Producing a Background Free Data Set for Measurement of the Charge Current Flux and Day-Night Asymmetry at the Sudbury Neutrino Observatory

Neil K. McCauley
The Queen's College, Oxford

Thesis submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy
at the University of Oxford
Michaelmas Term, 2001

Producing a Background Free Data Set for Measurement of the Charge Current Flux and Day-Night Asymmetry at the Sudbury Neutrino Observatory

Neil K. McCauley
The Queen's College, Oxford

Thesis submitted for the degree of Doctor of Philosophy

Michaelmas Term, 2001

Abstract

The SNO detector is a 1 kilo-tonne heavy water Čerenkov detector designed to solve the solar neutrino problem. The detector is situated 2km underground in the INCO Ltd. Creighton mine near Sudbury, Ontario. The heavy water is observed by approximately 9500 photo-multiplier tubes (PMTs) to detect Čerenkov light generated by solar neutrino interactions. Using heavy water SNO can detect neutrinos in three different ways. In this thesis aspects of the charge current (a reaction sensitive to electron neutrinos only) and elastic scattering flux analysis are presented. Some models predict a difference in the detected neutrino rate between day and night. Measurement of this via the day-night asymmetry can help solve the solar neutrino problem.

One of the principal problems for the solar neutrino analysis are the instrumental backgrounds; events caused by processes other than Čerenkov light. This thesis contains the descriptions of the backgrounds, the data selection cuts that have been designed to remove them and the effect of these cuts on SNO data. To demonstrate the effectiveness of the cuts, the fraction of good events removed and the residual background content of the data set after application of the cuts are measured.

Once the cuts have been applied to the data the charge current flux and day-night asymmetry can be measured. In this thesis data taken between the 2nd of November 1999 and the 1st July 2000, providing a total of 64.6 days of data during the day and 93.1 days of night data, are used to measure the day-night asymmetry. The results of this analysis for the asymmetry in the detected solar neutrino rate from the charge current and elastic scattering reactions are presented.

Acknowledgements

There are many people who deserve thanks for their help in the completion of this thesis. Firstly I would like to thank my supervisor Dave Wark for his support, guidance and enthusiasm throughout my time at Oxford. His willingness to share his knowledge and views on any subject from the parameters of the burst cut to the fortunes of West Bromwich Albion and the England cricket team has made working with him a privilege and a pleasure.

I would also like to thank Nick West for introducing me to SNOMAN and good programming practises. Steve Biller, Nick Jelley and Kate Frame have always provided me with entertaining discussions on various aspects of SNO and the wider world. The Oxford graduate students have all made my time in Oxford and Sudbury much more enjoyable. James Cameron and Gordon McGregor have both been through it all with me and each has always been ready to provide ideas and tips. Latterly Jeanne Wilson deserves thanks for taking up the slack as I began to concentrate on writing. I also thank Jessica Dunmore for the provision of the spinning tops. I must also thank Beverly Roger for all her help in organising my frequent and rarely simple travel plans.

I must also thank the members of the data cleaning, livetime, day-night and data selection working groups, without whom the content of this thesis would be remarkably different. Thanks must go to Josh Klein, Mark Neubauer, Vadim Rusu, Richard Van de Water, Xin Chen and Nathaniel Tagg who have provided some of cuts discussed and analysed in this thesis. Miles Smith and Scott Oser also deserve thanks for their thought and input on the measurement of the day-night asymmetry. Thanks must also go to Rich Helmer and the entire Penn group for convincing me that the combination of beer and wings is a good one.

Last but by no means least thanks must go to my parents for their love and support not only in the creation of this thesis but in all aspects of my life.

Statement of Originality

All the work, results and conclusions presented in this thesis are the work of the author except where stated in the text and below.

Because SNO is a large project the analysis of data taken by the detector is conducted in a collaborative fashion. In particular the two principle analysis presented in this thesis: data cleaning and day-night studies have both been carried out as a member of the corresponding working groups. For the analysis of data cleaning this means that not all the cuts discussed in the text were designed by the author. Amongst the instrumental cuts the AMB, QCluster, FGC, QvT, FTS, q/NHIT and ITC were designed by other members of the data cleaning group, as were the high level cuts, reconstruction and muon identification code. Modifications to some of these cuts based upon the sacrifice studies of the author were required however. The determination of contamination via a bifurcated analysis was pioneered by Vadim Rusu and Mark Chen. The bifurcated analysis in this thesis shows the authors . The results of the orphan study by the DAQ group, have been used to complete the discussion of dead time.

For the day-night analysis the technique used to determine the uncertainty due to the energy scale drift was developed by Scott Oser for the working group. Signal extraction is an essential part of the SNO analysis and is required to determine the separate asymmetries seen in the charge current and elastic scattering signals. In this thesis signal extraction code developed at Oxford by Gordon McGregor has been used to determine these asymmetries.

Contents

Chapter 1

The Solar Neutrino Problem

In 1930 Wolfgang Pauli postulated the existence of an extra particle to explain the apparent problem with energy conservation in nuclear β-decay. This particle came to be known as the neutrino. Over a quarter of a century was to pass before the first observation of the neutrino was made by Reines and Cowan *et al*[1]. They observed anti-neutrinos from a nuclear reactor by the inverse β-decay of the proton. It was not until the 1970s when neutrinos were first observed from the sun in the pioneering chlorine experiment of Davis *et al*[2]. This experiment was the first observation of extra-terrestrial neutrinos, and with the observation of a smaller flux than expected, the experiment was the founding of the thirty year old solar neutrino problem. It is only now that with the result of a variety of experiments that the final solution to this long standing problem is in sight.

1.1 The neutrino

The neutrino is the electro-weak partner of the charged lepton. However unlike all the other fundamental fermions the neutrino carries no charge, either electrical or coloured. Neutrinos therefore only interact with the $W^{\pm}$ boson to become their charged lepton partner, and with the Z^0 boson. As with all fundamental fermions there are three known flavours of neutrino: the electron neutrino (ν_e), the muon neutrino (ν_μ) and the tau neutrino (ν_τ). Each of these particles has been observed through its interactions with other matter. Observations on the width of the Z^0 boson[3] demonstrate that there are only three generations of active neutrinos with a mass less than half that of the Z^0 boson.

In the standard model of particle physics the neutrino is a massless particle, although there is now strong evidence (from neutrino oscillations discussed below) that this is not the case. Direct searches for neutrino mass based upon the kinematics of decays producing neutrinos have thus far been fruitless, so upper limits on the neutrino masses have been found instead. The latest results are: $m_{\nu_e} < 2.2\text{eV}$ [4][5], $m_{\nu_\mu} < 190\text{keV}$ [6]

and $m_{\nu_\tau} < 15.5$MeV[7].

For massless neutrinos each particle is in a fixed helicity state. Neutrinos are always left-handed, while anti-neutrinos are right-handed. However for massive neutrinos this raises an ambiguity. Neutrinos have no charge, so after a change of reference frame, the neutrino can become a neutrino of opposite helicity, or since it is now right-handed, it could become an anti-neutrino. In other words the neutrino could be its own anti-particle: a Majorana neutrino, rather than a Dirac neutrino. If the neutrino is a Majorana particle more sensitive measurements on the neutrino mass are possible by investigation of neutrinoless double β-decay. The most competitive limits come from experiments based upon ^{76}Ge giving: $m_{\nu_e} < 0.4$eV [8].

1.2 Neutrino oscillations

In the quark sector the eigenstates of the weak interaction are not the same as the eigenstates of the strong interaction; they are rotated with respect to one another. The mixing angles have been and continue to be measured by a variety of experiments. Leptons do not feel the strong interaction, but in a similar fashion there is no reason why the weak eigenstates should be the same as the mass eigenstates of the leptons, provided that neutrinos have mass. If this is the case then a neutrino created as an electron neutrino will have a finite probability of being detected as a different flavour of neutrino, after propagating some distance before detection. This phenomena, first suggested by Pontecorvo and Gribov[9][10], is known as neutrino oscillation.

1.2.1 Vacuum oscillations

The simplest form of neutrino oscillations are vacuum oscillations between two neutrinos, for example oscillations between ν_e and ν_μ. The mixing between the mass eigenstates and the weak eigenstates is controlled by a single parameter known as the vacuum mixing angle, θ_V, analogous to the Cabbibo angle in the quark sector. The eigenstates are related by a unitary transform:

$$\begin{pmatrix} |\nu_e> \\ |\nu_\mu> \end{pmatrix} = \begin{pmatrix} \cos(\theta_V) & \sin(\theta_V) \\ -\sin(\theta_V) & \cos(\theta_V) \end{pmatrix} \begin{pmatrix} |\nu_1> \\ |\nu_2> \end{pmatrix} \tag{1.1}$$

where ν_1 and ν_2 are the mass eigenstates, with masses m_1 and m_2. The analogy with the quark sector continues in the three neutrino case where the CKM matrix is replaced by the MNS matrix.

Continuing with the two neutrino oscillation case, the survival probability for the electron can be calculated:

$$p(\nu_e \to \nu_e) = 1 - \sin^2(2\theta_V)\sin^2(\frac{\pi R}{L_V}) \tag{1.2}$$

where R is the distance travelled by the neutrino in time t, and L_V is the vacuum oscillation length:

$$L_V = \frac{4\pi E\hbar}{\Delta m^2 c^3} \tag{1.3}$$

where E is the neutrino energy and:

$$\Delta m^2 = |m_1^2 - m_2^2|. \tag{1.4}$$

Using these equations it can be seen that neutrino oscillations, in the two neutrino case, are controlled by the two physical parameters $\sin^2(2\theta_V)$ and Δm^2. The objective in neutrino oscillation experiments is to measure these values. More generally we wish to measure the elements of the MNS matrix and the two Δm^2 values that relate the masses of the three neutrinos. Neutrino oscillations depend upon the value of $\frac{R}{E}$. The value of this ratio for an individual experiment limits the sensitivity of that experiment. Solar neutrino experiments are particularly sensitive to small values of Δm^2 due to the large distances involved.

1.2.2 Neutrino oscillations in matter

The above discussion applies only to neutrino oscillations in vacuum. When neutrinos propagate through matter additional complications arise that impact upon neutrino oscillations. These effects were first derived by Wolfenstein [11] and were applied to neutrino propagation through the varying density of the Sun by Mikheyev and Smirnov [12][13]. The large amount of mixing that this can give has since become known as the MSW effect.

Electron neutrinos can interact with matter by both charged and neutral current interactions, whereas for both μ and τ neutrinos only the neutral current channels are available. As neutrinos propagate through matter these forward-scattering processes contribute a mass terms to the mass matrix of the neutrinos. The mass term due to neutral current interactions is the same for all types of neutrino, and so has no effect on oscillations, while the charged current term only applies to the electron neutrino and this can have a significant effect on oscillation.

Following Bachall [14] the contribution to the mass matrix from this term is:

$$M_{matter} = \sqrt{2}G_F n_e P_e \tag{1.5}$$

where G_F is the Fermi coupling constant, n_e is the electron density and P_e is the projection operator for the electron neutrino. In this modified situation the form for neutrino oscillations is the same as the vacuum case, with a different mixing angle and oscillation length:

$$\tan(2\theta_M) = \frac{\tan(2\theta_V)}{[1 \pm L_V/L_e \sec(2\theta_V)]} \tag{1.6}$$

$$L_M = \frac{L_V}{\sqrt{(\sin^2(2\theta_V) + [L_V/L_e - \cos(2\theta_V)]^2)}} \tag{1.7}$$

where L_e is the electron-neutrino scattering length:

$$L_e = \frac{\sqrt{2}\pi\hbar c}{G_F n_e}. \tag{1.8}$$

The plus sign applies in equation 1.6 if $m_2 < m_1$ and the minus sign otherwise.

One particularly interesting feature of matter oscillations is that resonant oscillations can occur. Using equation 1.6 it can be seen that maximal mixing ($\theta_M = \frac{\pi}{4}$) is possible in matter even with tiny vacuum mixing angles, provided that $m_2 > m_1$. This resonance occurs when:

$$n_e = \frac{\Delta m^2 \cos(2\theta_V)}{2\sqrt{2}G_F E}. \tag{1.9}$$

This is of particular interest for solar neutrino studies due to the high electron density in the Sun, and the way that this can vary. Neutrinos are created in the centre of the Sun where the the electron density is highest. In this case it is possible that if $E >> \Delta m^2 \cos(2\theta_V)/G_F n_e$ then $\theta_M \approx \frac{\pi}{2}$ and the electron neutrino corresponds almost completely to the higher mass state (ν_2). As the neutrino propagates outwards from the centre of the Sun it passes through the resonance region where mixing is maximal. If the electron density is slowly varying (with respect to the oscillation length) the neutrino will remain in the ν_2 state as it passes out of the Sun. In particular if the vacuum mixing angle is small then the neutrinos leave the Sun in a state containing almost no electron neutrino. In other words the electron neutrinos are almost completely converted into another flavour by transit from the centre of the Sun. If the adiabatic condition does not hold then there is a certain probability for ν_2 to become ν_1 and neutrinos exit the Sun with a higher content of ν_e than in the adiabatic small mixing angle case. Because the value of the electron density at resonance is dependant on energy, the survival probability for electron neutrinos can be a strong function of energy. The survival probability for various regions of parameter space (see section 1.5.1) are shown in figure 1.1.

