyond the scope of this thesis. These “toy” contour plots are shown only as a rough demonstration of how the SNO data helps constrain the parameter space.

Returning to Figure 7.5, we see that as was expected the SNO data does not constrain the parameter space as well as the Super-Kamiokande data do. Let us now see what the effect of combining the two data sets is. The best fit flux and oscillation parameters for the combined Super-K real data and 1000 days of simulated SNO data is 1.12 times the nuance predicted flux and $\Delta m^2 = 2.45 \times 10^{-3}\mathrm{eV}^2$. The allowed approximate confidence intervals are shown in Figure 7.6.

This study considered 1000 days of SNO data, but in this thesis only the 337.25 days of the $\mathrm{D_2O}$ phase will be studied. The impact of the reduction in statistics is dramatic. With such a short exposure time statistical fluctuations in the highest and lowest zenith angle bins can move the best fit value of Δm^2 over several orders of magnitude. As was mentioned in section 2.5, data taking at the SNO experiment is expected to end on November 20, 2006. At that point the total through-going muon exposure time is expected to be near the 1000 days considered in the study above.

7.2.5 Systematic Uncertainties on the Neutrino Induced Muon Flux

As the reader may have guessed from the discussion in the last subsection, the measurement in this thesis will be entirely limited by statistics. There are a number of known systematic uncertainties on the muon flux, and before looking at the $\mathrm{D_2O}$ through-going muon data let us take a moment to consider their relative importance. The systematic uncertainties on the measurement of the neutrino induced muon flux are given in Table 7.3, along with the page number referencing where the systematic was described or measured. The maximum likelihood parameter estimation algorithm was used to evaluate two uncertainties that will be discussed below.

Some of the uncertainties (e.g., acceptance) are easily expressed in percentages, while

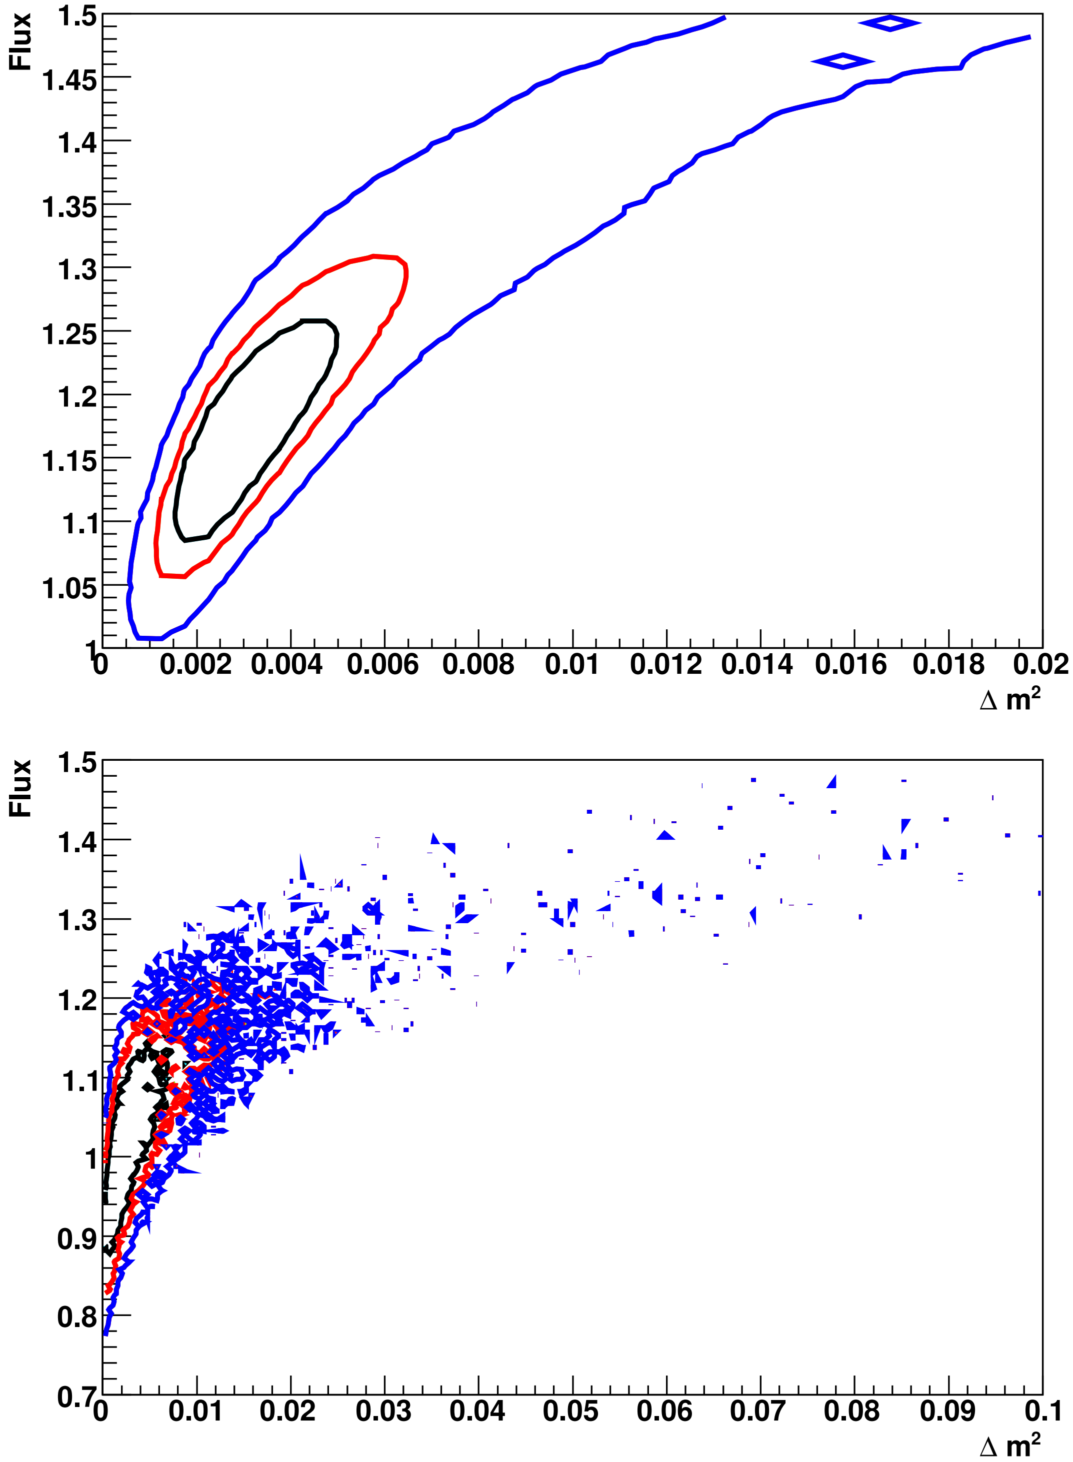

Figure 7.5: This figure shows the 68%, 90%, and 99% allowed regions in flux vs Δm^2 that the Super-Kamiokande data (top) and a simulated 1000 day sample of the SNO data (bottom) allow. Note the vastly different scale used to plot the SNO data. Due to statistical limitations the SNO data are not competitive on their own.

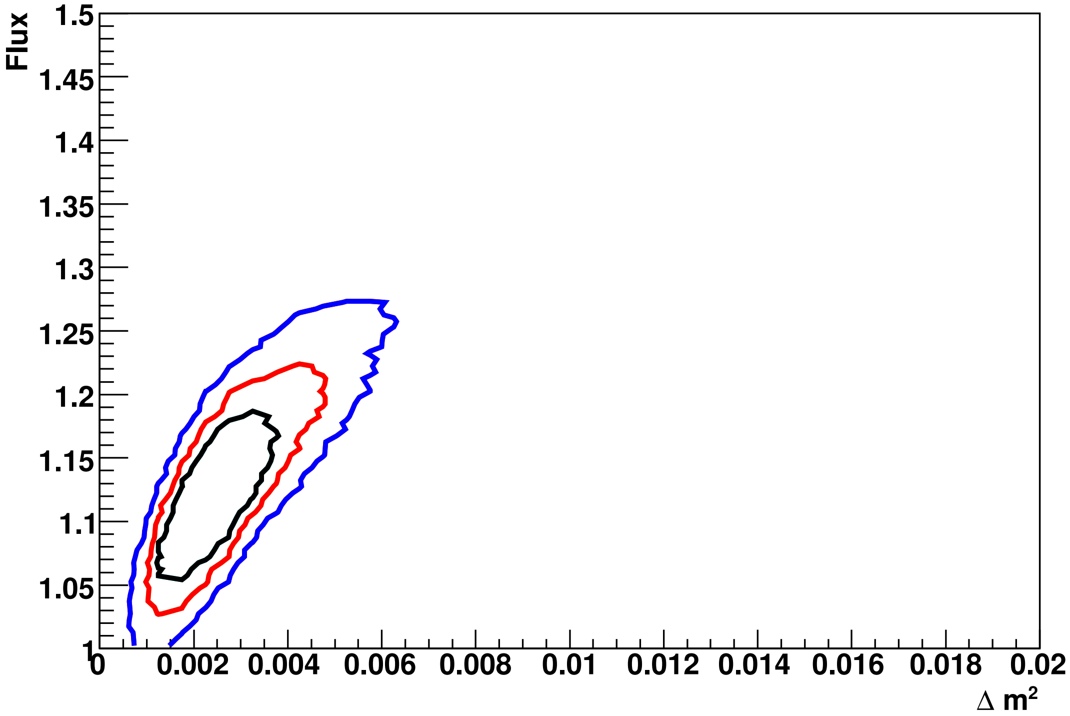

Figure 7.6: This figure shows the 68%, 90%, and 99% allowed regions in flux vs Δm^2 that are allowed by the combined Super-Kamiokande data and a simulated SNO 1000 day data set.

Systematic	See page	Parameter Value	Flux Uncertainty
Acceptance	153	$99.51 \pm 0.04\%$	0.04%
Fiducial Area	180	$2164200 \pm 5700\text{cm}^2$	0.3%
Angular Resolution	202	$0.32°$	$< 0.08\%$
Detector Asymmetry	173		2.4%
Stopping Muon Acceptance	202		0.3%
Live Time	156	$337.25 \pm 0.02\text{days}$	0.006%
Internal ν Bkgd	146	$9.9 \pm 3.0\text{events}$	2.7events
Cosmic μ Bkgd	149	$0.6 \pm 1\text{event}$	1event
Instrumental Bkgd	144	$< 0.004\text{events}$	$< 0.004\text{events}$
Total Uncertainty		$2.4\% + 3.2\text{events}$	3.1%

Table 7.3: The systematic uncertainties on the neutrino induced flux measurement are presented. In order to convert uncertainties in the number of events to approximate uncertainties as a percentage of the flux it was assumed that the data set would have 150. Note that the internal neutrino background will be slightly smaller than is reported here for any given non-zero value of Δm^2.

others (e.g., internal neutrino interaction background) were estimated as a number of events. It is desirable to present the total uncertainty on the flux as a single percentage, but this total actually depends upon the neutrino oscillation parameters! In order to convert the uncertainties on event numbers to approximate "percentage on the flux" uncertainties it was assumed that the D_2O data set would have 150 events.

Angular Resolution

The resolution of the FTI reconstruction algorithm has essentially no effect on the flux, at least in terms of smearing neutrino induced muons between different bins. The resolution was tested by smearing the simulated events used to generate the PDFs by five degrees (an "absurdly large" value), and looking at the difference in the expected flux.