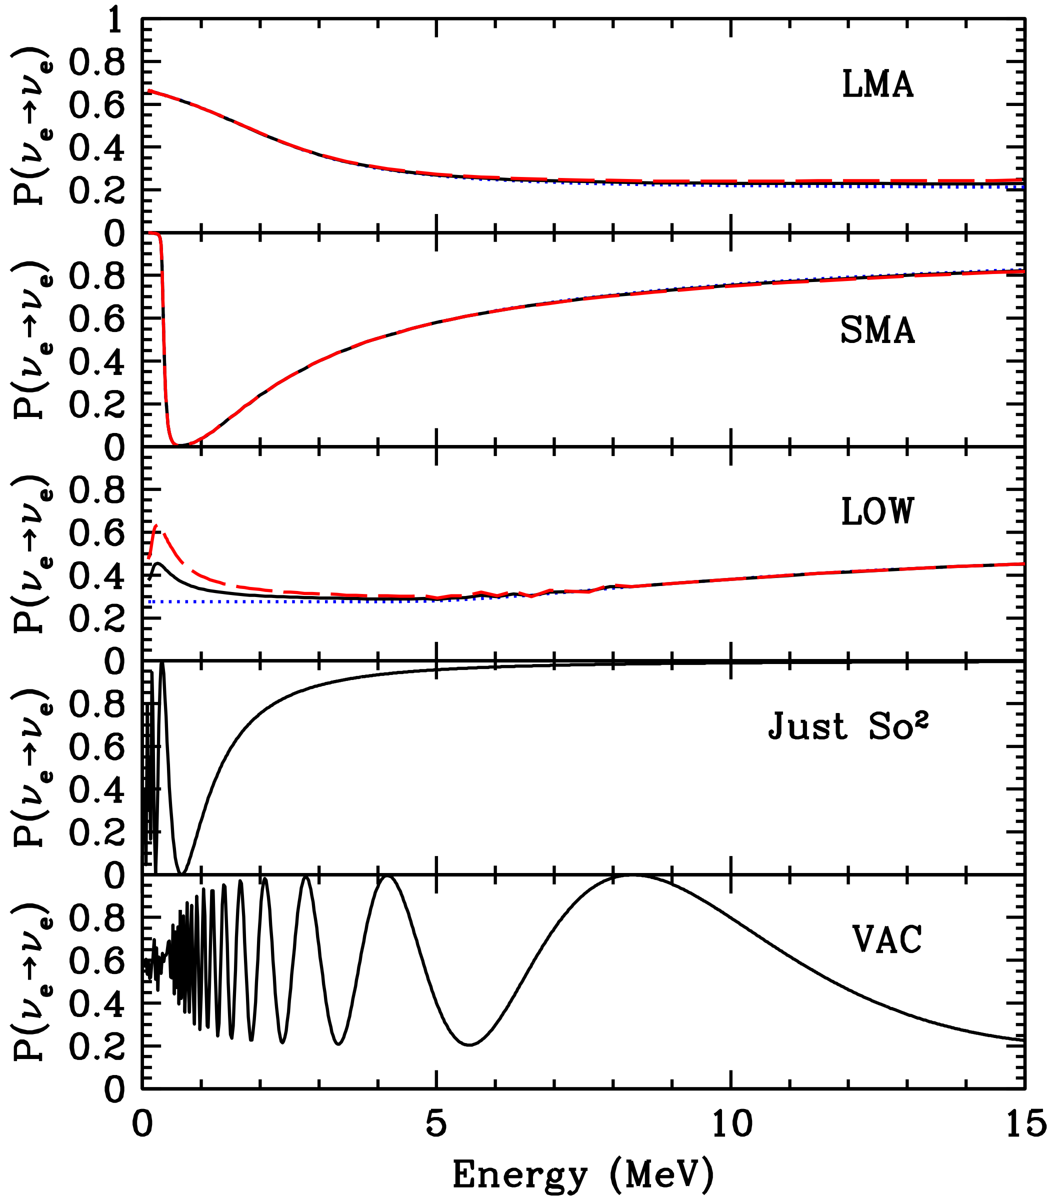

Figure 1.1: The survival probability as a function of energy for various regions of MSW parameter space (see section 1.5.1). The solid line represents the yearly average including Earth regeneration effects. The dotted line is for data taken during the day only, and the dashed line for data at night. Taken from [15].

Another feature of MSW oscillations of particular relevance to solar neutrino experiments is that the matter of the Earth can cause further oscillation of neutrinos. This means that the makeup of the neutrinos reaching a solar neutrino detector can be different during the night than during the day. Depending upon the detection method of the individual detector, this can appear as a day night asymmetry in the detected neutrino rate.

1.2.3 Sterile neutrinos

Neutrino oscillation experiments also give us the potential to infer the existence of a particle that could not otherwise be observed. These are sterile neutrinos; neutrinos that do not interact with either the $W^{\pm}$ or Z^0 bosons. They can however mix with the three known active neutrinos and further complicate possible oscillation scenarios. The lack of interactions of the sterile neutrino means that oscillations in matter also become more complicate. The neutral current interactions of active neutrinos effect neutrino oscillations involving sterile neutrinos. Sterile neutrinos should be kept in mind in the interpretation of results from neutrino oscillation experiments.

1.3 Standard solar models

The Sun is the nearest star to the Earth, and for this reason it is the star that we are able to measure with the highest precision[14]. For example the solar luminosity, mass and radius, as well as the chemical abundances at the Sun's surface are known to considerably more accuracy than for other stars. Therefore before we can hope to understand the myriad exotic stars seen in the night sky, we would hope to be able to understand our own star.

The Sun is a typical main sequence star steadily burning hydrogen into helium:

$$4\mathrm{p} \rightarrow^4 \mathrm{He} + 2\nu_e + 26.7\mathrm{MeV} \tag{1.10}$$

The energy released in this process is in the form of γ-rays and positrons, as well as the kinetic energy imparted to the outgoing particles. This process mostly occurs via the pp chain (figure 1.2), although a small amount of energy is released by the CNO cycle. The pp chain releases neutrinos from 5 different reactions as shown in figure 1.2. The neutrino spectrum from each of these reactions is set by standard nuclear physics, the only input of the solar model on the solar neutrino spectrum is therefore the relative number of neutrinos produced by each reaction. The two neutrinos produced in the creation of each ^{4}He carry away 0.6MeV on average, although higher energy neutrinos are produced; particularly the ^{8}B and hep neutrinos.

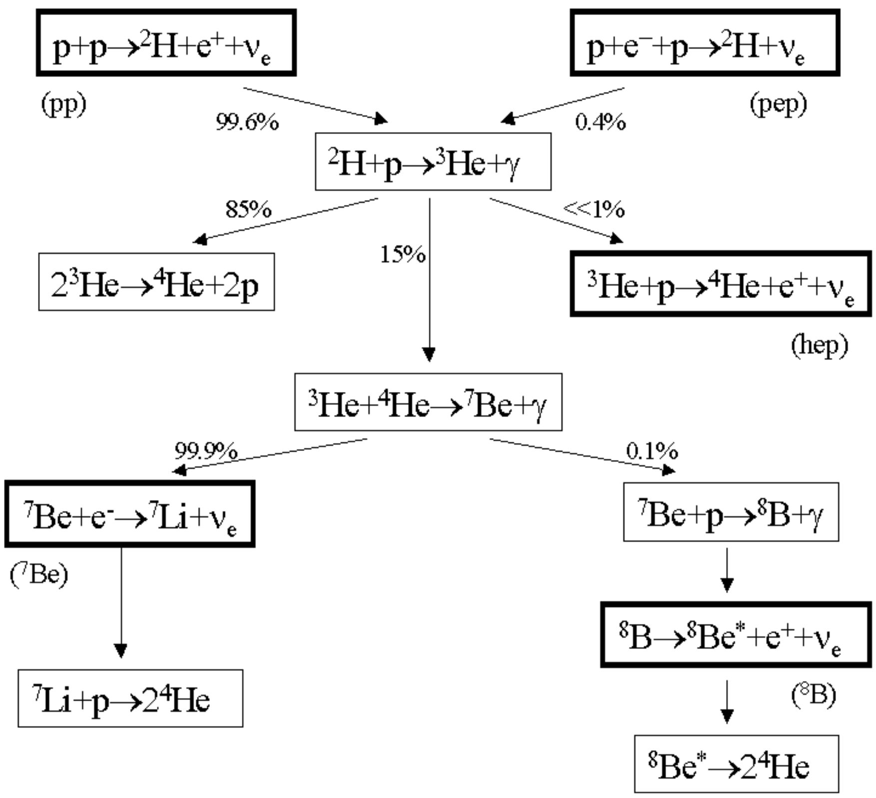

Figure 1.2: The nuclear reactions of the pp chain.

The rate-determining step of the pp cycle is the initial combination of two protons to create a deuteron. This reaction is highly dependant upon the temperature of the solar core, as the protons must have enough thermal energy to overcome their mutual electrostatic repulsion. Standard solar models [14], of which the latest (BP2001) is presented in [16] attempt to predict the neutrino flux from the various reactions based upon known physics and measurements of the Sun.

The Sun is modelled as a spherical body in hydrostatic equilibrium. Radiative and hydrostatic pressures exactly balance the force of gravity. This is a very good assumption because the timescale for the gravitational collapse of the Sun is of order one hour. Effects such as rotation and magnetic fields are neglected from the calculation. Solar models work by solving a coupled set of differential equations. These equations are evolved from the time the Sun joined the main sequence until the present day (and beyond). The initial abundances of the light elements and various other input parameters are then adjusted to reproduce the observed properties (such as the luminosity) of the Sun to provide a model of this particular star.

Energy is created in the core of the Sun and is lost from the surface. There is

Source	Flux $/10^{10}\mathrm{cm}^{-2}\mathrm{s}^{-1}$
pp	$5.95(1.00^{+0.01}_{-0.01})$
pep	$1.40 \times 10^{-2}(1.00^{+0.015}_{-0.015})$
hep	9.3×10^{-7}
^{7}Be	$4.77 \times 10^{-1}(1.00^{+0.10}_{-0.10})$
^{8}B	$5.05 \times 10^{-4}(1.00^{+0.20}_{-0.16})$
^{13}N	$5.48 \times 10^{-2}(1.00^{+0.21}_{-0.17})$
^{15}O	$4.80 \times 10^{-2}(1.00^{+0.25}_{-0.19})$
^{17}F	$5.63 \times 10^{-4}(1.00^{+0.25}_{-0.25})$

Table 1.1: Predicted (BP2001) neutrino fluxes at one astronomical unit from various reactions in the Sun. Taken from [16].

therefore a temperature gradient in the Sun and energy flows from the core to surface. In the core it is assumed that the flow of radiation dominates the energy flow. The solar model therefore depends upon the radiative opacity. This is set by the relative elemental abundances, which are variable as a function of time, as for example hydrogen is turned into helium. Under the assumption that the Sun was chemically homogenous at formation the initial chemical abundances of the heavier elements (such as carbon) can be input into the model. These are assumed to be equal to the present-day surface abundances. In the outer regions of the Sun energy transport is dominated by convection, this region of the Sun is therefore known as the convective zone. The time scale for a photon to randomly walk from the centre of the Sun is approximately 1000 years, while we observe neutrinos approximately 8 minutes after their creation.

The temperature profile of the stellar interior is then predicted by the model leading to predictions of the flux of neutrinos from each reaction in the pp chain and CNO cycle. For the latest standard solar model the neutrino fluxes are shown in table 1.1 and the expected neutrino energy spectrum along with experimental thresholds is shown in figure 1.3.

One of the undoubted successes of the standard solar model comes in the predictions for helioseismology. The Sun supports particular frequencies of sound waves (p-modes) between its surface and the base of the convective zone. This region defines a spherical shell in which the p-modes propagate. As with all systems containing waves in a box only certain waves are supported, although the Sun is much more complex than familiar everyday systems, and a huge variety of modes exist. By studying the effect these have on the surface of the Sun the interior of the Sun can also be studied. The standard solar model provides predictions for the speed of sound in the interior of the Sun, which in turn can be used to predict the modes of oscillation. The predicted sound speeds compared to

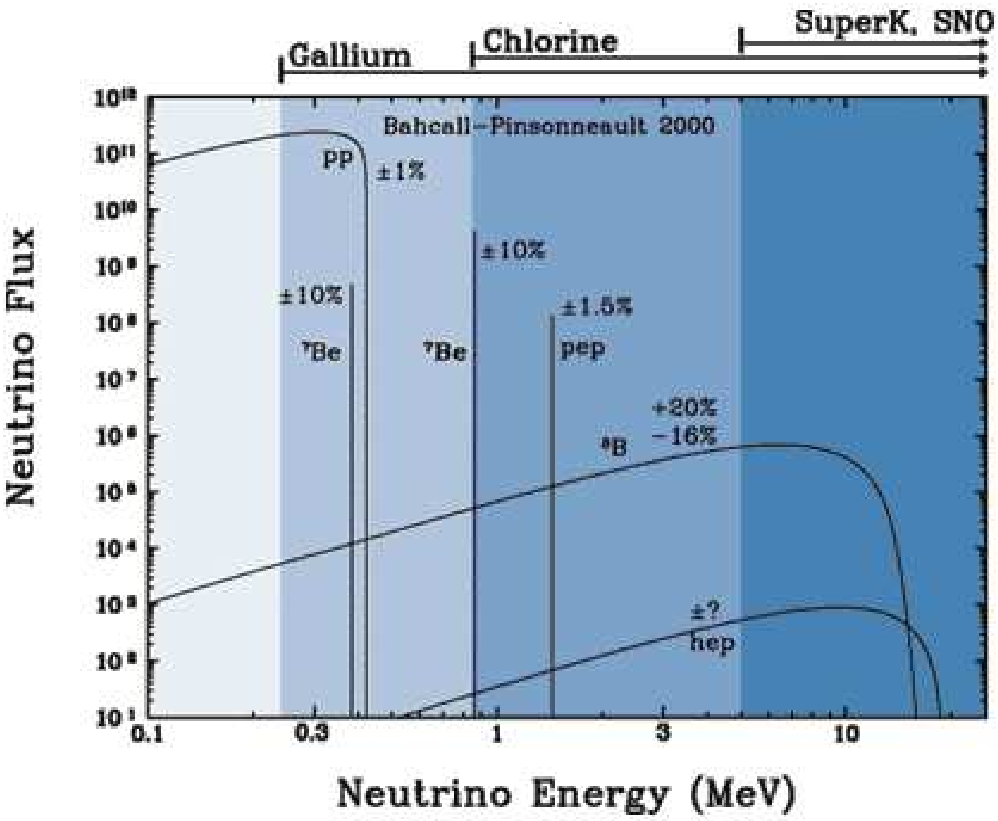

Figure 1.3: The energy spectrum of solar neutrinos predicted by the BP2001 standard solar model. Fluxes from continuous sources are in units of number per cm^2 per second per MeV at one astronomical unit. Line fluxes are in units of number per cm^2 per second. Taken from [17].

those measured using helioseismology are shown in figure 1.4. The predicted sound speeds depend upon the temperature of the Sun, as does neutrino production. Helioseismology limits the possible change in the temperature profile of the Sun and consequently provides constraints on the allowed values of the neutrino fluxes.

1.4 Solar neutrino experiments

Prior to the SNO experiment there have been several experiments carried out to measure the flux of neutrinos from the Sun. These experiments fall into two categories: radio-chemical and water Čerenkov experiments. Radio-chemical experiments measure the solar neutrino flux looking for the appearance of a new isotope in a chemical mixture caused by the interaction of electron neutrinos on a target nucleus. These experiments measure the total electron neutrino flux integrated above the energy threshold of the cho-

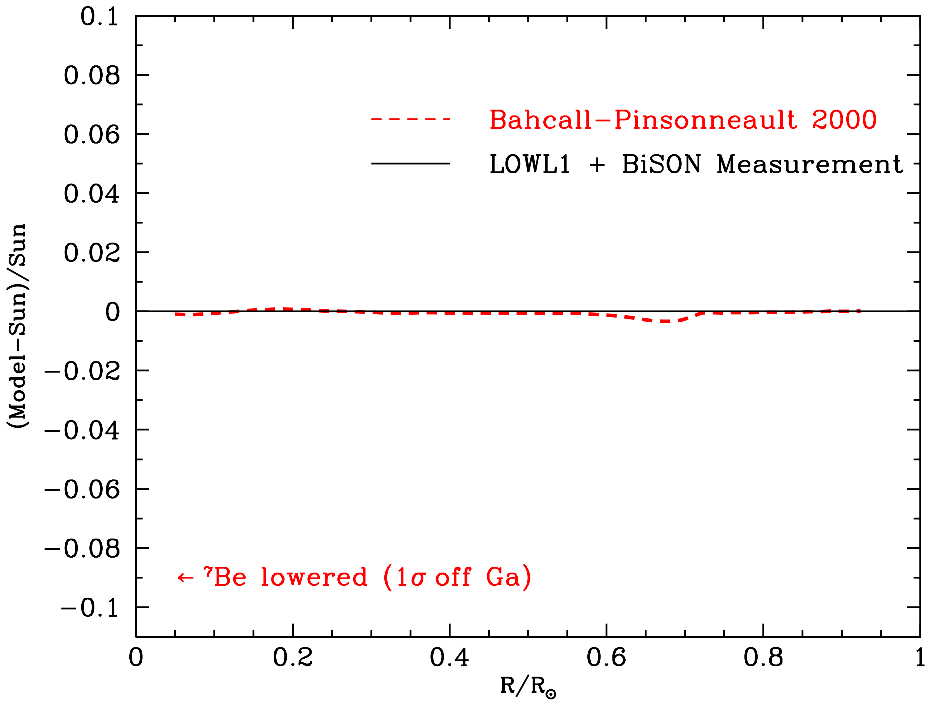

Figure 1.4: A comparison of helioseismology measurements of the sound speed as a function of radius to standard solar model predictions. Also shown is the scale of the disagreement expected if the model is altered to reduce the ^{7}Be flux such that the measurement of the gallium experiments agree with the solar model. Taken from [16].

sen reaction. Water Čerenkov detectors work in a different way. Neutrino interactions create electrons above the Čerenkov threshold, the light from which can then be detected by photo-multiplier tubes (PMTs). These experiments are real time and provided spectral information on the incoming neutrino flux, at the expense of a higher energy threshold compared to the radio-chemical experiments.

1.4.1 The chlorine experiment

The chlorine experiment [2], situated in the Homestake gold mine in South Dakota, is the original solar neutrino experiment. It was developed in the 1960's and the first solar neutrino observations were made using runs taken in 1970. The chlorine experiment is a radio-chemical experiment based upon the reaction:

$$\nu_e +^{37}\mathrm{Cl} \rightarrow e^- +^{37}\mathrm{Ar}. \tag{1.11}$$

The reaction has a neutrino energy threshold of 0.814MeV, so it is sensitive to all types of solar neutrinos except those from the pp reaction. The experiment is most sensitive to

^{8}B neutrinos due to an excited state of ^{37}Ar at 4.98MeV.

The detector consists of a tank of 615 metric tonnes of C_2Cl_4 surrounded by a water shield to protect the chlorine from neutrons from the rock wall. This is important as neutrons that reach the detector can cause the creation of ^{37}Ar. The detector is situated 1478m underground (equivalent to 4200m of water) to shield the detector from cosmic ray muons, which can also cause background by creation of ^{37}Ar.

To begin a measurement cycle a known amount of ^{36}Ar and ^{38}Ar are stirred into the tank. The tank is then left for one to three months for exposure to solar neutrinos. The argon in the tank including the ^{37}Ar created by solar neutrino interactions is then retrieved by bubbling helium through the tank from which the argon is then separated and placed in a proportional counter. The volume of argon obtained gives a measure of the extraction efficiency. ^{37}Ar decays by electron capture with a 35 day half-life. The number of ^{37}Ar atoms is counted by detection of Auger electrons emitted following the decay in the proportional counter.

Between 1970 and 1994 a total of 108 extractions were carried out and analysed. The result of these extractions gives a capture rate of:

$$2.56 \pm 0.16(stat.) \pm 0.16(sys.)\mathrm{SNU}. \tag{1.12}$$

The SNU is a convenient unit for radio-chemical solar neutrino experiments and corresponds to 10^{-36} captures per target atom per second. The expected neutrino capture rate is[16]:

$$7.6^{+1.3}_{-1.1}\mathrm{SNU}. \tag{1.13}$$

1.4.2 The water Čerenkov detectors

Water Čerenkov detectors were first built as proton decay experiments. However in 1985 the Kamiokande detector was upgraded with improved electronics and water purification systems to become Kamiokande II a detector which was able to detect solar neutrinos.

Water Čerenkov detectors detect neutrinos via the elastic scattering interaction:

$$\nu_x + e^- \to \nu_x + e^- \tag{1.14}$$

described in more detail in section 2.2.3. Naturally occurring radio-activity in the water, the PMTs, and the rock wall masks the neutrino signal at lower energies setting the energy threshold for the detector. A cleaner detector will in general run at a lower energy threshold. The typical energy threshold of water Čerenkov detectors is around 5MeV limiting them to the study of ^{8}B and hep neutrinos only.

Kamiokande

The Kamiokande II detector is located in the Kamioka mine, which is 1000 meters (2700 meters of water equivalent) underground. The detector consists of 3000 tons of water inside a cylindrical tank 15.6 m in diameter and 16.1 m high. The inner detector consists of 2140 tonnes and is surrounded by 948 20-inch PMTs. These provide a coverage[1] of around 20%. The outer volume is examined by 123 PMTs to veto cosmic rays. The outer volume also shields the detector from γ-rays and neutrons from the rock wall. In 1990 the detector was further upgraded to replace dead PMTs, add light reflectors (to increase the coverage to about 25%), install new electronics and reduce radon levels. This configuration is known as Kamiokande III. Reconstruction of the direction of events reveals a strong peak away from the Sun, providing direct evidence for the origin of the neutrinos. The combined result from Kamiokande II and III [18] made with a 7MeV threshold and a fiducial volume containing 680 tonnes is:

$$2.82^{+0.25}_{-0.24}(stat.) \pm 0.27(sys.) \times 10^6 cm^{-2} s^{-1}. \tag{1.15}$$

The ^{8}B neutrino flux from the standard solar model is[16]:

$$5.05^{+1.01}_{-0.81} \times 10^6 cm^{-2} s^{-1}. \tag{1.16}$$

Super-Kamiokande

In 1996 data taking at Super-Kamiokande (Super-K) began. Super-K is a much larger version of the Kamiokande experiment. Like Kamiokande the detector is cylindrical (39.3m in diameter and 41.4m high) consisting of 50 ktons of water, of which 32 ktons form the inner volume. The detector contains 11146 50cm PMTs looking inwards at this volume with a coverage of 40%. The outer detector, which is 2.6m thick, is again used to veto cosmic rays and to shield the detector from the rock wall. The region is viewed by 1885 20cm PMTs to act as veto tubes.

The analysis of data from Super-K proceeds in a similar fashion to that of Kamiokande data. Using 1258 days of data taken between May 1996 and October 2000 with an energy threshold of 5 MeV and a fiducial volume corresponding to 22.5 ktons Super-K make a measurement of the ^{8}B flux to extremely high statistical precision[19]:

$$2.32 \pm 0.03(stat.)^{+0.08}_{-0.07}(sys.) \times 10^6 cm^{-2} s^{-1}. \tag{1.17}$$

Super-K also measures the energy spectrum of electrons detected from solar neutrino interactions. This is found to be consistent with the expectations of an undistorted ^{8}B energy spectrum, with a χ^2 per degree of freedom of 19.0/18.

[1]The coverage is the fractional active area for photon detection.

Super-K makes neutrino measurements in real time, so the Super-K analysis has also been carried out in day and night bins to search for a day-night asymmetry. The day-night asymmetry measured by Super-K is[19]:

$$A_{ND} = \frac{\phi_N - \phi_D}{\frac{\phi_N + \phi_D}{2}} = 0.033 \pm 0.022(stat.)^{+0.013}_{-0.012}(sys.) \tag{1.18}$$

consistent with no asymmetry at the 1.3σ level.

1.4.3 The gallium experiments

While the water Čerenkov detectors and the chlorine experiment are sensitive to solar neutrinos, neither can see neutrinos from the pp interaction. The flux of pp neutrinos is very strongly constrained by the solar luminosity and is much more independent of the details of the solar model than the flux from other reactions. Experiments that measure the pp neutrinos require an energy threshold below 420 keV, the pp end-point. This can be achieved by the use of the gallium via the reaction

$$\nu_e +^{71} \mathrm{Ga} \rightarrow e^- +^{71} \mathrm{Ge}. \tag{1.19}$$

The energy threshold for this reaction is 233 keV.

To date three experiments have been built that use this reaction to detect solar neutrinos. SAGE (the Soviet American Gallium Experiment) [20] uses 57 tonnes[2] of liquid metallic Gallium as the target material. Liquid gallium is used to reduce problems with radio-active impurities, and it requires that the gallium is maintained at a temperature above the melting point of 29.8^oC. The experiment is located under Mount Andyrchi in the Baksan Valley. The rock overburden provides shielding from comic rays equivalent to 4700m of water.

Situated in the Gran Sasso laboratory, GALLEX (GALLium EXperiment) [21] and its successor GNO (Gallium Neutrino Observatory) [22][23] instead use a 100 tonne tank of gallium chloride solution containing 30 tonnes of gallium. The Gran Sasso laboratory provides shielding from comic rays equivalent to 3800m of water. This reduces the cosmic ray background in these experiments to 3.7 ± 1.1SNU, compared to a negligible contribution in SAGE.

As with the chlorine experiment the gallium experiments must extract the isotope produced by solar neutrino interactions before counting can occur. The gallium experiments use a chemical extraction technique to remove the germanium from the detector in the form of GeH_4. This is then placed inside proportional counters to count the number of atoms produced. The half life of ^{71}Ge is 11.4 days and this sets the time scale for each

[2]The experiment was increased in mass from 30 tonnes in mid-1991.

extraction. Each experiment runs for 2-4 weeks between extractions. ^{71}Ge decays by electron capture and decays are counted by detection of Auger electrons. In order to reduce systematics SAGE and GALLEX have also made use of calibration sources. In particular both have deployed extremely strong ^{51}Cr neutrino sources to calibrate the detectors.

The results from the gallium experiments are all consistent with each other. From 70 runs taken between January 1990 and October 1999 SAGE measures a capture rate of[24]:

$$75.4^{+7.0}_{-6.8}(stat.)^{+3.5}_{-3.0}(sys.)\text{SNU} \tag{1.20}$$

GALLEX using 65 runs taken between May 1991 and January 1997 measures[21]:

$$77.5 \pm 6.2(stat.)^{+4.3}_{-4.7}(sys.)\text{SNU} \tag{1.21}$$

and GNO using 19 runs between May 1998 and January 2000 measures[22]:

$$65.8^{+10.2}_{-9.6}(stat.)^{+3.4}_{-3.6}(sys.)\text{SNU}. \tag{1.22}$$

These measurements compare to the standard model expectation of[16]:

$$128^{+9}_{-7}\text{SNU}. \tag{1.23}$$

1.5 The solar neutrino problem

The results of all these experiments are significantly less than the predictions of the standard solar model, each experimental result lies between 30-60% of the prediction. This deficit is known as the solar neutrino problem and is best summarised in figure 1.5. There are three possible explanations for the solar neutrino problem:

- The experiments are wrong.

  The experimental measurements could be be incorrect. This seems unlikely however as the measurements have been made using six different experiments with three completely independent techniques. Each technique has very different analysis procedures and systematics. It therefore seems highly unlikely that the all experimental results are incorrect.

- The standard solar model is wrong.