Stopping Muon Acceptance

Muons that pass through a very large fraction of the detector are difficult to distinguish from through-going muons. The acceptance of stopping muons was discussed on page

153. To briefly re-cap, the probability of mistaking a stopping muon for a through-going muon was parameterized as a function of the fraction of the distance that the muon propagated. In an earlier study than the one presented in this thesis the author estimated the same parameterization with a different simulated data set. That data set had an impact parameter distribution that was not flat in b^2. Regardless of this fact, the estimated stopping acceptance parameterization obtained from that data set was fairly similar. The difference in the expected flux predicted using the two parameterizations was taken as the "worst case" estimate of the uncertainty on the through-going muon flux.

Uncertainties on the Conversion of the Neutrino Induced Muon Flux to a Neutrino Flux

Table 7.3 presents the systematic uncertainty on the measured neutrino induced muon flux. In converting the muon flux to a neutrino flux there are a few additional systematic uncertainties. Due to limitations in processing time, the predicted number of events at SNO has some uncertainty. The predicted neutrino induced muon flux at SNO without neutrino oscillations is 144.9 ± 0.9 events per year. The cross sections for neutrinos on rock have some uncertainty due to the models, and additional uncertainty due to uncertainty on the axial vector mass (M_A). The SNO neutrino induced muon flux uncertainty due to uncertainty in the cross section models has been estimated by Formaggio as $^{+7.8\%}_{-6\%}$ [76], while the uncertainty due to M_A has been estimated as 2.2%. The cross section uncertainty is energy dependent, and a minimum muon energy of 4.5GeV was assumed. The acceptance of the SNO detector to through-going muons reaches 50% at $\sim$3GeV, so given the energy distribution in Figure 7.2 the 4.5GeV cut is reasonable. Different muon propagation codes result in different flux predictions for a beam of muons passing some distance through rock. The uncertainty in the flux due to long range muon propagation at many TeV may be as large as 10-15% [77], but given the distribution in Figure 7.2 we are well below that high energy

Systematic	Parameter Uncertainty	Flux Uncertainty
Muon Flux Measurement	3.1%	3.1%
MC statistics	0.6%	0.6%
Neutrino Cross Sections		$^{+7.8\%}_{-6\%}$
Axial Mass Uncertainty	$15.\%$	2.2%
Muon Propagation		$< 2\%$
Seasonal Variation	$< 2.2\%$	1%
Total Uncertainty		$^{+8.9\%}_{-7.4\%}$

Table 7.4: The uncertainties on the conversion of the neutrino induced muon flux to a neutrino flux are summarized. The first entry is the uncertainty on the flux measurement itself from Table 7.3.

regime. A value of 2% is taken as a conservative upper estimate of the uncertainty due to the simulation of muon propagation. Finally, while the data taken in the $\mathrm{D_2O}$ phase was spread over one and a half years, the seasonal variation in the height of the atmosphere discussed on page 17 presumably has some residual effect. As an upper limit the uncertainty on the flux due to the residual seasonal variation will be taken as 1%.

7.2.6 Results

After applying all of the cuts defined in Chapter 5 the number of events classified as through-going muons in the $\mathrm{D_2O}$ phase is 149. The zenith angle distribution of these events is shown in Figure 7.7. Histograms normalized to 149 events show the predictions based on the simulated PDFs for no oscillations and oscillations with maximal mixing and $\Delta m^2 = 2.5 \times 10^{-3}\mathrm{eV}^2$ are shown for comparison.

The zenith angle data shown in Figure 7.7 was input to the maximum likelihood parameter estimation algorithm. The maximum likelihood was found at a flux of 1.2 times the nuance default, with $\Delta m^2 = 3 \times 10^{-3}\mathrm{eV}^2$, zero cosmic muon background events, nine internal neutrino background events, and $\Lambda = 0.96$.

A joint fit of the SNO and Super-K data returned a best fit likelihood at a flux value of

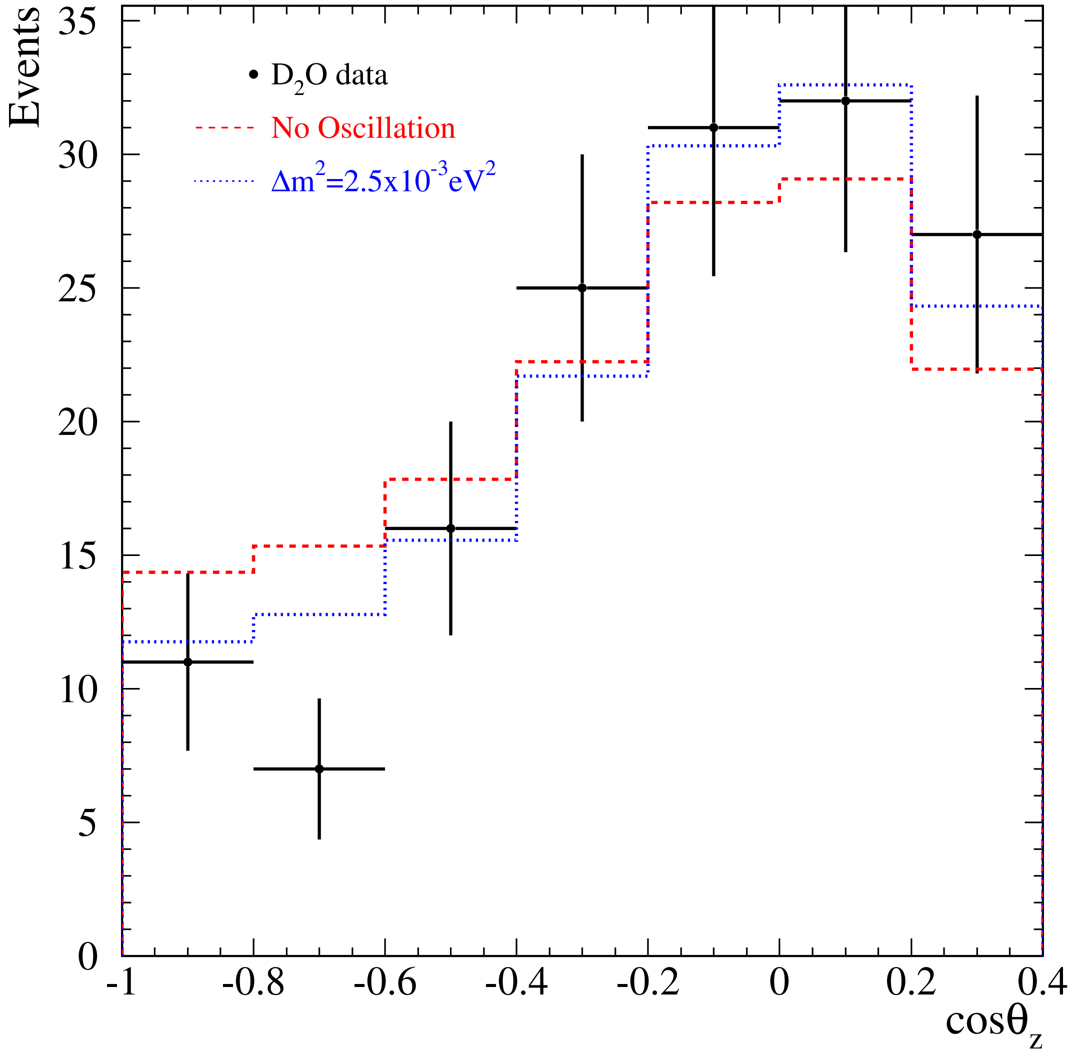

Figure 7.7: The $\cos\theta_z$ distribution of incoming neutrino induced muon events is shown. Histograms normalized to the observed 149 events show the predictions (including the internal neutrino background) for no oscillations and oscillations with maximal mixing and $\Delta \mathrm{m}^2 = 2.5 \times 10^{-3}\mathrm{eV}^2$.

1.16 times the nuance default, with $\Delta m^2 = 2.7 \times 10^{-3} eV^2$, zero cosmic muon background events, nine internal neutrino background events, and $\Lambda = 0.98$. The approximate confidence intervals for the joint fit are compared to the fit for Super-K only in Figure 7.8. It should once again be stressed that these contour plots are to be interpreted as “toy” studies, as the likelihood has not been minimized at each position.

Integral Muon and Neutrino Flux

Subtracting nine internal neutrino events from the 149 candidate through-going muons results in a total of 140 muons, so the measured integrated through-going muon flux is $2.54 \pm 0.21(\mathrm{stat}) \pm 0.08(\mathrm{syst}) \times 10^{-13} \mu/(\mathrm{cm}^2\,\mathrm{s}\,\mathrm{sr})$. This is significantly larger than the flux observed at Super-Kamiokande, $1.661 \pm 0.040(\mathrm{stat})^{+0.011}_{-0.013}(\mathrm{syst}) \times 10^{-13} \mu/(\mathrm{cm}^2\,\mathrm{s}\,\mathrm{sr})$). There are two reasons for the discrepancy. First, the integral flux is dominated by the events near the horizon, and SNO views more events in this range. Second, due to the fact that SNO is a smaller detector the energy threshold required for a muon to be considered through-going is smaller than the energy threshold in Super-K. The minimum possible kinetic energy that a muon would need to be considered through-going in SNO is $\sim$700MeV, while the 7m track length requirement in Super-K means that their minimum energy is $\sim$1600MeV. This effect makes the ratio of through-going to stopping muons larger in SNO than it is in Super-Kamiokande, so even if only events below the horizon are considered, the flux of *through-going* muons in SNO is still systematically larger.

7.3 Muon Depth to Vertical Intensity Relation

A standard measurement reported by underground experiments is a measurement of the “vertical muon intensity” as a function of distance propagated through “standard rock”. The experiments must convert their measured intensity in local rock to “vertical intensity”