  This solution to the solar neutrino problem also seems to be unlikely to be the correct one. The standard solar model relies upon well understood physics and uses cross sections that have been measured in terrestrial experiments. The standard

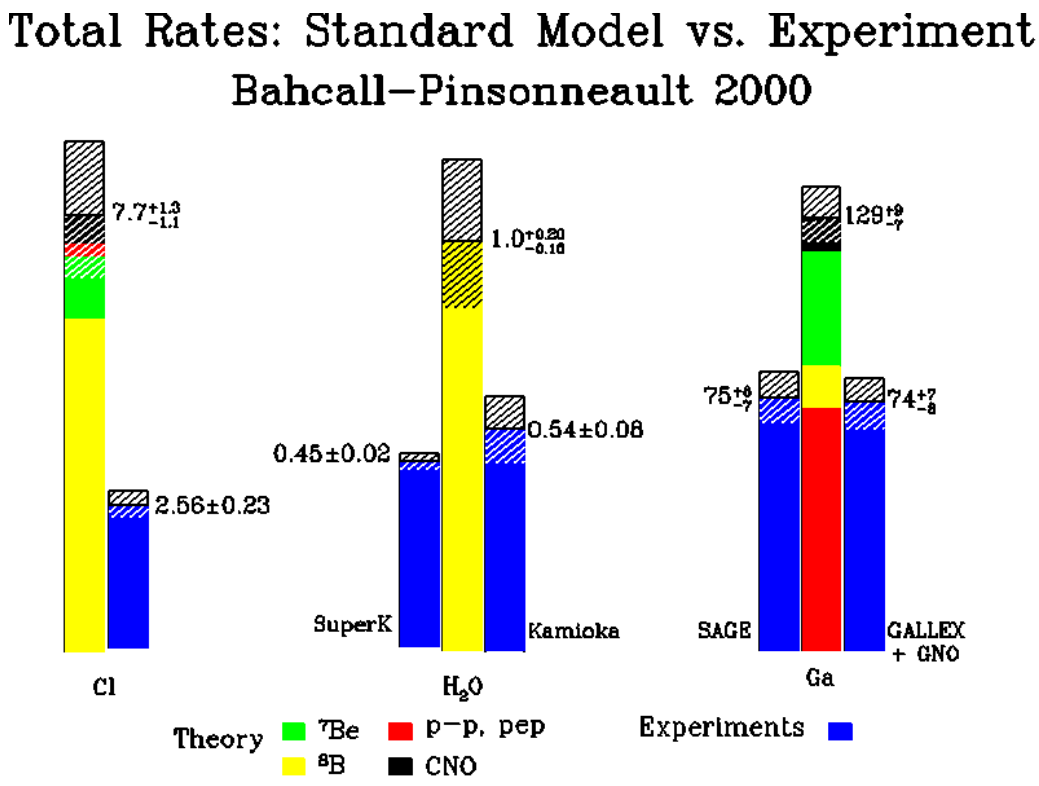

Figure 1.5: The rates of the various solar neutrino experiments compared to the theoretical predictions. The predictions are split into the separate predictions from the various contributory reactions to each experiment. The hatched regions represent the 1σ errors on the experiments and the solar model. Taken from [17].

solar model is also highly successful in describing helioseismology data. It turns out to be very difficult (if not impossible) to construct models of the Sun that agree with all the available data including the measured neutrino fluxes. Gallium experiments are particularly important in this context as the solar luminosity highly constrains the pp flux. As can be seen from figure 1.5, one possible solution to get around the constraint from the gallium experiments is to assume that there are no ^{7}Be neutrinos, although neutrinos from the ^{8}B reaction have been observed. In the pp chain ^{8}B is made from ^{7}Be it therefore seems very hard to suppress ^{7}Be and not ^{8}B; the lack of ^{7}Be neutrinos therefore seems to be somewhat of a paradox. A summary of this problem is shown in figure 1.6 where a fit to the neutrino fluxes has been carried out. In this analysis the shape of the neutrino spectra are constrained and each flux is a free parameter. The best fit results in a negative ^{7}Be flux and even pushing the fit to the limit these results are inconsistent with a number of non-standard solar models. This fit to the data is described in [25].

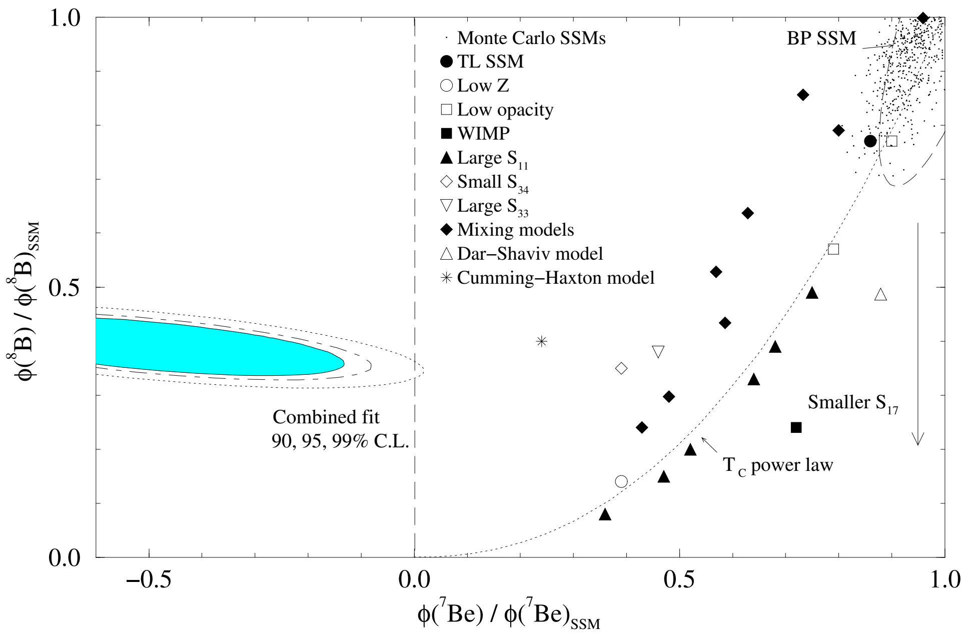

Figure 1.6: The results of the fit to the neutrino fluxes from [25]. Also shown are the 90% C.L. region from the Bachall-Pinsonneault standard solar model, the core temperature power law relating the two fluxes, and a number of non-standard solar models all of which are inconsistent with the data.

- Our understanding of neutrino physics is incorrect.

  This is the solution to the solar neutrino problem that seems to hold the most promise. In particular if neutrino oscillations do occur, then they can be used to explain the solar neutrino problem.

1.5.1 Neutrino oscillation solutions

Neutrinos are created as electron neutrinos in the centre of the Sun. The chlorine and gallium experiments are only sensitive to electron neutrinos, and although water Čerenkov detectors are sensitive to all active neutrinos they are much more sensitive to electron neutrinos than μ or τ neutrinos. Thus if some of the electron neutrinos created in the centre of the Sun change flavour, via neutrino oscillations, on transit to the various solar neutrino experiments then the solar neutrino problem can be explained.

The solar neutrino problem can be explained by vacuum oscillations alone. Vacuum solutions however require tuning of physical parameters to explain the observed deficit of neutrinos. In particular the oscillation length must be close to one astronomical unit.

The solar neutrino problem can also be explained by neutrino oscillations involving the MSW effect. Analyses are carried out using the solar neutrino data discussed, constraints from the standard solar model and data from the CHOOZ reactor experiment[26]. Three regions of the MSW plane ($\tan^2(\theta_V)$,Δm^2) can be found that explain the experimental data. These are the LMA (Large Mixing Angle, this solution occurs at high values of Δm^2), SMA (Small Mixing Angle) and LOW (this solution also has high values of the mixing angle but occurs at lower values of Δm^2 than the LMA solution) solutions. Vacuum solutions appear at even lower values of Δm^2. The results of one such MSW analysis of the solar neutrino problem for oscillation to active neutrinos are shown in figure 1.7 and are discussed in [27]. Similar analyses are carried out for oscillation to sterile neutrinos. These show similar results to those obtained when oscillations to active neutrinos are considered, except that the size and confidence of the LMA and LOW solutions are greatly reduced.

1.6 Conclusions

A number of experiments have been carried out that measure the rate of neutrinos from the Sun. Each of these experiments measures a flux of neutrinos that is significantly lower than that predicted by the otherwise successful standard solar model. This is known as the solar neutrino problem. The theory of neutrino oscillations including the MSW effect can explain the observed data, but there is no direct evidence that neutrino oscillations are the solution to the solar neutrino problem. The SNO detector, which can observe solar neutrinos via a number of different reactions, can provide this direct evidence, and will be the subject of the next chapter.

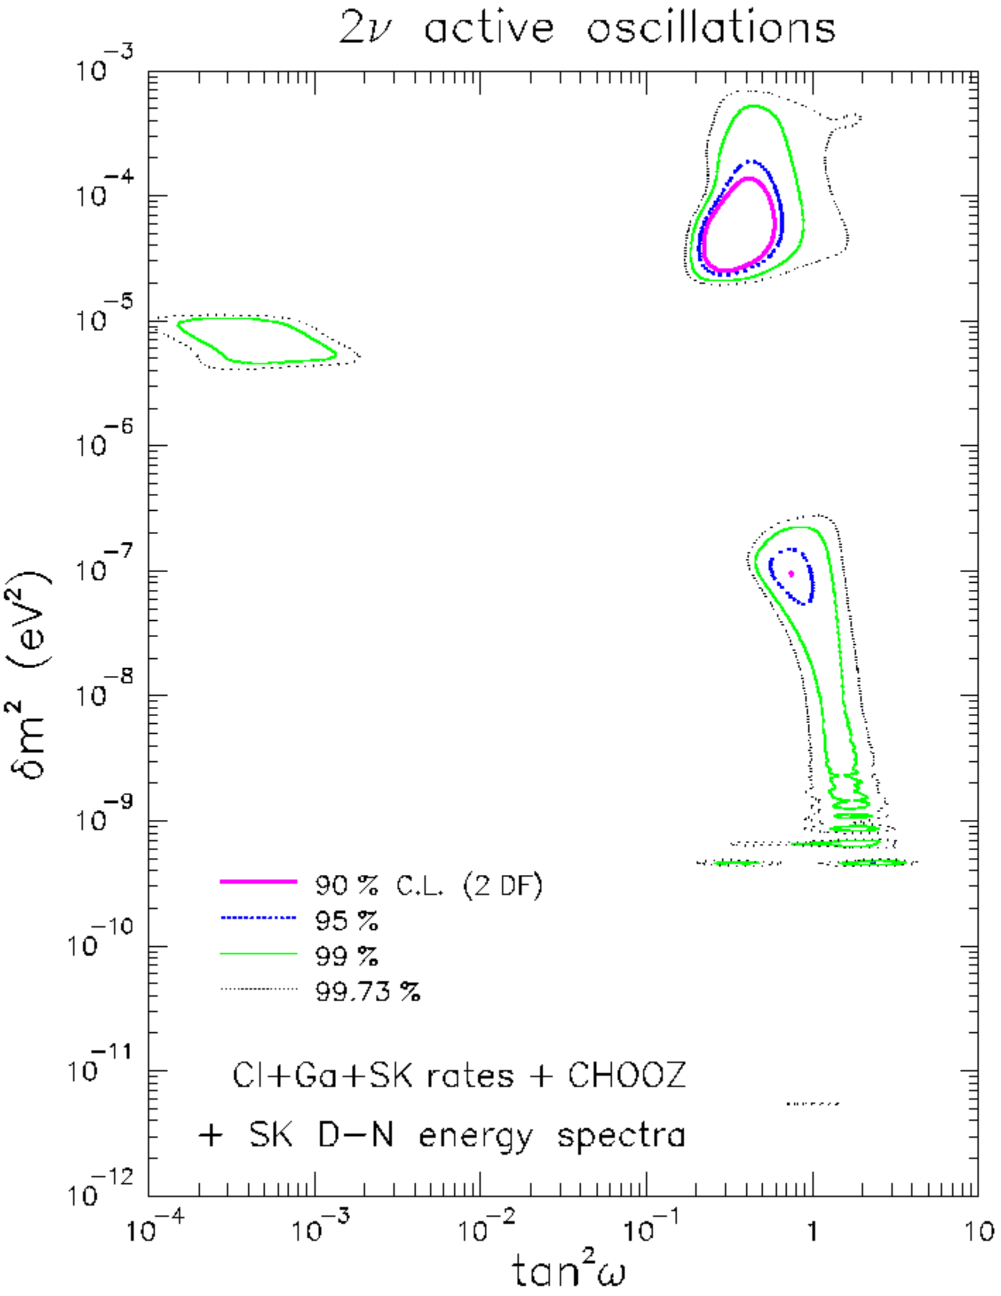

Figure 1.7: The results of the analysis of solar neutrino data including Super-K day-night energy spectra and the CHOOZ result. The plot shows the LMA, SMA, LOW and vacuum solutions in the Δm^2, $\tan^2 \omega$ plane, where ω is the mixing angle. Taken from [27].

Chapter 2

The SNO Detector

In order to attempt to solve the solar neutrino problem the SNO detector was built. SNO, based upon an idea by Herb Chen[28], is a 1kilo-tonne heavy water Čerenkov detector, in which deuterium is the primary target for solar neutrinos. The different neutrino reactions on deuterium allow SNO to make measurements that can definitively demonstrate neutrino flavour oscillations.

2.1 Čerenkov light

When a charged particle propagates through a medium faster than the speed of light in that medium, then it will emit Čerenkov light. The number of photons emitted is controlled by

$$\frac{d^2N}{dx d\lambda} = \frac{2\pi\alpha Z^2}{\lambda^2}\left(1 - \frac{1}{(n\beta)^2}\right) \tag{2.1}$$

where x is the distance travelled by the charged particle, λ is the wavelength of the Čerenkov photon, Z is the charge on the particle, α is the fine structure constant, $\beta = \frac{v}{c}$ and n is the refractive index of the medium. The emission of Čerenkov light occurs at a characteristic angle:

$$\cos(\theta_C) = \frac{1}{n\beta} \tag{2.2}$$

which in water for particles with $\beta \approx 1$, corresponds to 42^o. Thus an electron of a typical SNO energy (5-20MeV) emits a cone of light with a half angle of 42^o.

The Čerenkov threshold occurs when the speed of the particle corresponds exactly to the speed of light in the material. For electrons in water this corresponds to approximately 270keV. This provides a theoretical energy threshold for water Čerenkov detectors, although in practice the energy threshold is much higher due to low energy radio-active decays. For events well above threshold, the number of Čerenkov photons emitted is proportional to the energy of the particle with approximately 220 photons in the range between 200 and 600nm, emitted per MeV.

2.2 Neutrino interactions in SNO

There are two principle interactions of neutrinos with deuterium: the charged-current (CC) and neutral-current (NC) reactions. These reactions can be used together to show definitive proof of neutrino oscillations as the solution for the solar neutrino problem. SNO will also detect neutrinos via the elastic scattering interaction of neutrinos on atomic electrons, as with all water Čerenkov detectors. Due to the threshold of these reactions SNO is sensitive to neutrinos from the ^{8}B and hep reactions only.

2.2.1 The charged current interaction

Although in general the term charged current interaction refers to interactions involving the $W^{\pm}$ boson, in SNO it has a much more specific meaning. The charged current interaction is the reaction of the electron neutrino on the deuteron:

$$\mathrm{D} + \nu_e \rightarrow \mathrm{p} + \mathrm{p} + e^-. \tag{2.3}$$

This occurs via the interchange of a W^+ boson, and this results in the production of a relativistic electron. This electron therefore emits Čerenkov light in the water, which is detected by SNO. Because of their mass the protons produced are essentially stationary, meaning that the energy of the produced electron is correlated to that of the incoming neutrino minus the Q value of 1.44MeV. This reaction can therefore be used to measure the incoming neutrino energy spectrum as well as the flux of electron neutrinos.

The CC reaction is a Gamow-Teller process resulting in an angular distribution of the form:

$$f(\theta_\odot) = 1 - \frac{v}{3c}cos(\theta_\odot) \approx 1 - \frac{1}{3}cos(\theta_\odot). \tag{2.4}$$

This is a weak directional dependence for the charged current events, made harder to see by the non-zero angular resolution of the detector. There are several cross-section calculations for this process giving an error of approximately 3% in the total cross-section. This number provides an important theoretical limit for any CC measurement made by SNO. The theoretical cross sections are available both in tabulated [29] and functional [30] form.

2.2.2 The neutral current interaction

In SNO terminology, the neutral current interaction, as was the case for the charged current, refers to a specific reaction. This reaction is the neutrino disintegration of the deuteron.

$$\mathrm{D} + \nu_x \rightarrow \mathrm{p} + \mathrm{n} + \nu_x. \tag{2.5}$$

This reaction has the same cross section for all active neutrino types (ν_x). For this reason it can be used to measure the total flux of neutrinos above the 2.22 MeV energy threshold, irrespective of any active neutrino oscillations. The cross section is calculated in the same way as for the CC case, with a similar theoretical error. The theoretical error is however correlated between the two cases. This results in an uncertainty in the ratio of the cross sections of $\sim 1\%$. This reaction is the principle reason for the use of D_2O in SNO, as it can be used in combination with the charged current flux to provide very strong evidence for neutrino oscillations.

Neutral current interactions are detected via the neutron produced. Since the detection of neutrons in SNO is not straightforward, there are three different neutron detection techniques that will be used by SNO. The first of these is to use pure D_2O, and neutron capture on deuterium. This produces a 6.25MeV γ-ray which can be detected by the Čerenkov light produced from Compton scattered electrons. The next stage of operation, which began on the 28$^{\mathrm{th}}$ of May 2001, involves the addition of salt (NaCl) which adds ^{35}Cl. The addition of this isotope vastly increases the neutron detection efficency of the detector, with $\sim$ 80% of neutrons capturing on it. The excited state of ^{36}Cl produced then decays producing a number of γ-rays with a total energy of 8.6MeV. As was the case in pure D_2O these γ-rays Compton scatter electrons which produce Čerenkov light. The final approach to be used in SNO is independent of Čerenkov light. An array of ^{3}He proportional counters will be deployed through the D_2O region of the detector.

2.2.3 The elastic scattering interaction

The elastic scattering reaction is the same as that used in both the Kamiokande II and Super-K detectors. This reaction involves the elastic scattering of neutrinos from atomic electrons.

$$\nu_x + e^- \rightarrow \nu_x + e^-. \tag{2.6}$$

The reaction if very strongly peaked away from the direction of the incoming neutrinos (the sun), a property which follows directly from conservation of momentum. Unlike the CC and NC reaction, for which the calculation of the cross sections involve complex hadronic interactions, the cross section of the elastic scattering interaction follows directly from electro-weak theory.

$$\frac{d\sigma}{dE_e} = \frac{2G^2 m_e^2}{\pi}\left[g_L^2 + g_R^2\left(1 - \frac{E_e}{E_\nu}\right)^2 - g_L g_R\left(\frac{E_e}{E_\nu}\right)\right] \tag{2.7}$$

where $g_L = (\pm\frac{1}{2} + \sin^2\theta_W)$ and $g_R = \sin^2\theta_W$. The upper sign corresponds to ν_e scattering and the lower to $\nu_{\mu,\tau}$.

Using equation 2.7 the cross section for ν_e is approximately 6.5 times larger than for $\nu_{\mu,\tau}$, as the reaction can occur by W exchange for the ν_e as well as Z exchange available to all active neutrino types. Elastic scattering is therefore an unequal mixture of charged and neutral current signal, which can be used in combination with a charged current, or neutral current signal to measure neutrino oscillations.

The cross section for elastic scattering is much smaller than for the charged current reaction. This is only partially compensated for by the increased number of electrons compared to deuterons in the detector. There are therefore many more charged current events expected than events from the elastic scattering reaction.

2.2.4 Anti-neutrino interactions

In addition to the interactions of neutrinos already discussed, anti-neutrinos also have distinctive interactions in the SNO detector. The most distinctive reaction occurs on the deuteron, and is analogous to the charged current interaction.

$$\overline{\nu}_e + \mathrm{D} \rightarrow \mathrm{n} + \mathrm{n} + e^+. \tag{2.8}$$

The detection of one (or both) neutrons in coincidence with the positron provides a very strong signal for anti-neutrinos. There is a second reaction of anti-neutrinos that will occur in the SNO detector. This is the inverse β decay of the proton.

$$\overline{\nu}_e + \mathrm{p} \rightarrow \mathrm{n} + e^+. \tag{2.9}$$

This reaction will occur in the H_2O region, and will not be as distinctive as the CC analogue as the neutron is unlikely to be detected in this region.

2.3 Neutrino oscillation signatures in SNO

There are three distinct ways that SNO can identify neutrino oscillations as the solution to the solar neutrino problem. These effects are known as *smoking guns* for neutrino oscillations. Each of these tests can be carried out with the use of SNO data alone, but the results of other experiments can be included if appropriate.

2.3.1 Appearance of $\nu_{\mu,\tau}$

SNO can measure the flux of electron neutrinos by the CC reaction. The NC reaction is however sensitive to the total flux of active neutrinos. The comparison of the NC flux to the CC flux can therefore reveal a $\nu_{\mu,\tau}$ contribution to the total neutrino flux from the sun. If this is the case then the appearance of $\nu_{\mu,\tau}$ in the total neutrino flux demonstrates that neutrino oscillations occur in the solar neutrino sector. Experimentally the measurement of the $\frac{CC}{NC}$ ratio can be used to show that this is the case. The $\frac{CC}{NC}$ ratio is an extremely powerful measure of neutrino oscillations as the expected deviation from the no oscillation hypothesis is a many σ effect for all active neutrino oscillation solutions to the solar neutrino problem.

The elastic scattering flux also contains a contribution from $\nu_{\mu,\tau}$, although in this case the cross section is not equal to that for ν_e but is approximately 6.5 times smaller. The $\frac{CC}{ES}$ ratio can still however be used to demonstrate the occurrence of neutrino oscillations if the experimental error bars can be reduced sufficiently. In SNO many years of running are required to reduce the error bars on the ES flux sufficiently for this to be a powerful discriminant for most of parameter space. Super-K on the other hand measures solar neutrinos to very high statistical precision using solely the ES reaction[19]. A comparison of the CC flux measured at SNO to the Super-K measurement of the ES flux can therefore potentially be used to demonstrate neutrino oscillations at a much earlier stage than using SNO data alone.

2.3.2 Energy distortions

In figure 1.1 the survival probability as a function of energy is shown for the various oscillation solutions to the solar neutrino problem. SNO will initially measure the CC energy spectrum above a threshold in the region of 5-7 MeV for the electron energy. This corresponds to an examination of the neutrino energy spectrum above a threshold of 6.5-8.5 MeV. Particularly for the SMA and vacuum solutions large energy distortions are predicted in this region. The observation of a distortion of the energy spectrum is another “smoking gun” for neutrino oscillations.

The charged current interaction has an almost one to one correspondence between electron and neutrino energy, in comparison to the elastic scattering reaction which is an integral process. For this reason the CC spectrum is more sensitive to spectral distortions than ES spectrum, meaning that SNO has the opportunity to observe spectral distortions with less data than detectors using the ES reaction only, particularly if the energy threshold for analysis can be reduced.

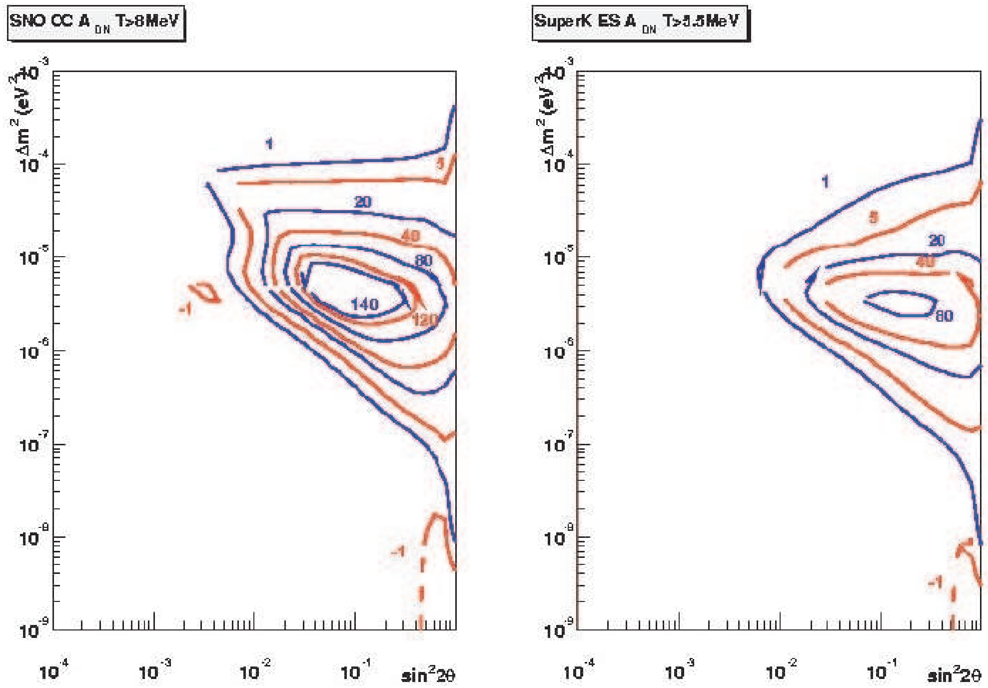

Figure 2.1: Contours of constant day-night asymmetry A_{ND} on the MSW plane for the CC signal above 8MeV and the ES signal with a 5MeV threshold. A_{ND} ranges between $\pm 200\%$. Courtesy of [31].

2.3.3 Day-night asymmetries

The final clear signature for neutrino oscillations, in particular for MSW oscillations is the observation of a day-night asymmetry. A day-night asymmetry can occur as the neutrinos undergo further oscillation as they travel through the material of the Earth. The density and size of the Earth is such that significant day night effects are expected for regions of the MSW plane populated by solutions to the solar neutrino problem. Contours of constant day-night asymmetry are shown in figure 2.1. The plot shows that the most likely solution in the regions of parameter space relevant to solar neutrinos is a positive asymmetry. A positive asymmetry occurs when the night flux is higher than that measured during the day, due to the regeneration of ν_e as the neutrinos pass through the Earth.

The distance travelled through the Earth is dependant upon the solar zenith angle (θ_Z). The density profile of the Earth also changes as a function of the θ_Z particularly at the core mantle boundary at $\cos(\theta_Z) < -0.837$. There is therefore particular interest in the measured neutrino flux as a function of the position of the sun, and these studies can give us more handles on the exact MSW solution than observation of the day-night asymmetry alone.

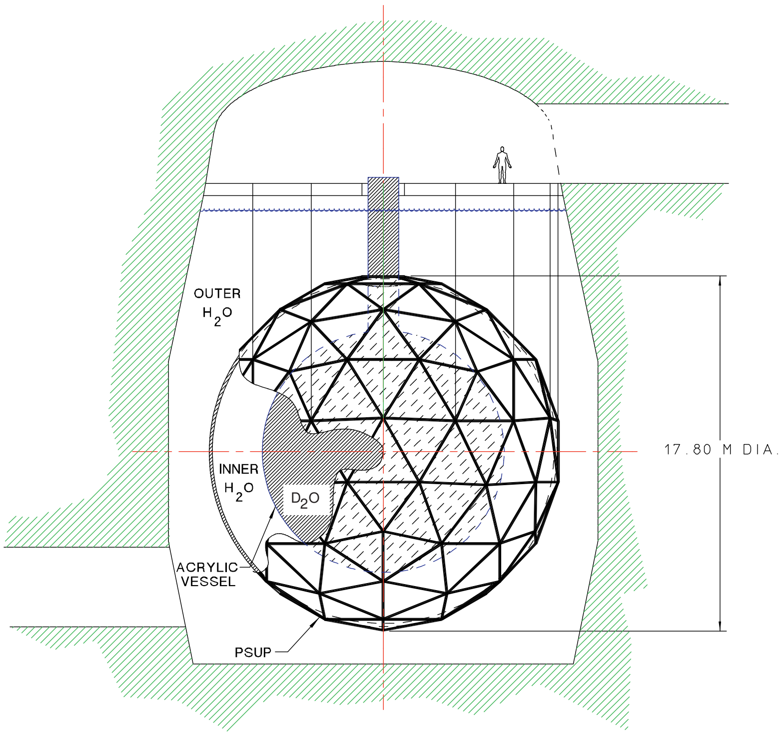

Figure 2.2: A schematic of the SNO detector. Courtesy of [32]

It is also possible in SNO to measure the day-night asymmetry of the neutral current flux. This can be used to provide clear evidence of oscillation to sterile neutrinos, if this process occurs. Observation of a day-night asymmetry in the ES signal alone cannot do this due to the large CC component. However the day-night asymmetry of these three signals is not independent, and it will be an important experimental cross check to ensure that the day night asymmetries observed in the three signals are consistent with one another.

2.4 The SNO detector

2.4.1 The detector

A schematic of the SNO detector[32] is shown in figure 2.2. SNO consists of 1000 tonnes of ultra-pure D_2O, contained within a 6m radius, almost spherical acrylic vessel (AV). This is surrounded by the PMT support structure (the PSUP), a 17.8m diameter geodesic sphere. This contains 9438 inward looking photo-multiplier tubes (PMTs), and a further 91 PMTs looking outwards (OWLs). The entire detector is contained within a barrel shaped cavity, 22m in diameter and 34m in height. The detector is surrounded by 7000 tonnes of pure H_2O, both between the acrylic vessel and the PSUP and the PSUP and the rock wall.

The detector is located in the INCO, Ltd., Creighton mine, near Sudbury, Ontario. It is situated on the 6800 foot level and the centre of the detector is 2039m below the surface. This corresponds to shielding equivalent to approximately 6000m of water. This provides a large amount of shielding against cosmic ray muons, resulting in a through-going muon rate of 3 muons every hour. Creighton mine is a working mine, and this has consequences for the operation of SNO, in particular with periodic blasting, principally during the middle of the night.