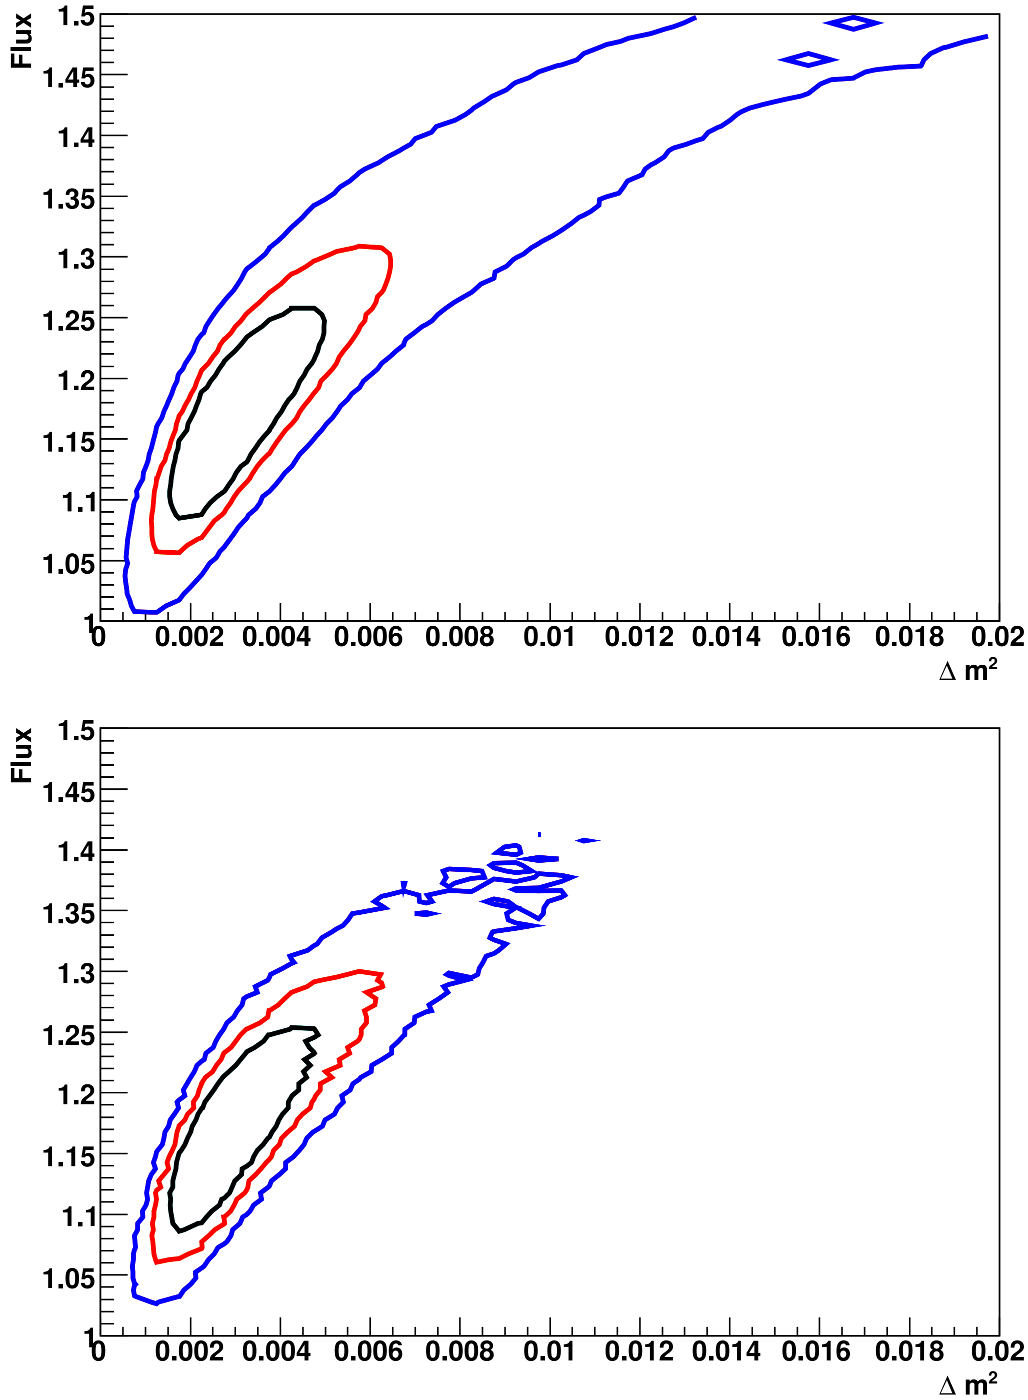

Figure 7.8: This figure shows the 68%, 90%, and 99% allowed regions in flux vs Δm^2 that are allowed by the combined SNO and Super-Kamiokande data. The top frame is for Super-Kamiokande only, and is identical to the top frame of Figure 7.6. The bottom frame shows the result of combining the SNO and Super-K data.

in standard rock in order for the comparison to be meaningful. As was discussed in Chapter 1, the flux of cosmic ray muons is larger for muons traveling horizontal relative to Earth's surface than for muons traveling vertically downward. For $\theta_z < 60°$ the flux goes as $\sec\theta_z$. The conversion from measured muon intensity to vertical muon intensity is thus simply:

$$\mathrm{I}^{\mathrm{v}}_{\mu}(\cos\theta_z) = \mathrm{I}_{\mu}(\cos\theta_z) \times \cos\theta_z \tag{7.5}$$

Standard rock is defined as $\mathrm{CaCO_3}$ with a density of 2.65g/cm^3. Differences in the local rock density and Z/A from that of standard rock are accounted for by adjusting each event's slant depth. In SNO this conversion has been estimated [78] as

$$x = (1.015 \pm 0.01)x_{\mathrm{SNO}} + \frac{x^2_{\mathrm{SNO}}}{4 \times 10^5\,\mathrm{m\,w.e.}} \tag{7.6}$$

Using Equation 7.6 we can define a set of bins in "standard rock depth," which due to SNO's flat overburden can easily be transformed into bins in $\cos\theta_z$. We will represent the muon intensity measured for muons which passed through a range of standard rock $x_i < x < x_{i+1}$ as $\mathrm{I}_{\mu}(x_i)$. In this case the muon intensity is defined as:

$$\mathrm{I}_{\mu}(x_i) = \frac{N_i}{\mathrm{T}\,\Omega_i\,\varepsilon\,\mathrm{A}} \tag{7.7}$$

where N_i is the number of muons observed in bin i, T is the live time, Ω_i is the solid angle of bin i, ε is the product of the acceptance (a_{TG}) and area normalization (N_{ATG}) that were defined in 5.3.8, and A is the fiducial area of the detector. Equation 7.7 can be be changed from muon intensity to vertical muon intensity by using Equation 7.5 to weight each muon event:

$$\mathrm{I}^{\mathrm{v}}_{\mu}(x_i) = \frac{1}{\mathrm{T}\,\Omega_i\,\varepsilon\,\mathrm{A}} \sum_{j=1}^{n} \cos\theta_{zj} \tag{7.8}$$

where j has been introduced as an event index that sums over all muon events in bin i.

Due to the increasing extinction of muons with depth, the statistics available at large slant depths becomes small. For this reason the range of depth used to create each bin was varied. The bin definitions are shown along with the measured vertical intensities in Table C.1 of Appendix C. The table also reports the mean distance traveled by muons and the number of muons observed in each bin. It should once again be noted that the correction to vertical is only valid out to an angle of $\sim 60°$, so the data beyond 12.5km w.e. should be treated with skepticism.

The vertical intensity data for the SNO $\mathrm{D_2O}$ phase are shown in Figure 7.9, with data from two experiments located in the Gran Sasso lab overlaid. The SNO data are shown with statistical error bars only. The most important systematic is related to the multiplicity of muon events. The muon reconstruction algorithm SNO does not measure the number of muons in the event, but both the LVD [79] and MACRO [80] experiments do. Their reported intensity represents the total number of muons, while the SNO intensity represents the number of muon events. The difference between the intensity of single muon events and total intensity was measured as 10% by LVD [79]. Estimating the expected multiplicity of muon events at SNO would be a very non-trivial amount of work, so 10% is taken as a systematic uncertainty on the SNO measurement. Figure 7.10 shows the same data, with SNO's average neutrino induced muon flux added for comparison. The neutrino induced muon data have not been scaled to vertical intensity.

The conversion from muon direction to depth assumed that SNO was at a depth of 2092m, but in reality the muons enter the SNO cavity at a shorter distance. The requirement that they pass all the way through SNO adds an additional ~4-18m w.e. of depth. The overburden of SNO is not perfectly flat, and no information about the variation of the overburden is used in calculating the muon distances. This uncertainty is equivalent to 0.3% on the depth of the bin centroids. Furthermore, in this analysis the curvature of the

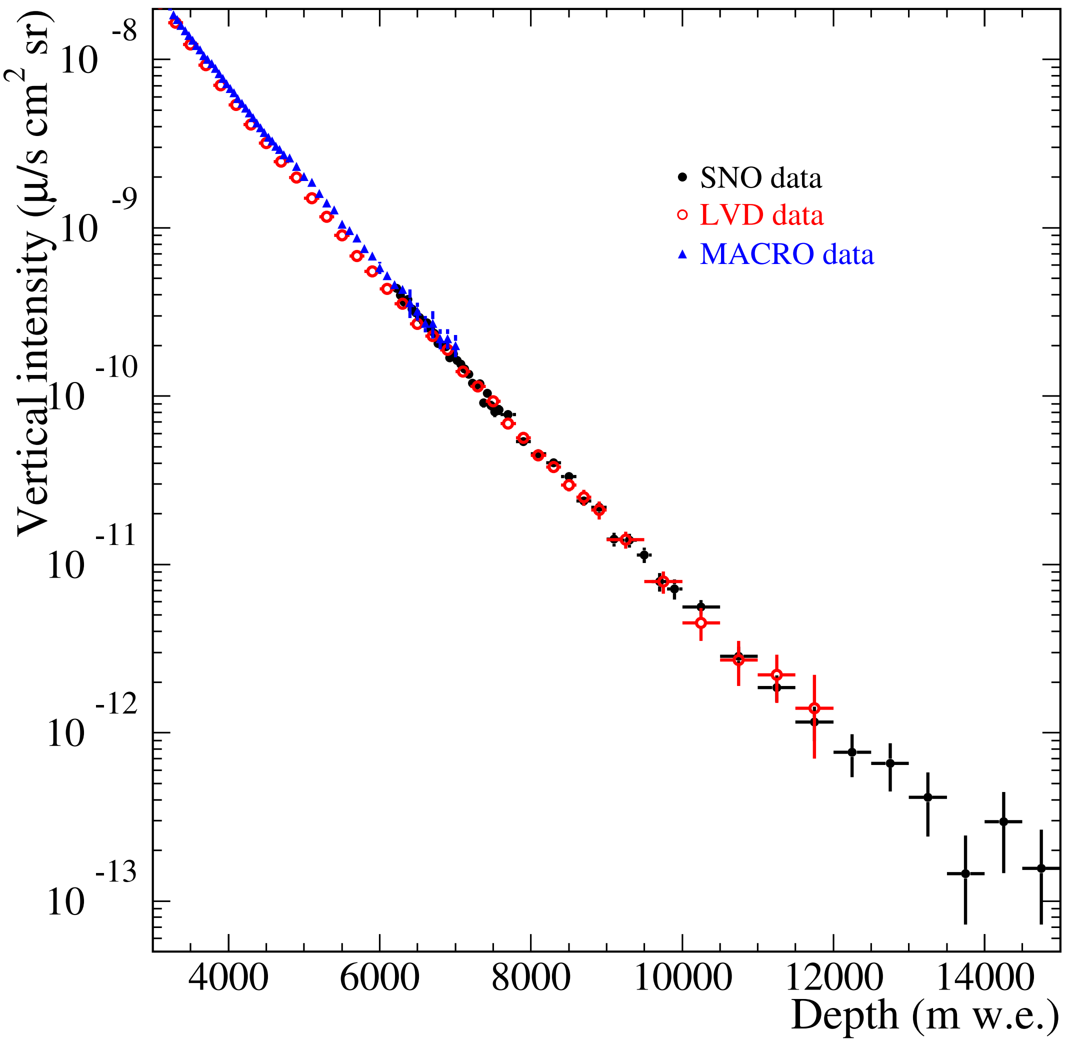

Figure 7.9: The vertical muon depth intensity relation is shown for SNO, with LVD [79] and MACRO [80] data overlaid for comparison.

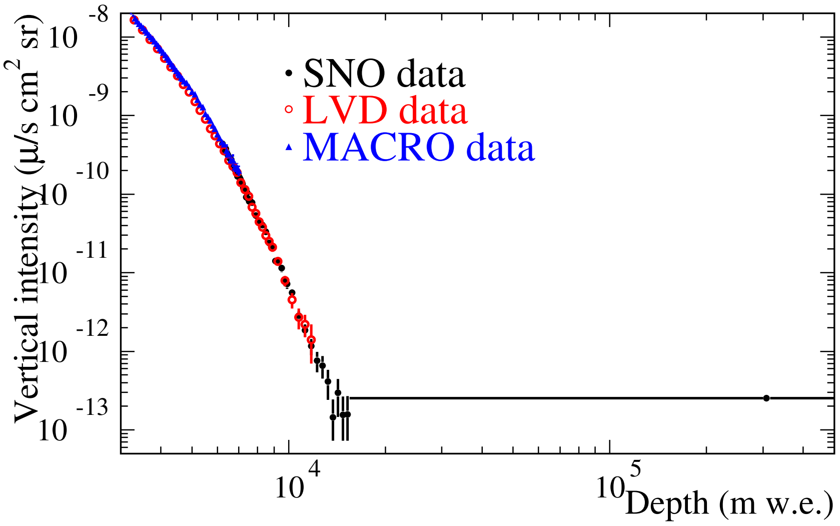

Figure 7.10: The vertical muon depth intensity relation is shown on a log-log scale for SNO, with LVD [79] and MACRO [80] data overlaid. The SNO neutrino induced muon data are also shown as a single wide bin. The neutrino induced muon intensity has not been scaled to vertical.