The primary target material in SNO is the D_2O, due to the neutrino reactions on the deuteron, and the neutron transport properties of D_2O. The D_2O is ultra-pure, both radio-chemically and isotopically. The hydrogen fraction is 0.08% by mass [33]. Even this small fraction of hydrogen contributes significantly to the transport of neutrons within the D_2O. Radio-chemical purity is required to reduce the amount of uranium and thorium in the detector. Decays from elements at the bottom of the decay chains of these elements create a low energy Čerenkov background, known as *the wall* and some decays produce a γ-ray above 2.2MeV which can photo-disintegrate the deuteron.

The spherical part of the acrylic vessel is built out of 122 separate low radio-activity panels. These all transmit ultra-violet light, important as an even larger fraction of the detectable Čerenkov light would otherwise be attenuated by the vessel. Most of these panels are nominally 5.6cm thick, with the exception of the belly plates. The belly plates are 10 equatorial panels, containing groves for the support ropes. The vessel is supported by 10 Vectran ropes that attach to the belly plates, and suspend the vessel from above. This method of suspension is both robust and has a minimal effect on Čerenkov light produced inside the vessel. The acrylic vessel also has a neck to allow access to the D_2O region for both calibration sources and deployment of the neutral current detectors. The neck is visible in figure 2.2.

There are two H_2O regions in the detector. The inner region lies between the AV and the PSUP, and the outer lies between the PSUP and the rock wall. These regions are primarily used to shield the detector and the D_2O from external radiation. The outer volume shields the detector from high energy γ-rays and neutrons emanating from the rock wall of the cavity. The inner region shields the D_2O from γ-rays from both the PSUP and PMTs. The inner H_2O region can also be used to extend the signal region for the elastic scattering reaction of solar neutrinos. The outer region works in conjuction with the OWL tubes to reject cosmic ray muons passing through the detector.

The PMT support structure houses the 9438 inward looking PMTs which are the primary method of data collection in SNO. The PMTs used by SNO are Hamamatsu R1408 photomultipliers, with a photocathode diameter of 20cm. Each tube is built out of Schott 8246 glass, which has very low radioactive content and was developed in collaboration with Oxford for the SNO experiment. Each of these PMTs is surrounded by a concentrator to increase the coverage of the detector. The coverage is the fractional solid angle subtended by the PMTs, and with the addition of concentrators this is increased from 31% to 59%. The reflectivity of the concentrator reduces this to an average of 54%. As well as the increased coverage the concentrators also aid the detection of light from a region inside 7m while preventing each tube from seeing its nearest neighbours[34]. This helps reduce the Čerenkov background due to radioactive decays in the PMTs.

As well as housing the PMTs the PSUP also houses the 91 outward looking PMTs (OWLs) used as veto tubes. The OWLs are not however the only tubes used as veto tubes in SNO. A second class of outward looking tubes were added during the commissioning phase of the detector. These are known as the BUTTs,[1]. These tubes were added to the outer H_2O volume to allow PMTs to be tested in the SNO environment. Twenty-two of these tubes are connected to the electronics and can therefore be used as veto tubes. There are a further 4 PMTs added to the neck region to further veto events. The use of all these tubes will be discussed in chapter 4.

2.4.2 The SNO electronics

Directly above the detector is the region of the lab called the Deck. This region of the detector houses the SNO electronics [35],[36],[37],[38]. The electronics in SNO are responsible for two principal things: the measurement of the signals from the PMTs and for triggering the detector.

The collection of the PMT signals

Every PMT in SNO (including non-standard PMTs) is connected to a channel of the electronics. The electronics are divided in the following way. There are 19 *crates*, containing 16 *cards* of 32 *channels* each. This provides 9728 electronic channels in SNO, enough for all the PMTs with some spares that can be used for other things, such as connection to PMTs inside calibration sources. Each card contains 4 *paddle cards* controlling 8 channels including the high voltage to the PMTs.

When the signal from the PMT exceeds a threshold the discriminator fires and

[1]BUTTs is short for Berkeley Underwater Tube Testing

the time to amplitude converter (TAC) system is started. This continues until either a trigger signal is received, or the system runs out of time and is reset. If the trigger signal is received then the time of the PMT, measured by the charge on the TAC capacitor, is stored on one of 16 cells. It is also important to measure the charge deposited on each PMT. The charge signal is split into two paths, with different gains in a ratio of 16:1. The charge signal is then integrated over two different time windows: 70 and 400 ns to provide coverage for direct and late light. Three charge signals QHS, QHL and QLX[2] are each stored on a separate bank of 16 cells. Only one of the two low gain charges is stored, and the long or short time window depends upon the settings of the DAQ software, typically QLL is used.

The 4 stored charges are read out by analogue to digital converters (ADCs) and are converted to 12 bit digital signals. These are then built together along with the GTID (see below) and the channel number to create a PMT bundle, that is passed to the DAQ software.

The SNO trigger system

In order for the detector to collect data a system is needed that controls when the PMT data is read out. This system is called the trigger system, and is conceptually very simple. If an interesting event occurs inside the detector the trigger system must realise this and cause the PMT data to be read out. There are two principle types of trigger used by SNO: the NHIT and ESUM triggers.

The NHIT triggers are based upon the coincidence of a certain number of inward looking PMT hits within a specified time window. There are two time windows used in the trigger system, one nominally 20ns and the other nominally 100ns wide. The latter is used as the principal trigger, as the time window is closely matched to the transit time of light across the detector. A number of triggers are available using this system, three using the 100ns window: NHIT100lo/med/hi, and two using the 20ns NHIT20/NHIT20LB. For typical neutrino running the NHIT100med/hi and NHIT20 triggers are used, and the NHIT100lo is prescaled such that one event in every 1000 is recorded. The typical NHIT100med trigger threshold is set between 16 and 18 hits, corresponding to an energy of approximately 2MeV. In order to avoid problems that can occur with particular channels the NHIT trigger on any channel can be disabled while the data from that PMT is still included in each event.

The second type of trigger is the ESUM trigger. This is based upon the linear sum of the signals from every inward looking PMT. Thus this trigger will fire if a number of

[2]The nomenclature for the charges is as follows: Q Gain Time. There are the High and Low gain measurements over the Short and Long integration times.

PMTs fire and the charge exceeds threshold, or if a large amount of charge is deposited on a single PMT. Unlike with the NHIT triggers it is not possible to remove PMTs from analogue sum that is fed into the ESUM trigger system. As with the NHIT triggers there is more than one ESUM trigger as well, a low and high gain version: ESUMlo/hi. In this case counter-intuitively if ESUMlo fires then a larger ESUM signal has been recorded than if only ESUMhi fires.

As well as these two principal triggers a number of other triggers exist. The pulsed global trigger (PulseGT) is fired at a rate of 5Hz to record the state of the detector at that time. This is excellent for diagnostics and to record the noise rate on the PMTs. There are a number of external triggers for calibration sources, and the pedestal trigger, for the forced readout of particular channels. Pedestal events are used for calibration of the electronics. Finally there are two OWL ESUM triggers and an OWL NHIT trigger as well that are used for muon studies.

When a trigger that is masked in (enabled) fires this causes the electronics to generate a global trigger. This has two consequences: it causes readout of the PMT signals in the form of PMT bundles, and the generation of an event header from the master trigger card (MTC). The event header contains the non-PMT information about the event. This includes two sources of time information: a 10MHz clock, synchronised with a GPS receiver on the surface, and a 50MHz clock based upon a local oscillator. The trigger that caused the event is also recorded in the event header along with any other triggers that fired within 10ns. Any triggers that occur after this time are recorded as missed triggers. Perhaps most importantly the firing of the trigger causes a 24 bit counter to be incremented that provides a unique id number to each event. This number is known as the global trigger id (GTID) and is written to both the event header and the PMT bundles. The PMT bundles and the event header are then sent out of the electronics to be read out by the DAQ.

2.4.3 Data acquisition

In SNO data acquisition (DAQ) consists of two separate systems: SHARC (SNO Hardware Acquisition and Readout Control) software controls the hardware of the SNO detector and the builder which combines the signals from the detector to build events.

SHARC provides a graphical interface for detector control. It can be used to control detector settings, such as channel thresholds or trigger settings, and status of PMTs including high voltage control. The status of these settings is stored in data files known as DQxx banks, written by SHARC. SHARC is also used to start different runs[3] of

[3] A run is a convenient sized period of data taking. The splitting of the data into runs is an important part of bookkeeping.

various types including electronic calibration runs. From the point of view of the detector operator, SHARC is also used to help monitor the state of the detector, providing useful detector diagnostic information to the user.

The operation of the builder is more relevant to the contents of this thesis. The builder is a program that takes the event headers and PMT bundles output by the detector and combines them to form events. This is done on the basis of the GTID contained in each. The builder can be overwhelmed with data during periods of extremely high rate, usually due to problems with the detector. When this happens it can fail to combine all PMT bundles with their event headers. The unattached PMT bundles are known as orphans.

2.4.4 The water systems

Low energy decays due to uranium and thorium chain isotopes can cause problems for the operation of SNO. These decays can emit Čerenkov light providing a low energy background wall that cannot be avoided in the analysis of solar neutrinos. More importantly decays that create a γ-ray of more than 2.2MeV can photo-disintegrate a deuteron, providing a signal in the detector that is indistinguishable from the neutral current interaction. It is therefore very important to measure and control the radio-active content of the detector. For this reason all detector components were built with materials that had as low a radioactive content as possible, and a clean room was built underground for the detector.

To achieve this cleanliness for the D_2O and H_2O, two separate water systems were designed to clean the water and measure the content of uranium and thorium. The heavy water is circulated through ultra-filtration membranes and a reverse osmosis (RO) system to clean the water. The water is also de-gassed in order to remove radon. The light water system passes the water through a RO system and a de-gasser, as well as a UV unit to kill any bacteria present on the water. Nitrogen gas is then added to the water to avoid problems with the high-voltage connectors on the PMTs. The light water is continually circulated through this system, while the heavy water system is only operated at certain times. To prevent radon in the lab air from being dissolved in the water both the D_2O and H_2O have nitrogen cover gas systems. Nitrogen boiled off from a liquid nitrogen dewer is continually passed over these regions.

In order to measure the uranium and thorium content of both the D_2O and the H_2O assays are preformed. There are three principal assay systems: MnO_x coated beads and a HTiO based seeded ultra filtration system[39] are used to measure the content of uranium and thorium and their decay chains. The radon content of the water is measured using a process de-gasser. Target levels for the concentration of uranium and thorium in the D_2O

have been set to a contribution equivalent to 5% of the predicted NC signal, while targets in the H_2O are based upon the contribution to the background wall. Water assay results suggest that these targets have been met or surpassed, with the exception of thorium in the H_2O, which is slightly above target levels, but is consistent within the errors on the measurement[40].

2.5 SNOMAN

SNOMAN (SNO Monte carlo and ANalysis)[41],[42], [43] is the SNO collaboration's principle Monte Carlo and analysis package. It has been developed using FORTRAN 77, and the ZEBRA memory manager [44].

SNOMAN provides two principle functions to the collaboration. The first of these is as the Monte Carlo of the detector. SNOMAN models the propagation of electrons, γ-rays, muons, neutrons and Čerenkov photons, using code based upon EGS4[45], MCNP[46] as well as in-house code. The propagation of these particles through the complex SNO geometry, a detailed simulation of the response of the PMTs, and the simulation of the trigger system, provides the basic Monte Carlo for SNO. The Monte Carlo is described in more detail in [47],[48] and [49] amongst others.

SNOMAN also contains a number of important analysis tools such as PMT calibration tools, reconstruction tools, event filters and various other tools. These will be discussed as and when they are used in this thesis. The analyses carried out in this thesis have all been developed within the SNOMAN environment, and the analysis tools developed are installed in the code.

2.6 Calibration of SNO

The calibration of SNO can be split into two distinct categories. The first of these is the calibration of the individual PMTs, taking the ADC values recorded and turning them into meaningful measurements of time and charge. The calibration of events is also important, turning the raw measured values from an event into meaningful quantities, and to configure SNOMAN to reproduce the data.

2.6.1 The calibration of PMTs

There are three separate steps used in the calibration of the signals from the PMTs. These are the ECA, HCA and PCA.

The ECA is the electronics calibration[50]. This converts the digital signal from the TAC to a time in nano-seconds, using the TAC slope information. The TAC slope is the mapping between the ADC value and the real hit time, so-called due to the close to linear behaviour over the majority of times. The TAC slope calibration is carried out by a process of linear interpolation between calibration points along the slope. There are two ways that a channel can fail the time calibration, even if the channel has valid constants. The first of these is if the time falls in a region of the TAC slope where the interpolation points are inconsistent[4]. The second is if the time falls during the TAC curl. This is a region of the TAC slope at the very earliest times, that is double valued due to a feature of the electronics. The ECA also corrects the charge measurements for the pedestal, the value measured at zero charge. This provides a charge output that is still digital, but that now begins at 0 counts instead of a variable number of order 500. This is the usual charge used in SNO and is measured in units of counts above pedestal (CAP). The ECA also contains inbuilt validation, which can be used to remove pathological channels from certain analyses, and is very useful for monitoring the health of the detector. ECA calibration constants are obtained by special runs where the electronics is read out in a controlled way, using forced events known as pedestals.

The next step in the PMT calibration procedure is the HCA[51], or history calibrator. This is another feature of the electronics where it was found that the ECA calibration of time is rate dependant. More precisely the recent pattern of hits on a particular channel causes the TAC slope to be different to that measured during calibration, where the slope is measured at high rate. The HCA corrects the ECA time for this effect. There have been two separate hardware fixes implemented to reduce this problem. The magnitude of the rate effect was found to be temperature dependant, so the first fix was to turn off the principle heat source on the affected chip. The heat source was a feature known as the twiddle bits[5], they were disabled on the 17$^{\text{th}}$ of August 2000. Then on the 5$^{\text{th}}$ of December 2000 the second fix was installed. These were fans to improve the cooling of the crates and to make the temperature across the crate more uniform. These two fixes combined to significantly reduce the rate effect, but did not eliminate it, hence the need for the HCA. The typical size of the HCA correction was reduced from 5ns to less than 1ns, and the large tails in the distribution were dramatically reduced.

The final step in the calibration chain is the PCA[52] or the PMT calibration. The PCA is comprised of two parts, the time calibration and the charge calibration. Unlike ECA and HCA constants, which are obtained from pedestal events, the PCA uses calibration source data from the laserball[6]. The time calibration serves two purposes.