Earth has been ignored. At 12.5 km w.e. at the SNO site the change in distance due to the curvature of the Earth is on the order of 0.07%, so it is not important compared to the other systematic uncertainties. Finally, the average muon distance traveled is not necessarily the same as the central value of the bin, but varies on the order of 0.1%.

The vertical muon intensity curve is often fit with the function

$$\mathrm{I}^{\mathrm{v}}_{\mu}(x) = A\left(\frac{x_o}{x}\right)^{\alpha} e^{(-x/x_o)} \tag{7.9}$$

The values of $A = 1.73 \pm 0.14 \times 10^{-6}$, $x_o = 2.19 \pm 0.03\mathrm{km\ w.e.}$, and $\alpha = 5.21 \pm 0.05$ (statistical uncertainties only) fit the SNO data well with a χ^2/dof of 48.4/48. These fit values are quite different from those measured by other experiments, but it should be noted that there is a large correlation between the parameters. Fixing $\alpha = 2$ results in a fit of $A = 2.89 \pm 0.03 \times 10^{-6}$, $x_o = 1.14 \pm 0.02\mathrm{km\ w.e.}$, but with a far worse χ^2/dof of 69.9/49. This is still not in agreement with LVD's measurement of $A = 2.15 \pm 0.08 \times 10^{-6}$,

$x_o = 1.16^{+0.06}_{-0.03}\mathrm{km\ w.e.}$. One can see that due to the correlations between A and α, the parameters of this function may change well outside of their estimated uncertainties without having a great effect on the predicted intensity over six orders of magnitude.

With the fit to Equation 7.9 we can now convert the uncertainty in the mean position of the bin to a flux uncertainty. The uncertainty varies with the slant depth as $\delta \mathrm{I}^{\mathrm{v}}_{\mu}(x) = e^{0.01x/x_o} \times 1.016\%$. At 7km w.e. this is ∼5%, at 10km w.e. this is ∼6.4%. The total systematic uncertainty on the vertical muon intensity is thus in the range 11-12%.

7.4 Muon Event Rate as a Function of Time

As was mentioned in 1.2.5, there is a well known seasonal variation effect in the cosmic ray muon rate. This variation correlates well with atmospheric temperature [81], and is due to the changing pion decay length to interaction length ratio with changes in the height of the top of the atmosphere. At higher energies pions have longer lifetimes in the Earth's reference frame, and for this reason higher energy pions "see" a steeper atmospheric density gradient than lower energy pions do, and higher energy pions experience more interactions. The size of the seasonal muon rate variation is therefore larger at larger energies. The muon event rate as a function of time since the start of SNO data taking is shown for muon data from the $\mathrm{D_2O}$ and salt phases in Figure 7.11. The data were fit to a flat distribution plus a sinusoid with a period of one year. The result of the fit was a baseline rate of 62.3 ± 0.5 events per day, with a sinusoid amplitude of $2.6 \pm 0.9\%$ and an offset of 4.2 ± 0.8 months. The fit returned a χ^2/dof of 44.2/43. For comparison, a flat fit returned χ^2/dof=60.4/45.

As was been discussed above, the size of the seasonal rate effect is expected to increase with increasing muon energy. For example, the amplitude of the effect at Super-Kamiokande is $\approx 1\%$ [82]. The flat overburden of SNO makes it possible to look selectively in different muon energy windows by studying the rates observed in different zenith angle

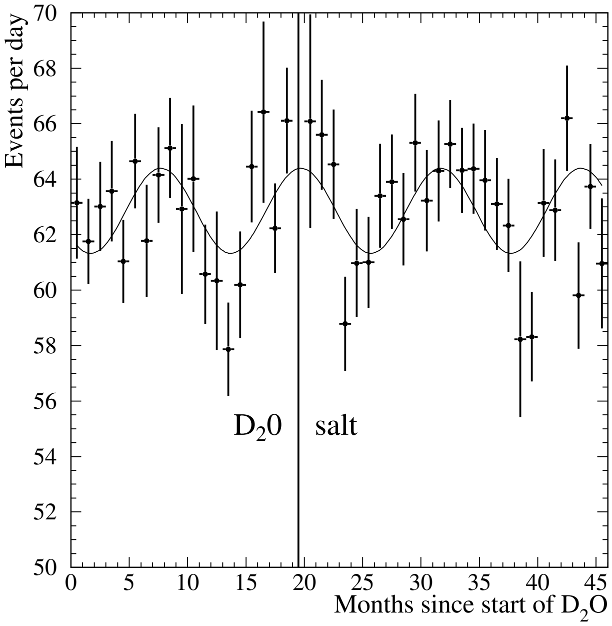

Figure 7.11: This figure shows the through-going muon event rate as a function of time since the D_2O phase data taking started on November 2, 1999. A sinusoid fit with a fixed period of one year is overlaid. Note that the zero is suppressed.

bins. The SNO rate data were separated into four bins, and the resultant fits are shown in Figure 7.12 and summarized in Table 7.5. In each monthly bin the uncertainty in the number of events was taken to be the square root of the number of events. The mean number of events in the lowest $\cos\theta_z$ window was too small to make this Gaussian approximation to the Poisson distribution uncertainty, so the uncertainties in the smallest $\cos\theta_z$ bin are overestimated. The observations summarized in Table 7.5 offer support for the hypothesis that the seasonal rate variation has an energy dependence, but due to statistical limitations the data cannot be considered definative evidence.

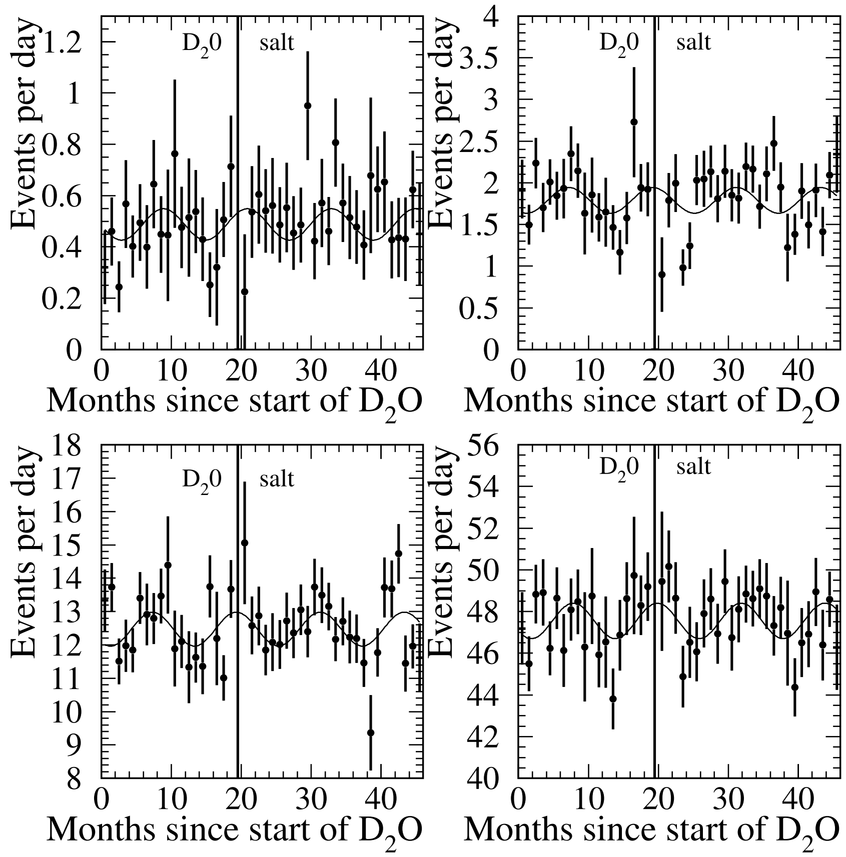

Figure 7.12: The through-going muon event rate is shown as a function of time for four different zenith angle bins. The zenith bins are $0.4 < \cos\theta_z < 0.6$ top left, $0.6 < \cos\theta_z < 0.7$ top right, $0.7 < \cos\theta_z < 0.85$ bottom left, and $0.85 < \cos\theta_z < 1$ bottom right. Note that the zero is suppressed for the two bottom plots.

$\cos\theta_z$ Bin	Evts/day	Amp (%)	χ^2/dof	Prob(%)	Flat χ^2/dof	Prob(%)
All Sky	62.9 ± 0.3	2.4 ± 0.6	$44.2/43$	42.1	$60.4/45$	6.2
$0.4 - 0.6$	0.488 ± 0.024	12.6 ± 6.8	$27.0/43$	97.3	$30.3/45$	95.4
$0.6 - 0.7$	1.8 ± 0.5	8.6 ± 3.6	$58.4/43$	5.9	$64.1/45$	3.2
$0.7 - 0.85$	12.47 ± 0.12	4.1 ± 1.4	$49.4/43$	23.3	$58.2/45$	9.0
$0.85 - 1$	47.6 ± 0.2	1.8 ± 0.7	$32.7/43$	87.3	$39.3/45$	71.1

Table 7.5: The result of fits to the data in Figure 7.12 are shown in this table. The second column lists the fit mean rate and the third column lists the fit amplitude in percentage. The χ^2/dof and probability of observing a larger χ^2/dof are listed for both the sinusoid fit and the fit to the flat hypothesis.

7.5 Neutron Production as a Function of Muon dE/dx

7.5.1 Neutron Production Hypothesis

The technique by which dE/dx is estimated for through-going muons was discussed in section 5.3.4. At the start of that section muon energy loss in water was also discussed. Briefly, muons lose energy at a fairly constant rate due to continuous ionization of the surrounding media, and in occasional bursts due to radiative processes. If a muon radiates a bremsstrahlung photon above an energy of 2.2MeV, it is possible that the photon will photodissociate a deuteron in the heavy water. The probability of dissociating a deuteron should scale with the radiative component of dE/dx. Because the ionization component of dE/dx is fairly constant, the excess of estimated dE/dx over the 200MeV/m lost through ionization should be proportional to the radiative energy loss rate, $\frac{dE_r}{dx}$. Therefore, one can reasonably expect that neutron production (and hence detection) should be proportional to the estimated value of dE/dx.

7.5.2 Muon and Follower Selection

Muons were selected for this analysis using the cuts listed in Table 7.6. While fewer cuts are used here than were used for the other analyses in this chapter, the number of events inside 830cm in the $\mathrm{D_2O}$ data set is the same as was found in Table 5.2. As was discussed in section 2.5, neutrons are detected in the $\mathrm{D_2O}$ phase by their capture on deuterium and the subsequent production of a 6.25MeV gamma. In the salt phase neutrons capture on Cl, resulting in the production of multiple gammas with a total energy of 8.6MeV. The capture time of a free neutron in the $\mathrm{D_2O}$ phase of the experiment is ~ 40ms, while the capture time in the salt phase is ~ 5.5ms [83]. In this section a "neutron follower" is defined as a physics event that occurs within some time window after a muon selected with the cuts

Name	Cut Value
Burst	See text
Neck tubes	< 4 fire
Impact parameter	< 840 cm
dE/dx	> 200 MeV/m
CITR	> 0.85
TRICH	> 0.7
Linear Discriminant	> 0.6

Table 7.6: Cuts used to select muons for the neutron follower analysis. For a description of the cuts see sections 5.3.3 and 5.3.4.