[4]This happens if the second of three consecutive interpolation points does not have a TAC value between the other two. This can occur due to real PMT hits during during the calibration procedure.

[5]The Twiddle bits were used to help equalise the duration of the TACs across the different channels in the detector

[6]The laserball is a source of optical photons derived from a laser and passed through a diffuser ball. The laserball is described briefly in chapter 5

Firstly the cable delays (which are not constant tube to tube) are calibrated out, ensuring that every PMT time is relative to the same t_0. The second process removes the effect of the discriminator walk. The larger a charge deposited on a PMT the quicker the rise time of the pulse, resulting in an earlier crossing of the discriminator. The size of the walk correction spans the range 0-5ns. The PCA also calibrates the charge response (gain) of the PMTs. This converts the charge from all PMTs into the same (arbitrary) units.

Once these levels of calibration have been applied the PMT timing is calibrated to approximately 1.5ns in agreement with ex-situ measurements of the PMTs.

2.6.2 The calibration of events

As well as the calibration of the signals from the PMTs the detector must be calibrated such that the Monte Carlo reproduces the data and that the response of the detector is understood. This form of calibration falls into two broad categories: optical and energy calibration.

The optical calibration is used in the Monte Carlo to model the propagation and detection of Čerenkov light. There are two principle effects that have been measured, the attenuation of the various optical media: the D_2O, H_2O and the acrylic that forms the vessel. The angular response of the PMT concentrator combination has also been measured. The optical calibration is carried out using optical photons of various wavelengths using the laserball. Due to limitations in the current technique, optical calibration has only been carried out using data taken inside the D_2O volume. This means that the angular response has only been measured for a subset of all possible angles, and that the response around the concentrator cut-off[7] has not been mapped out. The use of this data set also means that the co-variance between the attenuation of the acrylic vessel and the H_2O is not broken by this analysis, so ex-situ measurement must also be used. These problems are important when choice of fiducial volume is made. The optical calibration procedure, results and implementation are described in [53] and [54].

There are two measures of energy used by SNO, the total number of PMT hits: NHITs, and the energy measured by the RSP process. Energy calibration for NHITs is the simpler of the two calibration procedures, and relies on the reproduction of the expected spectra by the Monte Carlo. This is reliant on the measurement of the optical parameters and the value of the PMT collection efficency, provided that the creation and propagation of photons has been implemented correctly. The PMT collection efficency is set by the ^{16}N source at the centre of the detector. The Monte Carlo of the ^{16}N and other calibration sources in various positions can be compared to the data to measure the

[7] For angles above the cut-off the PMTs are blinded by the concentrator. It is this feature that prevents PMTs from being has seen by their near neighbours.

systematic error due to the energy scale.

The use of NHIT as a primary energy estimator is not ideal, as the NHIT of an event not only depends upon the energy of an event, but also the position and direction of the event as well. This variation is controlled by the optical properties of the detector, and is therefore modelled by the Monte Carlo, but there is some desire to remove these effects from the energy resolution function. This is achieved by the use of response processor (RSP), which corrects for these effects. Other than the correction for optical effects, the RSP process also differs from NHIT, by a selective use of PMT hits. Only hits in a prompt light window are used, as these are the only hits for which the current optical correction can be applied. Once the number of prompt hits has been corrected for the optical response of the detector, it is then converted into a most likely energy on the basis of calibration data and tube accounting. A similar analysis between calibration data and Monte Carlo is then carried out to determine the energy scale and its error. More detail on RSP and energy calibration can be found in [55], [56] and [57].

2.7 The measurement of the charged current flux

On the 18$^{\text{th}}$ of June 2001 the first results were published by SNO presenting a charged current flux measurement [58]. Many parts of this analysis will be discussed later in this thesis, or have already been touched upon in this chapter.

A dataset of 241 days of data was used (the same data that will be considered in majority of this thesis) and events that reconstruct inside a fiducial volume of 550cm and above a kinetic energy threshold of 6.75MeV were selected for the analysis. The choice of 6.75 MeV was made to remove the majority of the neutrons. This allows the charged current analysis to proceed without consideration of the extra effects, particularly the neutral current background due to photo-disintegration, that must be considered for a neutral current analysis.

The data analysis proceeds via a signal extraction procedure. This procedure is used to determine the number of charged current, elastic scattering and neutron events in the data set. This is carried out by a maximum likelihood technique using probability density functions (PDFs) for the three signal classes. These PDFs are functions of three variables: energy, reconstructed position (r^3) and the cosine of the angle between the direction of the event and the direction from the sun ($\cos(\theta_\odot)$). The results of this analysis give the following neutrino fluxes:

$$\Phi^{CC}_{SNO}(^8B) = 1.75 \pm 0.07(stat.)^{+0.12}_{-0.11}(sys.) \pm 0.05(theor.) \times 10^6 cm^{-2} s^{-1} \tag{2.10}$$

$$\Phi^{ES}_{SNO}(^8B) = 2.39 \pm 0.34(stat.)^{+0.16}_{-0.14}(sys.) \times 10^6 cm^{-2} s^{-1} \tag{2.11}$$

where the theoretical uncertainty comes from the CC cross section. The SNO measurement of elastic scattering has a large statistical error bar, but it is consistent with the high precision measurement made by Super-K[19]:

$$\Phi_{SK}^{ES}(^8B) = 2.32 \pm 0.03(stat.)^{+0.08}_{-0.07}(sys.) \times 10^6 cm^{-2}s^{-1}. \tag{2.12}$$

The SNO CC flux is significantly (3.35σ) lower than the elastic scattering flux measured by Super-K, although this measurement is dominated by systematics. This provides evidence for neutrino oscillations by a $\nu_{\mu,\tau}$ contribution to the elastic scattering signal. This can be stated in the following way:

$$\Phi_{SNO}^{CC} = \Phi_e \tag{2.13}$$

$$\Phi_{SK}^{ES} = \Phi_e + \epsilon\Phi_{\mu,\tau} \tag{2.14}$$

where $\epsilon = \frac{1}{6.481}$ and $\Phi_{\mu,\tau}$ is the flux of μ and τ type neutrinos. These equations can be solved to find Φ_e and $\Phi_{\mu,\tau}$. The results are best presented on the Φ_e,$\Phi_{\mu,\tau}$ plane, see figure 2.3. The SNO result is a vertical bar and the elastic scattering result is a diagonal band due to the contribution of μ and τ type neutrinos. The intersection of these two bands defines the region of parameter space given by the contours. The contours on this plot define the 68.27%, 95% and 99% confidence levels. The no-oscillation hypothesis defined by $\Phi_{\mu,\tau} = 0$ is ruled out at the 99.96% confidence level. The central region of the plane corresponds to a total active neutrino flux of:

$$\Phi_{SNO+SK}^{Tot}(^8B) = 5.44 \pm 0.99 \times 10^6 cm^{-2}s^{-1} \tag{2.15}$$

in close agreement with the latest prediction[16]:

$$\Phi_{BP2001}^{Tot}(^8B) = 5.05^{+1.01}_{-0.81} \times 10^6 cm^{-2}s^{-1}. \tag{2.16}$$

As well as a measurement of the CC flux, SNO also presented a measurement of the CC energy spectrum. This can be measured using the same approach as for the flux measurement, except that the PDF based upon the expected CC energy spectrum is not used. The extracted spectrum is consistent with the expected spectrum within statistics.

In section 1.5.1 the analysis of all neutrino data to find acceptable solutions in the MSW plane was briefly discussed. The SNO measurement of the charged current flux can be added to this analysis to further constrain $\tan^2(\theta_V)$ and Δm^2. This has been carried out in [27] and the results are shown in figure 2.4. Comparing the allowed solutions to those without the CC flux measurement (figure 1.7) it can be seen that the SMA and most of the vacuum solutions are now ruled out at $> 3\sigma$. A similar analysis also rules out all sterile neutrino solutions at the same confidence level.

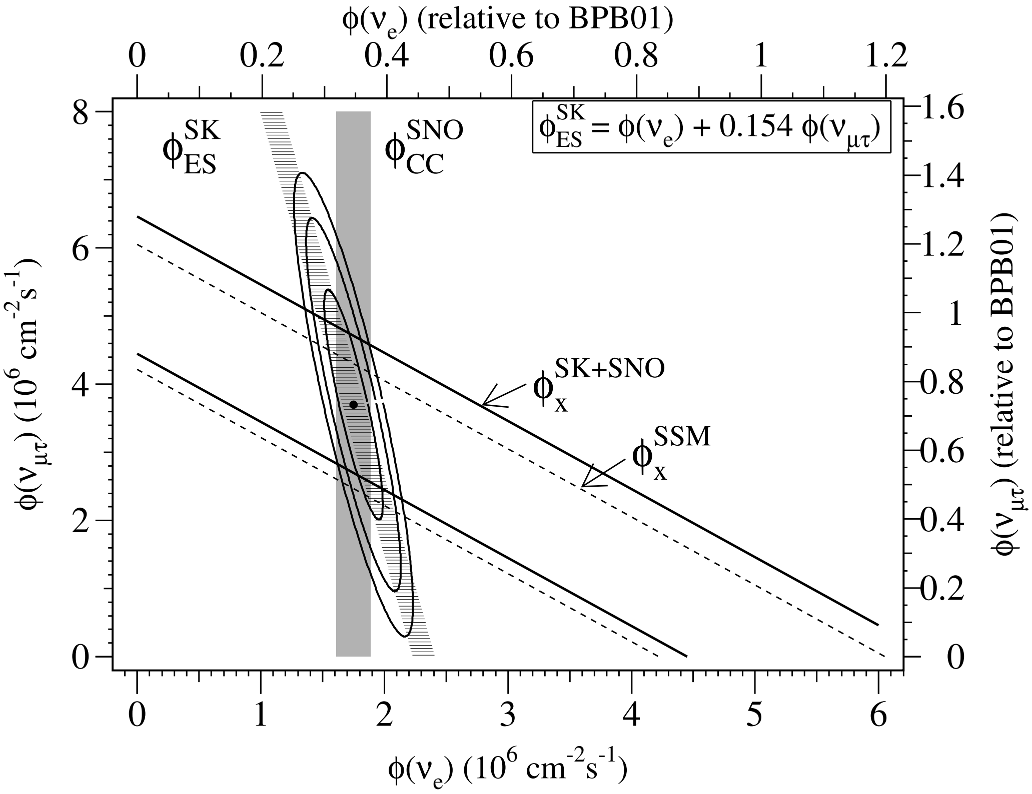

Figure 2.3: The flux of ^8B neutrinos detected as electron type compared to those detected as μ and τ type derived from the SNO CC and Super-K ES results. Courtesy of [58].

2.8 Conclusions

The SNO detector is a unique solar neutrino detector, able to detect solar neutrinos in a number of different ways. SNO has already published a measurement of the charged current flux, which in combination with the elastic scattering flux from Super-K gives strong evidence for neutrino oscillations. The next step for the experiment is to reduce the energy threshold to include a larger sample of neutrons and extract a measurement of the neutral current flux from pure D_2O. This measurement can potentially be used to provide evidence for neutrino oscillations from SNO data alone. One measurement of the charged current flux not reported in the first SNO paper, that can be of interest is the day-night asymmetry. The extraction of this value from the same data used to measure the charged current flux will be the subject of the final chapter of this thesis. The remainder of this thesis will discuss one of the principle supporting analyses required for the measurement of the solar neutrino flux by any reaction: the removal of the instrumental backgrounds.

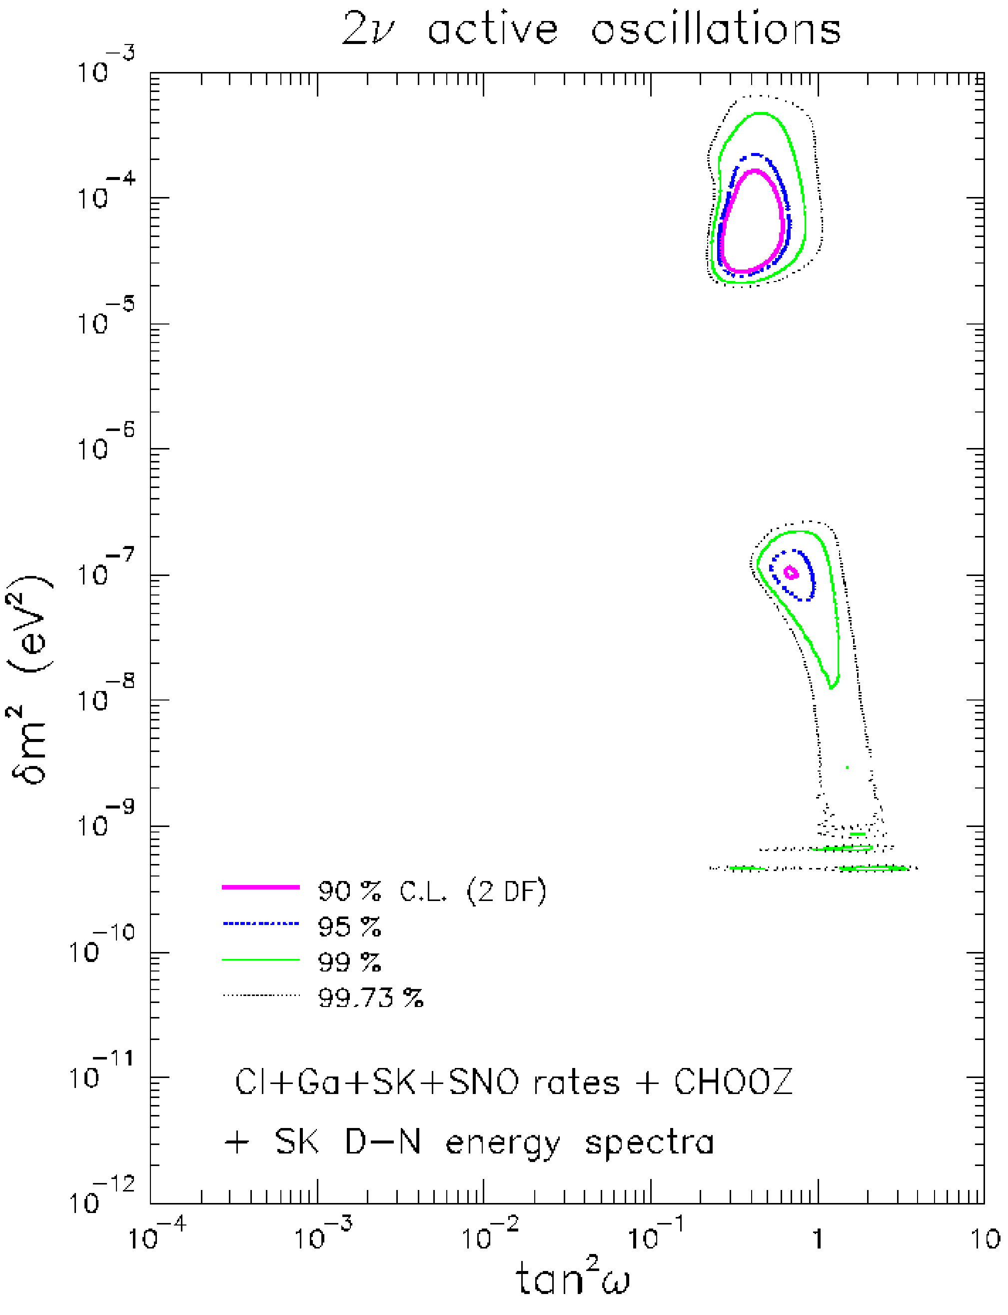

Figure 2.4: The results of the analysis of solar neutrino data including the SNO charged current flux, Super-K day-night energy spectra and the CHOOZ result Taken from [27].