Name	Cut Value
D_2O Time Since Muon Event	< 200 ms
Salt Time Since Muon Event	< 50 ms
Total Energy	> 4 MeV
Event Radius	< 550 cm
PMT "In Time Ratio"	> 0.55
Event Isotropy Parameter	$-0.12 < \beta_{14} < 0.95$

Table 7.7: Cuts used to select neutrons for the neutron follower analysis. An explanation of the PMT in time ratio cut and the isotropy selection criteria can be found in [58] and [84] respectively.

in Table 7.6. Follower events were saved if they fell within a 400ms time window, and considered neutrons if they passed the selection cuts and fell within 50ms after the muon in the D_2O phase and within 5ms after the muon in the salt phase. The parameters used to select physics events at $\sim$MeV energies have not been discussed in this thesis, but are listed for reference in Table 7.7. Briefly, the "ITR" cut works in a similar manner to the CITR parameter in the muon analysis, by measuring what fraction of the hit PMTs fire within a certain time window. The "isotropy parameter[3]" provides a measure of how isotropic the distribution of light in the event is. The ITR and isotropy event selection criteria are described in the theses of McCauley [58] and Dunmore [84] respectively.

[3] β_{14} in SNO jargon.

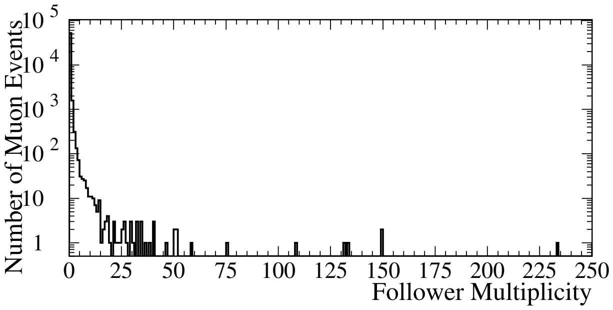

Figure 7.13: The follower multiplicity distribution is shown for all muon events in the D_2O and salt phases. The events with very large numbers of followers are believed to be due to photonuclear interactions.

7.5.3 Muon Follower Multiplicity

The distribution of the "neutron multiplicity", the number of events that pass the cuts following a muon event, is shown for both phases in Figure 7.13. The muon selection criteria were loosened when making this plot, because it was found that the known very high multiplicity events were cut by the muon data selection cuts. These events are characterized by extremely large light levels in comparison to normal muons, so it is not surprising that the data selection cuts removed them. The neutron multiplicity is not expected to be Poisson distributed, as photonuclear interactions experienced by muons near or above the muon critical energy in water may result in the production of large numbers of neutrons. The large fluctuations in neutron multiplicity are due to fluctuations in the production of neutral pions, as the electromagnetic decay of the pions can result in a large number of gammas with energies above 2.2MeV. The reconstructed estimate of muon dE/dx for the very high multiplicity events will be shown later in this section. The multiplicity distributions at low multiplicity and with all of the cuts listed in Table 7.6 (with the exception of the multiplicity cut) is shown for each phase individually in Figure 7.14. It can be seen that the average multiplicity is larger in the salt phase. This is due to the increased detection efficiency for neutrons.

The effect of the low energy cuts upon the number of followers is shown in Figure

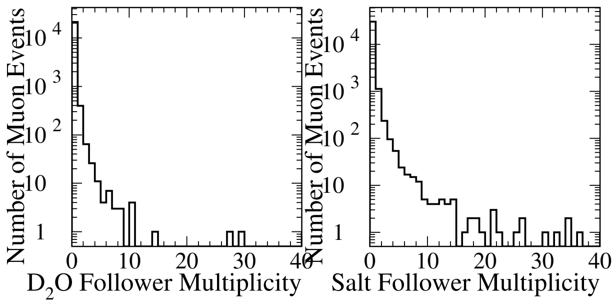

Figure 7.14: The follower multiplicity distribution is shown for events with fewer than 40 followers in each phase.

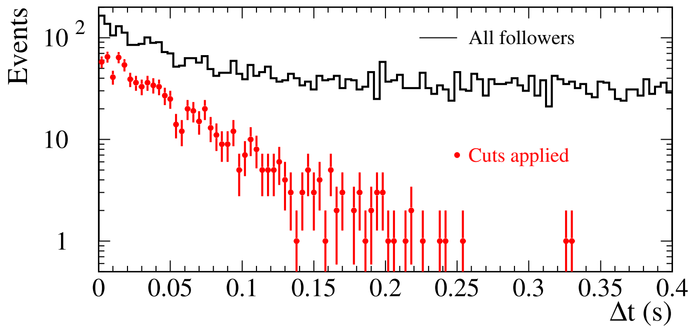

Figure 7.15: The time distribution of events following muons in the $\mathrm{D_2O}$ phase is shown for events with loose selection cuts (black histogram) and the cuts listed in Table 7.7 (red points). The majority of the events were removed by the energy cut.

7.15. The histogram shows the time distribution of all events following muons passing a relaxed set of cuts, and the red data points show the distribution for events passing the cuts in Table 7.7. Almost all of the events that were removed were cut by the 4MeV total energy requirement.

The distribution of follower times relative to the muon event is shown for both phases of the experiment in Figure 7.16. Both of the curves are well fit with an exponential distribution, with a χ^2/dof of 47.0/47 for the $\mathrm{D_2O}$ phase and 7.5/10 for the salt phase[4]. The distribution of neutron followers as a function of $(r/600cm)^3$ where r is the radius of the event is shown in Figure 7.17.

Let us now compare the follower events to the muons that preceded them. Since the

[4]An exponential fit to the salt data with finer binning returned a χ^2/dof of 65.6/79.

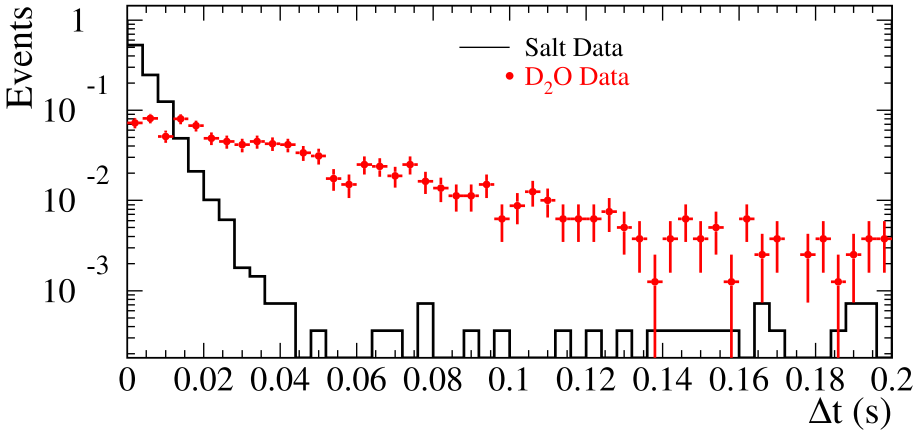

Figure 7.16: The distribution of muon follower times is shown for the $\mathrm{D_2O}$ (red points) and salt (black histogram) phases. The neutron capture time was greatly reduced with the introduction of the salt.

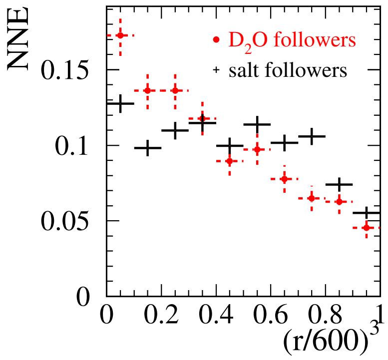

Figure 7.17: The reconstructed position of muon followers is shown for the $\mathrm{D_2O}$ (red dashed points) and salt (black points) phases. The histograms have been normalized to equal unit area, and the abscissa of this plot is the volume weighted radius, $(r/600cm)^3$. Neutrons that wander into the $\mathrm{H_2O}$ are quickly captured on hydrogen, leading to the slope observed in both phases. The increased capture cross section in the salt phase reduces the effect.

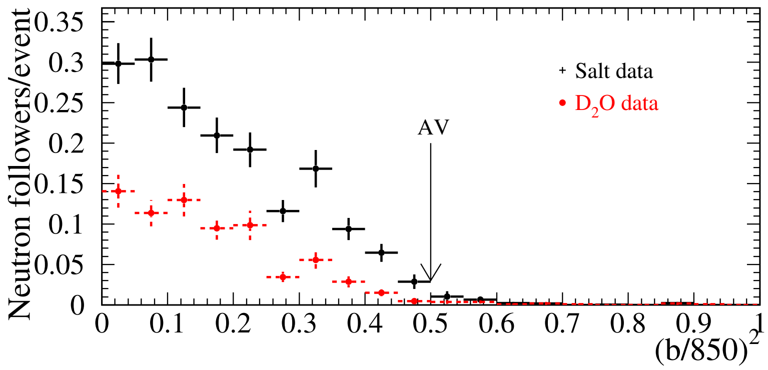

Figure 7.18: The neutron follower multiplicity is shown as a function of the reconstructed muon impact parameter squared. The figure demonstrates that neutron followers are only present after muon events in which the passed through (or at least very near to) the D_2O .

neutron events are constrained to be within 550cm of the center of the detector it is reasonable to expect that only muons that pass through the D_2O are likely to produce followers. The number of followers is shown as a function of muon impact parameter in Figure 7.18. The figure demonstrates that the behavior is as expected. It is very clear that the low energy physics events that follow muon events are neutrons, so finally we will check whether or not the neutron follower multiplicity is actually correlated with the reconstructed dE/dx. Since we now know that there are very few neutron followers for muons that do not pass through the D_2O, we will only consider muons that have an impact parameter of less than 600cm.

The neutron follower multiplicity is shown as a function of dE/dx over a 200MeV/m range for the D_2O and salt phases in Figure 7.19. In order to reduce the importance of very high multiplicity events, only events with fewer than 11 followers were considered. A linear fit was performed for each of the profile histograms in the figure, returning a χ^2/dof of 18.1/16 in the D_2O phase and 17.6/18 in the salt phase. The fit parameters were

$$N_f = -0.134 \pm 0.016 + (6.1 \pm 0.6) \times 10^{-4} \times \mathrm{dE/dx} \tag{7.10}$$

$$N_f = -0.295 \pm 0.018 + (1.40 \pm 0.08) \times 10^{-3} \times \mathrm{dE/dx} \tag{7.11}$$

for the D_2O and salt data shown on page 222 respectively, where N_f is the number of followers. This observation is therefore consistent with the hypothesis that the radiative energy loss of muons is directly responsible for the production of neutrons through the photodisintegration of deuterons.

The linearity observed in Figure 7.19 breaks down if high multiplicity events are included, and also at very large values of dE/dx. The follower distribution over a larger range and without multiplicity cuts is shown in Figure 7.20. Roughly 88% of the muons which pass through the D_2O in SNO have a fit dE/dx of less than 400MeV/m, so the statistics beyond 400MeV/m are poor.

7.5.4 Summary

In this section we have seen that neutron production in the range of 200-400MeV is linearly related to the estimated value of dE/dx. In section 5.3.4 we saw that the estimated dE/dx is related closely to the true value of dE/dx for most muon events, so this result indicates that neutron production through photodissociation is indeed proportional to muon energy loss.

7.6 Conclusions

This thesis has performed a study of through-going muon events at the Sudbury Neutrino Observatory. The two sources of muons at the SNO site are the "downward going" cosmic muons, which occur only in the range $0.4 < \cos\theta_z < 1$, and neutrino induced muons, which occur over the entire range of $\cos\theta_z$. In order to observe these events it was necessary to develop techniques for the calibration, reconstruction, and selection of through-going muons. The SNO experiment does not possess a calibration source for these ~GeV-TeV muons, so the analyses presented in this thesis depend upon simulation to determine efficiencies and reconstruction uncertainties. The simulation was tested by varying a set of the

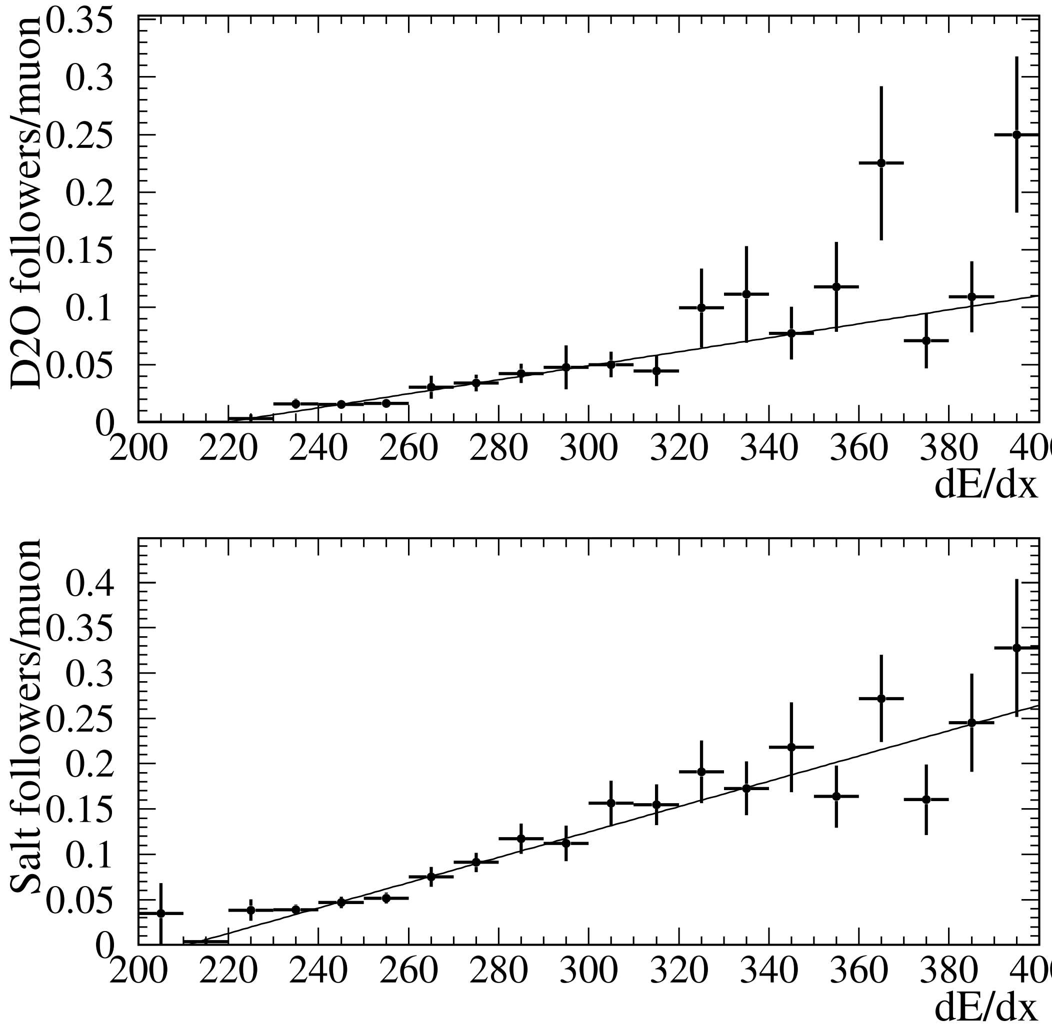

Figure 7.19: The number of muon followers is shown as a function of dE/dx for muon events that pass through the D_2O, along with a linear fit. Events with more than ten followers were excluded in order to avoid bias due to large multiplicity events.

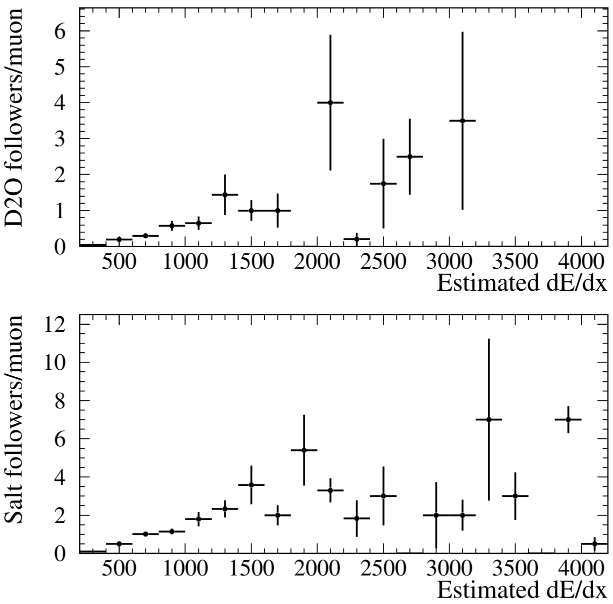

Figure 7.20: The number of muon followers is shown as a function of dE/dx over a far larger range than was shown in Figure 7.19. Only muons which passed through the D_2O are included, but no cut was applied on the maximum number of muon followers.

simulation input parameters, such as the amount of Rayleigh scattering. These tests and comparisons of simulated cosmic muons to the cosmic muons in the data indicate that the simulation does a good job of modeling the detector.

The techniques for selecting and reconstructing muons were used to study neutrino induced muons in the D_2O phase of the experiment. The observations are most consistent with two neutrino oscillations with a neutrino squared mass difference of $\Delta m^2 \sim 3 \times 10^{-3} eV^2$, in good agreement with the Super-Kamiokande results. With less than one year of data the statistical significance of this measurement is small, and while the data are most consistent with oscillations they can not rule out the no-oscillation hypothesis. In order to demonstrate the effect that the addition of SNO data can have on global neutrino analyses the SNO and Super-K data were fit together, and in a set of "toy" contour plots it was shown that the SNO data should help to constrain the high end of the allowed Δm^2 region.

The depth to vertical intensity relation was measured for the cosmic muons in the SNO

D_2O phase, and the results were in good agreement with previously published measurements. A study of the cosmic muon event rate as a function of time showed strong consistency with the expected seasonal rate variation. The amplitude of the seasonal oscillation at SNO was observed to be larger than has measured at other underground sites, and this is believed to be due to SNO's extreme depth. Finally, the muon energy loss estimation that was developed as a part of muon event selection was used to study the relationship between deuteron photodissociation and muon dE/dx. A strong correlation between the neutron multiplicity and the muon dE/dx was observed for muons with energy losses below 400MeV/m.

This thesis also serves as a touchstone document for the SNO through-going muon analysis. The techniques developed in this thesis will be used by the SNO collaboration to perform a neutrino induced muon flux measurement with an expected 1000 live days of data. While SNO will not have statistics comparable to the Super-Kamiokande through-going muon sample it is expected that this additional data will help to constrain the possible atmospheric neutrino oscillation parameter space. More importantly, it will help improve estimates of the atmospheric neutrino flux, which will be the dominant background to the cosmic neutrino signals that the future km^3 neutrino detectors will attempt to observe.

Appendix A

Glossary of Terms

A number of terms and abbreviations are defined once and then used repeatedly through this thesis. This glossary serves as a reference for such terms.

AV The **A**crylic **V**essel separates the D_2O from the H_2O.

Concentrator cut off A light concentrator surrounds each PMT to increase the effective photocathode coverage. The concentrator also prevents the PMT from observing some of the light water, to reduce the detection efficiency of H_2O and PMT radioactive backgrounds.

Critical energy The critical energy is defined as the energy at which a particle loses the same amount of energy through ionization and radiative processes.

CRMS Time residual RMS of PMTs inside the Cherenkov cone, see page 129.

Deck The detector deck is above the detector cavity. The front end electronics and calibration equipment are housed on the deck.

Fiducial Area The area of the detector in which the analysis is performed. See page 125.

FTI A muon reconstruction algorithm in SNOMAN, see 4.3

FTR A muon reconstruction algorithm in SNOMAN, used as a seed for FTI. See 4.2

FTM A time-only muon reconstruction algorithm in SNOMAN, described in [52]

Gain Scaled Unit unit of charge, see page 42

Laserball Optical calibration source used in the SNO detector, produces light at specified wavelengths with a few hundred picosecond time spread. See Page 40.

MC **M**onte **C**arlo

NHit **N**umber of PMTs **H**it in an event

NNE **N**ormalized **N**umber of **E**vents

PDF **P**robability **D**enstity **F**unction

Photon Bomb Monte Carlo technique used in SNOMAN in which a set of photons are produced with user specified positions, energies, directions, and times. See page 56

PMT **P**hoto **M**ultiplier **T**ube

Prepulse PMT pulses created by photons which strike the dynode rather than the photocathode, see page 33

PSUP **P**hototube **SUP**ort structure

QHS High gain short integrate PMT charge measurement, see page 36

QLL Low gain long integrate PMT charge measurement, see page 36

RMS **R**oot **M**ean **S**quared

Run Integer number representing an interval of time during which the detector was online and in a known condition

SNO **S**udbury **N**eutrino **O**bservatory

SNOMAN **SNO** **M**onte Carlo and **AN**alysis software package

Time residual The hit time of a PMT with the photon propagation time subtracted from it.

WEB PMT **W**et **E**nd High Voltage **B**reakdown

Appendix B

Possible Improvements to this Analysis

At several points in this thesis minor problems were discovered and discussed. These size of these problems was not large enough to affect the statistically limited neutrino flux and oscillation measurement in 7.2. These problems are summarized here, along with the author's suggestions for changes that could be made for future analysis based upon the work presented in this thesis. This appendix is intended to summarize these suggestions for the SNO collaboration muon working group, and is unlikely to be of interest to others reading the thesis.

B.1 Reconstruction

The interpolation tables used by the FTI algorithm were generated with a version of SNOMAN that did not simulate the outgoing angles of bremsstrahlung photons correctly. The "refusal" of the fitter to reconstruct events near some table points is believed to be due to events with high energy bremsstrahlung being included in the tables. While re-creating the tables is not expected to have any impact on any of the physics analyses the author feels it should be fixed for aesthetics.

The choice of speed of light and Cherenkov angle used to define the PMT time residuals was not particularly well motivated. In principle the appropriate choice of speed of light could be different for muons that pass through the acrylic zero, one, or two times. The ultimate modification would be to take out the charge based multiple photon timing correction and replace it with a discriminator walk correction and an integral over the possible arrival times of each photon. This, however, is almost certainly more computationally intensive than is reasonable.

It would be interesting to modify the hypothesis in FTI from a single muon with a given entrance position and direction to two (or more) muons with different entrance positions but with the same directions. Some multiple muon events have muon separations large enough to be seen by eye, and it is likely that FTI would be able to fit them. If anyone attempts this it is seriously recommended that an impact parameter dependent minimum number of photons detected from any individual muon be applied to the fit.

On a very technical level, in the second temperature lowering in the SNOMAN routine fti_fit_event_sim_an it appears that the best fit point is not replacing the worst fit point. The variable spx_worst_i should be calculated as it is earlier in the code. This change will not have an impact on the result, but it is the way the code was intended to have been written.

B.2 Data Selection

The TRICH parameter could be modified so that it knows about the concentrator cut off. This would reduce the size of the variation of TRICH with impact parameter that can be seen in Figure 5.4. On a similar note, the SNOMAN routine mes_muon_timing should be changed so that only tubes with both good charges AND good times are considered.

It seems likely that at least some of the difference in the CRMS parameter must be due to the single photon transit time distributions being more broad in the simulation than they are in reality. It is possible that reducing the difference between data and simulated muons would allow for better rejection of internal neutrino interactions. If a future analyst creates a narrower single photon transit time distribution it is absolutely essential that the multiple photon timing correction be re-created to match the new transit time distribution.

On the subject of internal neutrino interactions, it is strongly recommended that a cut be designed to remove these events. One possibility for such a cut would be to fit for the "speed of light" that minimizes the spread of the time residual distribution given the fit entrance and exit positions. This fit speed of light should be very close to c for real muons, but would presumably be far larger for internal neutrino interactions.

There are two modifications that could be made to the dE/dx estimation. First, one could change dE/dx to use Equation 5.8 instead of Equation 5.9. Second, it has been observed that due to the light absorption of the neck there is a few % up-down asymmetry in dE/dx. Since dE/dx is used primarily as a cut and since the cut is not set too close to the start of the distribution this has a small effect on the acceptance. Still, it might be worth thinking about how this effect could be accounted for.

Finally, the angular dependence of the prepulse could be added to the simulation, but it is unlikely that this would have an impact on any analyses.

B.3 Flux Analysis

The largest detector uncertainty on the neutrino induced muon flux was due to the acceptance changing as a function of zenith angle. This is almost entirely due to a small number of events that pass close to the neck. Propagating the events generated by nuance through SNOMAN would allow this acceptance to be measured accurately and reduce this uncertainty below the 1% level.

Appendix C

Muon Vertical Depth Intensity Data

The following table contains the vertical depth intensity data discussed in section 7.3.

The table continues on the following page.

Min Depth (m w.e)	Mean Depth (m w.e.)	Events	Vertical Intensity
6200	6225	1367	$(4.36 \pm 0.12) \times 10^{-10}$
6250	6275	1239	$(3.98 \pm 0.11) \times 10^{-10}$
6300	6325	1121	$(3.63 \pm 0.11) \times 10^{-10}$
6350	6375	1151	$(3.76 \pm 0.11) \times 10^{-10}$
6400	6425	1003	$(3.30 \pm 0.10) \times 10^{-10}$
6450	6475	941	$(3.12 \pm 0.10) \times 10^{-10}$
6500	6525	876	$(2.93 \pm 0.10) \times 10^{-10}$
6550	6575	828	$(2.79 \pm 0.10) \times 10^{-10}$
6600	6625	803	$(2.73 \pm 0.10) \times 10^{-10}$
6650	6675	731	$(2.50 \pm 0.09) \times 10^{-10}$
6700	6725	684	$(2.36 \pm 0.09) \times 10^{-10}$
6750	6774	593	$(2.06 \pm 0.09) \times 10^{-10}$
6800	6825	575	$(2.01 \pm 0.08) \times 10^{-10}$
6850	6875	557	$(1.96 \pm 0.08) \times 10^{-10}$
6900	6924	477	$(1.69 \pm 0.08) \times 10^{-10}$
6950	6976	498	$(1.78 \pm 0.08) \times 10^{-10}$
7000	7025	452	$(1.63 \pm 0.08) \times 10^{-10}$
7050	7075	428	$(1.55 \pm 0.08) \times 10^{-10}$
7100	7125	397	$(1.45 \pm 0.07) \times 10^{-10}$

Min Depth (m w.e)	Mean Depth (m w.e.)	Events	Vertical Inensity
7150	7176	366	$(1.35 \pm 0.07) \times 10^{-10}$
7200	7225	320	$(1.19 \pm 0.07) \times 10^{-10}$
7250	7274	310	$(1.16 \pm 0.07) \times 10^{-10}$
7300	7325	315	$(1.18 \pm 0.07) \times 10^{-10}$
7350	7376	241	$(9.1 \pm 0.6) \times 10^{-11}$
7400	7424	274	$(10.4 \pm 0.6) \times 10^{-11}$
7450	7475	229	$(8.79 \pm 0.6) \times 10^{-11}$
7500	7523	209	$(8.08 \pm 0.6) \times 10^{-11}$
7550	7574	213	$(8.29 \pm 0.6) \times 10^{-11}$
7600	7699	788	$(7.79 \pm 0.3) \times 10^{-11}$
7800	7892	528	$(5.37 \pm 0.2) \times 10^{-11}$
8000	8097	435	$(4.53 \pm 0.2) \times 10^{-11}$
8200	8294	375	$(4.01 \pm 0.2) \times 10^{-11}$
8400	8494	303	$(3.32 \pm 0.19) \times 10^{-11}$
8600	8696	213	$(2.39 \pm 0.16) \times 10^{-11}$
8800	8894	190	$(2.18 \pm 0.16) \times 10^{-11}$
9000	9099	120	$(1.41 \pm 0.13) \times 10^{-11}$
9200	9300	116	$(1.39 \pm 0.13) \times 10^{-11}$
9400	9504	93	$(1.14 \pm 0.12) \times 10^{-11}$
9600	9682	63	$(7.9 \pm 1.0) \times 10^{-12}$
9800	9888	56	$(7.2 \pm 1.0) \times 10^{-12}$
10000	10220	105	$(5.6 \pm 0.5) \times 10^{-12}$
10500	10713	51	$(2.8 \pm 0.4) \times 10^{-12}$
11000	11254	32	$(1.9 \pm 0.3) \times 10^{-12}$
11500	11752	19	$(1.2 \pm 0.3) \times 10^{-12}$
12000	12239	12	$(7.6 \pm 2.2) \times 10^{-13}$
12500	12806	10	$(6.6 \pm 2.1) \times 10^{-13}$
13000	13299	6	$(4.1 \pm 1.7) \times 10^{-13}$
13500	13601	2	$(1.5 \pm 1.0) \times 10^{-13}$
14000	14305	4	$(3.0 \pm 1.5) \times 10^{-13}$
14500	14523	2	$(1.6 \pm 1.1) \times 10^{-13}$
15000	15350	2	$(1.6 \pm 1.1) \times 10^{-13}$

Table C.1: This table contains the measured muon vertical depth intensities measured at SNO. The measurement of these values was discussed in section 7.3. The listed uncertainties are statistical only, for a discussion of systematic uncertainties please see the text.

Appendix D

Neutrino Runs Used in the Throughgoing Analysis

D_2O phase runs:

10000 10002 10003 10005 10008 10015 10020 10023 10025 10030 10031 10034 10035 10036 10038 10040 10124 10125 10129 10130 10133 10141 10142 10149 10151 10161 10162 10163 10169 10170 10171 10172 10174 10177 10189 10190 10197 10219 10221 10224 10236 10237 10534 10536 10549 10551 10554 10555 10611 10638 10649 10650 10651 10655 10659 10675 10677 10678 10680 10686 10687 10700 10701 10704 10705 10706 10708 10709 10710 10714 10734 10735 10736 10737 10738 10739 10740 10741 10742 10743 10744 10747 10748 10749 10756 10762 10770 10773 10775 10776 10779 10781 10782 10783 10784 10797 10801 10803 10804 10805 10806 10811 10813 10815 10821 10826 10828 10836 10841 10843 10869 10871 10873 10876 10878 10879 10881 10882 10883 10884 10885 10886 10887 10891 10894 10922 10923 10924 10925 10927 10932 10933 10935 10936 10938 10939 10942 10943 10944 10946 10948 10949 10950 10951 10953 10954 10955 10956 10959 10961 10962 10963 10970 10972 10975 10976 11269 11271 11272 11281 11286 11289 11291 11303 11310 11312 11313 11347 11366

11368 11371 11377 11381 11383 11384 11389 11390 11393 11397 11399 11400 11402
11406 11407 11415 11417 11429 11431 11433 11436 11437 11443 11444 11446 11462
11466 11474 11479 11481 11489 11490 11493 11498 11502 11504 11506 11508 11510
11512 11525 11528 11530 11531 11532 11533 11537 11539 11541 11543 11544 11550
11553 11554 11557 11558 11561 11568 11570 11575 11579 11582 11591 11621 11650
11652 11655 11657 11670 11676 11679 11681 11682 11690 11698 11701 11703 11705
11706 11730 11732 11733 11764 11783 11802 11804 11805 11816 11819 11820 11824
11828 11829 11831 11859 11864 11867 11875 11890 11899 11901 11903 11911 11915
11919 11924 11925 11927 11928 11976 11977 11978 11981 11985 11988 11990 11991
12038 12054 12059 12082 12125 12131 12150 12157 12159 12162 12165 12167 12168
12172 12173 12178 12181 12183 12187 12190 12192 12197 12201 12203 12207 12222
12224 12226 12227 12229 12233 12234 12237 12238 12240 12243 12257 12289 12290
12329 12330 12506 12568 12571 12575 12576 12577 12582 12588 12590 12598 12614
12615 12618 13121 13292 13294 13302 13331 13334 13335 13338 13340 13341 13351
13389 13392 13396 13401 13405 13408 13415 13418 13421 13423 13426 13428 13431
13432 13434 13441 13444 13446 13451 13729 13733 13746 13774 13811 13874 13880
13882 13886 13895 14006 14008 14031 14033 14077 14078 14080 14083 14177 14185
14186 14190 14196 14252 14255 14264 14287 14289 14291 14292 14293 14301 14304
14308 14311 14315 14316 14377 14378 14386 14388 14389 14390 14393 14394 14398
14399 14402 14404 14405 14409 14410 14411 14413 14415 14417 14418 14422 14425
14429 14431 14438 14450 14451 14464 14466 14493 14495 14496 14652 14677 14680
14684 14685 14750 14757 14762 14764 14768 14770 14773 14775 14777 14781 14782
14787 14807 14814 14878 14883 14915 14958 14961 14962 14969 14970 15005 15011
15012 15014 15018 15020 15021 15022 15025 15027 15028 15029 15034 15051 15058
15065 15067 15078 15083 15105 15111 15112 15117 15119 15120 15129 15132 15147
15153 15165 15180 15214 15228 15268 15269 15270 15271 15272 15276 15279 15309

15340 15352 15370 15538 15563 15567 15595 15598 15600 15601 15604 15610 15611 15612 15615 15617 15618 15620 15624 15625 15640 15641 15643 15647 15651 15652 15653 15654 15655 15656 15657 15662 15669 15670 15671 15672 15673 15679 15684 15696 15698 15733 15724 15745 15746 15747 15748 15750 15752 15755 15762 15767 15768 15789 15791 15792 15794 15799 15802 15806 15808 15810 15819 15820 15821 15826 15828 15829 15830 15842 15843 15844 15862 15865 15869 15870 15871 15872 15873 15874 15877 15884 15905 15907 15941 15943 15947 15948 15949 15958 15978 15997 15998 16002 16003 16013

Salt phase runs:

20684 20691 20694 20697 20699 20700 20704 20705 20706 20739 20741 20742 20744 20747 20750 20751 20752 20754 20757 20759 20773 20776 20777 20779 20781 20785 20786 20787 20789 20796 20800 20807 20808 20809 20814 20841 20846 20852 20854 20865 20866 20870 20873 20874 20877 20900 20906 20911 20930 20934 20936 20937 20947 20955 20964 20965 20967 20968 20969 20978 20979 20980 20993 20994 20998 20999 21002 21005 21263 21501 21541 21562 21566 21594 21598 21604 21609 21610 21611 21612 21615 21616 21617 21620 21621 21622 21624 21628 21629 21643 21647 21651 21652 21654 21659 21660 21663 21667 21668 21672 21674 21681 21682 21693 21697 21705 21706 21707 21708 21711 21713 21715 21717 21727 21730 21736 21740 21750 21758 21777 21781 21784 21785 21786 21793 21794 21795 21796 21797 21798 21807 21808 21809 21810 21828 21835 21836 21838 21846 21854 21863 21864 21865 21870 21871 21873 21880 21884 21898 21901 21903 21911 21912 21913 21922 22001 22006 22009 22011 22013 22023 22027 22029 22030 22031 22063 22065 22066 22078 22086 22088 22090 22092 22126 22329 22331 22354 22358 22361 22366 22369 22375 22378 22380 22381 22382 22383 22397 22399 22400 22401 22402 22405 22406 22410 22414 22417 22418 22419 22420 22422 22423 22426 22430 22434 22435 22440 22441 22444

Bibliography

[1] C. L. Cowan, F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire, Science **124**, 103 (1956).

[2] R. Davis Jr, D. S. Harmer, and K. C. Hoffman, Phys. Rev. Lett. **20**, 1205 (1968).

[3] G. Danby *et al.*, Phys. Rev. Lett. **9**, 36 (1962).

[4] K. S. Hirata *et al.* (KAMIOKANDE-II), Phys. Lett. **B205**, 416 (1988).

[5] K. Hirata *et al.* (KAMIOKANDE-II), Phys. Rev. Lett. **58**, 1490 (1987).

[6] A. Mann, *Shadow of a Star* (W.H. Freeman and Company, New York, 1997).

[7] D. Decamp *et al.* (ALEPH), Z. Phys. **C53**, 1 (1992).

[8] Q. R. Ahmad *et al.* (SNO), Phys. Rev. Lett. **89**, 011301 (2002), `nucl-ex/0204008`.

[9] T. Araki *et al.* (KamLAND), Phys. Rev. Lett. **94**, 081801 (2005), `hep-ex/0406035`.

[10] Y. Ashie *et al.* (Super-Kamiokande), Phys. Rev. **D71**, 112005 (2005), `hep-ex/0501064`.

[11] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, Phys. Rev. **105**(4), 1413 (1957).

[12] S. Weinberg, Phys. Rev. Lett. **19**, 1264 (1967).

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

[14] C. Giunti and M. Laveder (2003), `hep-ph/0310238`.

[15] B. T. Cleveland *et al.*, Astrophys. J. **496**, 505 (1998).

[16] J. N. Bahcall, M. H. Pinsonneault, and S. Basu, Astrophys. J. **555**, 990 (2001), `astro-ph/0010346`.

[17] A. S. Brun, S. Turck-Chieze, and J. P. Zahn, Astrophys. J. **525**, 1032 (1999), `astro-ph/9906382`.

[18] O. L. G. Peres and A. Y. Smirnov, Phys. Lett. **B456**, 204 (1999), `hep-ph/9902312`.

[19] B. Aharmin *et al.* (SNO), Phys. Rev. C **72**(055502) (2005).

[20] T. K. Gaisser, *Cosmic rays and particle physics* (Cambridge University Press, Cambridge, UK, 1990).

[21] L. Anchordoqui, T. Paul, S. Reucroft, and J. Swain, Int. J. Mod. Phys. **A18**, 2229 (2003), `hep-ph/0206072`.

[22] R. P. Kokoulin, Nucl. Phys. Proc. Suppl. **70**, 475 (1999).

[23] D. Mei and A. Hime, Phys. Rev. **D73**, 053004 (2006), `astro-ph/0512125`.

[24] M. Ambrosio *et al.*, Phys. Rev. D **56**(3), 1418 (1997).

[25] T. K. Gaisser and M. Honda, Ann. Rev. Nucl. Part. Sci. **52**, 153 (2002), `hep-ph/0203272`.

[26] G. D. Barr, T. K. Gaisser, P. Lipari, S. Robbins, and T. Stanev, Physical Review D **70**, 023006 (2004).

[27] M. Ackermann and et al., in *International Cosmic Ray Conference* (2005), pp. 107–+.

[28] W. Seligman, *A Next-to-Leading-Order QCD Analysis of Neutrino-Iron Structure Functions at the Tevatron*, Ph.D. thesis, Columbia University (1997).

[29] Y. Ashie *et al.* (Super-Kamiokande), Phys. Rev. Lett. **93**, 101801 (2004), `hep-ex/0404034`.

[30] Y. Fukuda *et al.* (Super-Kamiokande), Phys. Lett. **B467**, 185 (1999), `hep-ex/9908049`.

[31] M. C. Gonzalez-Garcia, M. Maltoni, and J. Rojo (2006), `hep-ph/0608319`.

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

[33] M. S. Neubauer, *Evidence for Electron Neutrino Flavor Change through Measurement of the* $^{8}\mathrm{B}$ *Solar Neutrino Flux at the Sudbury Neutrino Observatory*, Ph.D. thesis, University of Pennsylvania, Penn (2001).

[34] P. Wittich, *First Measurement of the Flux of Solar Neutrinos from the Sun at the Sudbury Neutrino Observatory*, Ph.D. thesis, University of Pennsylvania (2000).

[35] S. Redmond, personal communication.

[36] M. L. Dunford, personal communication.

[37] C. Tunnel and S. Seibert, *Measuring the Trigger Efficiency in the NCD-Phase and the All-Phase Implications*, Tech. Rep., University of Texas at Austin (2006).

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

[39] B. A. Moffat, *The Optical Calibration of the Sudbury Neutrino Observatory*, Ph.D. thesis, Queen's University, Kingston, Canada (2001).

[40] J. R. N. Cameron, *The Photomultiplier Tube Calibration of the Sudbury Neutrino Observatory*, Ph.D. thesis, Wolfson College, Oxford University, Oxford (2001).

[41] J. Klein, *Realistically Simulated Charge Spectrum for SNO and the Prediction of Tube-by-tube Efficiencies*, Tech. Rep., University of Texas at Austin (2005).

[42] A. S. Hamer, *Energy Calibration of SNO for measurement of the Charged-Current Neutrino Reaction*, Ph.D. thesis, Queen's University, Kingston, Ontario (1999).

[43] S. Biller *et al.*, *SNO Electronic Calibration Constants*, Tech. Rep., SNO Collaboration (2001).

[44] M. L. Dunford, *Measurement of the Total* ^{8}B *Solar Neutrino Flux at the Sudbury Neutrino Observatory*, Ph.D. thesis, University of Pennsylvania (2006).

[45] M. G. Boulay, *Direct Evidence for Weak Flavor Mixing with the Sudbury Neutrino Observatory*, Ph.D. thesis, Queen's University (2001).

[46] The SNO Collaboration, *The SNOMAN User's Manual*, version 5.02 ed. (2006).

[47] S. J. Brice, *Monte Carlo and Analysis Techniques for the Sudbury Neutrino Observatory*, Ph.D. thesis, Balliol College, Oxford University, Oxford (1996).

[48] R. Van Berg, personal communication.

[49] C. Kyba, *An improved gain calibration for QLX*, Tech. Rep., SNO Collaboration (2004).

[50] C. Kyba, *An Improved SNOMAN PMT Transit Time Distribution with Realistic Prepulse*, Tech. Rep., SNO Collaboration (2003).

[51] E. Frank, *A Study of Atmospheric Neutrino Interactions in the Kamiokande-II Detectors*, Ph.D. thesis, University of Pennsylvania (1992).

[52] N. Tagg, *The 8Li Calibration Source and Through-going Muon Analysis in the Sudbury Neutrino Observatory*, Ph.D. thesis, University of Guelph (2001).

[53] A. Kibayashi, *Atmospheric Neutrino Oscillation Analysis Using FLUKA 3D Flux and Super-Kamiokande Datas*, Ph.D. thesis, University of Hawaii (2002).

[54] N. McCauley, *The Blindness Scheme for the Muon Analysis*, Tech. Rep., SNO Collaboration (2004).

[55] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in Fortran* (Cambridge University Press, 1992), 2nd ed.

[56] J. R. Wilson, *A Measurement of the* $^{8}\mathrm{B}$ *Solar Neutrino Energy Spectrum at the Sudbury Neutrino Observatory*, Ph.D. thesis, Jesus College, Oxford University, Oxford (2004).

[57] J. Klein, *A Short, Sharp Shock: Results from the* $\mathrm{D_2O}$ *Rn Spike*, Tech. Rep., University of Texas at Austin (2003).

[58] N. K. McCauley, *Producing a Background Free Data Set for Meausurment of the Charge Current Flux and Day-Night Asymmetry at the Sudbury Neutrino Observatory*, Ph.D. thesis, The Queen's College, Oxford University, Oxford (2001).

[59] T. Tsui, personal communication.

[60] W.-M. Yao, C. Amsler, D. Asner, R. Barnett, J. Beringer, P. Burchat, C. Carone, C. Caso, O. Dahl, G. D'Ambrosio, A. DeGouvea, M. Doser, *et al.*, Journal of Physics G **33**, 1+ (2006), URL `http://pdg.lbl.gov`.

[61] G. O. Gann, personal communication.

[62] J. R. Klein, personal communication.

[63] G. J. Feldman and R. D. Cousins, Physical Review D **57**, 3873 (1998).

[64] Hamamatsu Photonics, *Photomultiplier Tubes* (Hamamatsu Photonics K. K., 2006), 3rd ed.

[65] M. D. Lay, *Creation and Detection of Cerenkov light in the Sudbury Neutrino Observatory*, Ph.D. thesis, Wolfson College, Oxford University (1994).

[66] D. Casper, Nucl. Phys. Proc. Suppl. **112**, 161 (2002), `hep-ph/0208030`.

[67] A. Bodek and U. K. Yang, Nucl. Phys. Proc. Suppl. **112**, 70 (2002), `hep-ex/0203009`.

[68] R. A. Smith and E. J. Moniz, Nucl. Phys. **B43**, 605 (1972).

[69] D. Rein and L. M. Sehgal, Ann. Phys. **133**, 79 (1981).

[70] P. Lipari and T. Stanev, Phys. Rev. D **44**(11), 3543 (1991).

[71] D. E. Groom, N. V. Mokhov, and S. I. Striganov, Atom. Data Nucl. Data Tabl. **78**, 183 (2001).

[72] M. Honda, T. Kajita, K. Kasahara, and S. Midorikawa, Phys. Rev. **D70**, 043008 (2004), `astro-ph/0404457`.

[73] G. Battistoni, A. Ferrari, T. Montaruli, and P. R. Sala (2003), `hep-ph/0305208`.

[74] D. Casper, Personal Communication.

[75] G. Cowan, *Statistical Data Analysis* (Oxford University Press, 1998).

[76] J. A. Formaggio, personal communication.

[77] E. V. Bugaev, I. A. Sokalski, and S. I. Klimushin (2000), `hep-ph/0010323`.

[78] C. Waltham, personal communication.

[79] M. Aglietta *et al.* (LVD), Phys. Rev. **D58**, 092005 (1998), `hep-ex/9806001`.

[80] M. Ambrosio *et al.*, Phys. Rev. D **52**(7), 3793 (1995).

[81] K. Munakata *et al.* (Matsushiro), J. Phys. Soc. Jap. **60**, 2808 (1991).

[82] G. Guillian *et al.* (Super-Kamiokande) (2006), `astro-ph/0508468`.

[83] V. L. Rusu, *Measurement of the Solar* ^{8}B *Neutrino Spectrum in the Sudbury Neutrino Observatory*, Ph.D. thesis, University of Pennsylvania (2003).

[84] J. A. Dunmore, *The Separation of CC and NC Events in the Sudbury Neutrino Observatory*, Ph.D. thesis, Hertford College, Oxford University, Oxford (2004).