Chapter 3

The Instrumental Backgrounds of SNO

Before the detector was turned on for the first time in 1998, the signal region in SNO was believed to be almost background free. At low energy there is the background wall due to βs and γs from decay products of uranium and thorium. At higher energy only a small number of background events were believed to exist mainly due to muon spallation products and a small number of high energy γ-rays from the PSUP that penetrate into the D_2O.

Unfortunately this state of affairs turned out to not be the case. SNO has a large problem with instrumental backgrounds. Instrumental backgrounds are events caused by the detector, that are not due to radioactive decays inside SNO. There are many causes of backgrounds, but electronic breakdown, static discharge, and electronic pickup are the most common. Events due to radioactive decays, such as PMT or D_2O $\beta\gamma$s are known as physics backgrounds.

The first hint of a problem was seen in 1998, when the detector was turned on for testing without the water inside (a time known as airfill). The tests led to the observation of flashers, which appear throughout the solar neutrino NHIT regime. Once the water had been added to the detector the full scale of the problem with instrumental backgrounds was seen. Flashers, while being the most numerous class of background, are only one of a number of distinct background classes due to the detector. The full scale of the problem can be seen from the raw NHIT spectrum for the neutrino data set, shown in figure 3.1. In this figure, the raw NHIT spectrum is compared to the expectation from the latest standard solar model [16].

The instrumental backgrounds have to be removed from the data almost completely before the neutrino signal can be analysed. This process is known as data cleaning. Backgrounds events are removed by data selection cuts. These have been designed to leave as little background in neutrino data sets as possible, a statistic known as contamination.

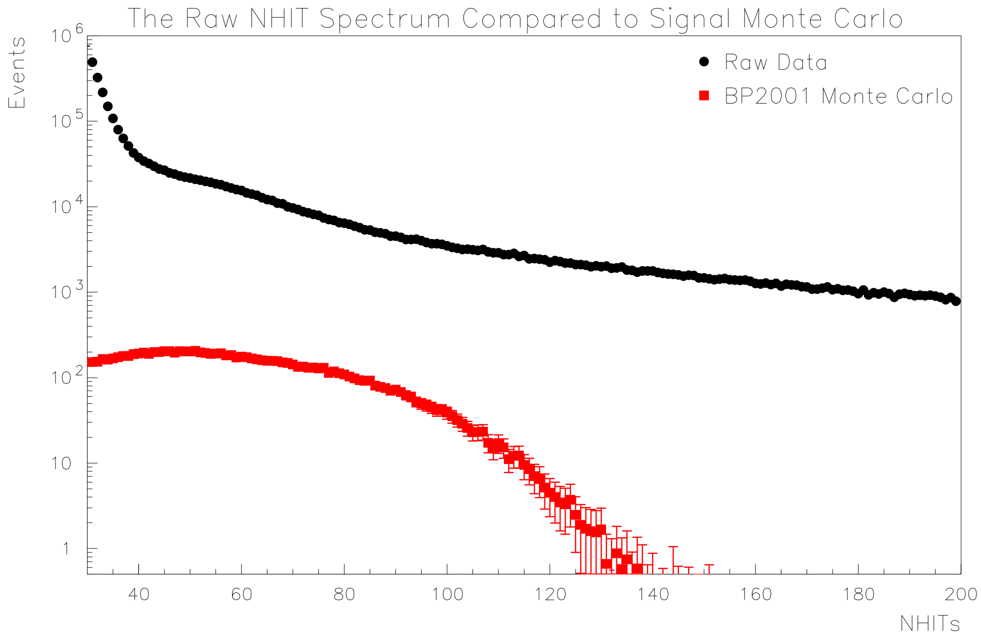

Figure 3.1: The raw NHIT spectrum for the neutrino data set compared to the expectation from the Monte Carlo of BP2001, assuming no oscillations.

They are also designed to accept as much signal as possible. The fraction of signal removed by the cuts is known as sacrifice. These two statistics feed directly into the measurement of the solar neutrino flux and as such they are the final deliverables of the data cleaning process[1].

3.1 The different classes of instrumental background

Before the suite of cuts can be developed the different classes of events from detector backgrounds need to be identified. This was done by hand scanning the data and examination of bursts of events. This approach is of course unreliable, and impractical as an approach to data cleaning, but with some experience it is possible to identify different classes of event and look for new ones. Once a class of events has been identified, work can begin on cuts to try and remove it. Knowledge of the different classes of detector background is also important for further study of the cuts, and to look for possible contamination of the neutrino data set.

The study of instrumental backgrounds can also in some cases lead to a permanent fix for a particular instrumental background. Once a background has been identified and

[1] Along with a suite of cuts to define neutrino data.

understood it has been possible in some cases, to design a hardware fix that will stop the events occurring. The action of preventing backgrounds before they occur is of course preferable to the removal of the events in software. If the event class is not present in the data stream, then there is no need for a cut to remove them. Unfortunately the prevention of background is only possible for some classes of events, partly because some background classes are still very poorly understood. Even when a fix has been implemented, this was not always carried out before the start of the neutrino data set, meaning that background events persist in the data to be analysed.

Background classes tend to fall into one of four particular categories:

- PMT backgrounds

  This is the dominant class of instrumental backgrounds where the PMTs themselves are responsible for the events. As PMTs work under high voltage, breakdown can occur and emit light. This is particularly likely to happen as a tube dies, but does happen throughout the life of the PMT.

- External backgrounds

  The PSUP is not a closed sphere, for this reason light can reach the inside of the detector from the outside and cause background events.

- Electronics Backgrounds.

  Both internal and external electronic pickup can occur and cause events.

- Acrylic Backgrounds.

  The final major source of background is the acrylic vessel. This contains pipes from the water systems, which can cause the emission of light in to the detector. The acrylic vessel also seems to emit isotropic light from various places. Due to the apparent energy of the events these events are also classed as instrumental backgrounds.

Each class of detector background deserves some discussion before the cuts and tags to be used are considered. First a typical neutrino event will discussed, followed by each background in the historical order of discovery.

3.1.1 A typical neutrino event

Figure 3.2 shows a typical neutrino event. This event provides a good opportunity to introduce the SNO event display. In the top left corner is a geodesic map of the PSUP also showing the acrylic vessel. The location of hit PMTs are shown on this map as the

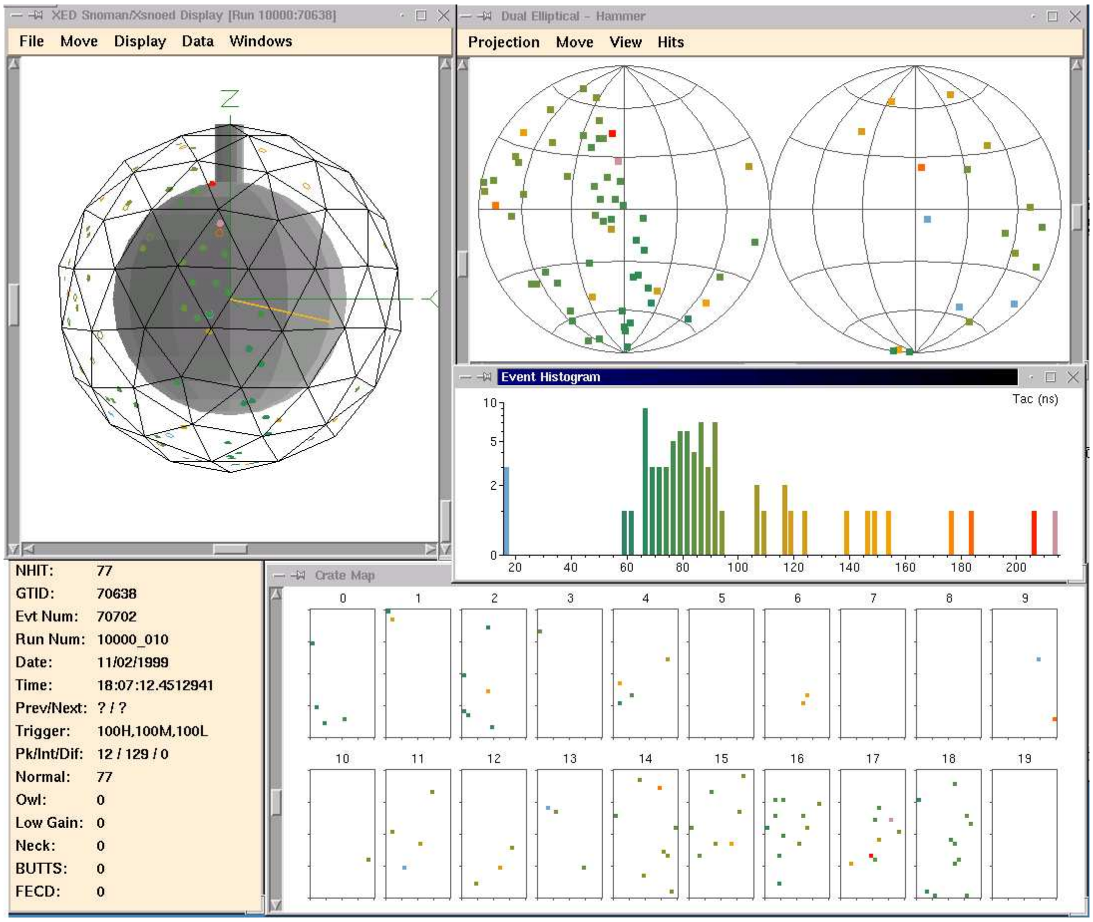

Figure 3.2: A typical neutrino event from SNO. In this display fully calibrated time is shown on the histogram and represented in the colours on the various projections.

coloured points. The colour of the points is related to the value of the time of the hit, and corresponds to the hits on the same bin in the histogram on the centre right. In this event the histogram is displaying fully calibrated time, although PMT charge can instead be displayed. The top right shows a projection of the PMT hits onto a two-dimensional map. This is just one of the many ways that a sphere can be projected into two dimensions. The bottom right shows the hits again this time in electronics space. Each box corresponds to one of the 19 electronics crates. Finally in the bottom left is a summary of some basic event information, such as the number of PMT hits, the run and event number and various other event information.

The neutrino event shown in figure 3.2 displays the usual characteristics of good events. The hit pattern clearly shows the ring of hits expected for Čerenkov light. The time distribution of hits is fairly narrow for the hits in the ring, with those hits outside occurring later. The width of the time distribution depends upon where in the detector the event occurred. The hit pattern in electronics space is relatively random, showing that for neutrino events there are not correlated hits due effects in the electronics. The only correlations come from the non random mapping between real space and electronics space. Overall there are no striking features in this event suggesting that it is anything other than an event caused by Čerenkov light.

3.1.2 Flashers

Flashers were the first of the detector backgrounds to be noticed back during the airfill days. These events signalled that we would have a problem once the detector was up and running in a production mode and that a substantial effort on data cleaning would be required.

Flashers, as the name suggests, occur when one of the PMTs in the detector gives off a large pulse of light. This is believed to occur due to static discharge somewhere inside the dynode stack, although the precise mechanism is not known. The size of this pulse varies and the NHIT spectrum of flashers occurs throughout the solar neutrino spectrum, and stretches upwards to much higher values of NHIT, where they can interfere with the muon analysis.

Figure 3.3 shows a typical flasher event. This particular flasher falls into a sub-category of flashers known as golden flashers as it shows all the characteristics of flasher events. The main feature of the flasher shown in figure 3.3 is the early cluster on the one side of the detector, and the later light on the other. The flashing PMT is contained within the cluster and light propagates across the detector. The hit pattern of the light opposite the flashing PMT is very typical of flasher events. This ellipse of hits is probably due to shielding of light within the host PMT. The flashing PMT usually records a very high charge, usually so high that the measurement of the charge will rail the ADC[2]. The very high charge is responsible for the cluster in the flasher event, as adjacent tubes in electronics space fire due to crosstalk[3] from the flashing tube. The high charge measure-

[2]The charge measurement is railed when the charge value is the maximum possible value given the number of bits used to record it. In SNO this corresponds to a value of 4095 counts. Often a railed charge will record a value that is instead well below the pedestal for that channel, typically more than 250 counts below. This occurs when the charge deposited is even higher than that measured at 4095. When this occurs the charge measurement is said to be on the negative rail, a charge of 4095 is known as the positive rail. Hits in the negative rail are know to be due to very high charges as the high gain charge will often be in this category, while the low gain charge, is still below the positive rail.

[3]Crosstalk occurs when the signal due to a high charge deposited in one tube, causes nearby tubes to fire the discriminator due to pickup. Nearby in this case means nearby in electronics space, typically tubes on the same paddle card will fire, although in rare cases crosstalk will cross these boundaries, as the pickup can occur in the cables. The charge measured on a tube that fires due to crosstalk is consistent with zero, resulting in there being no valid PCA time for the tube, hence the use of ECA time in figure 3.3.

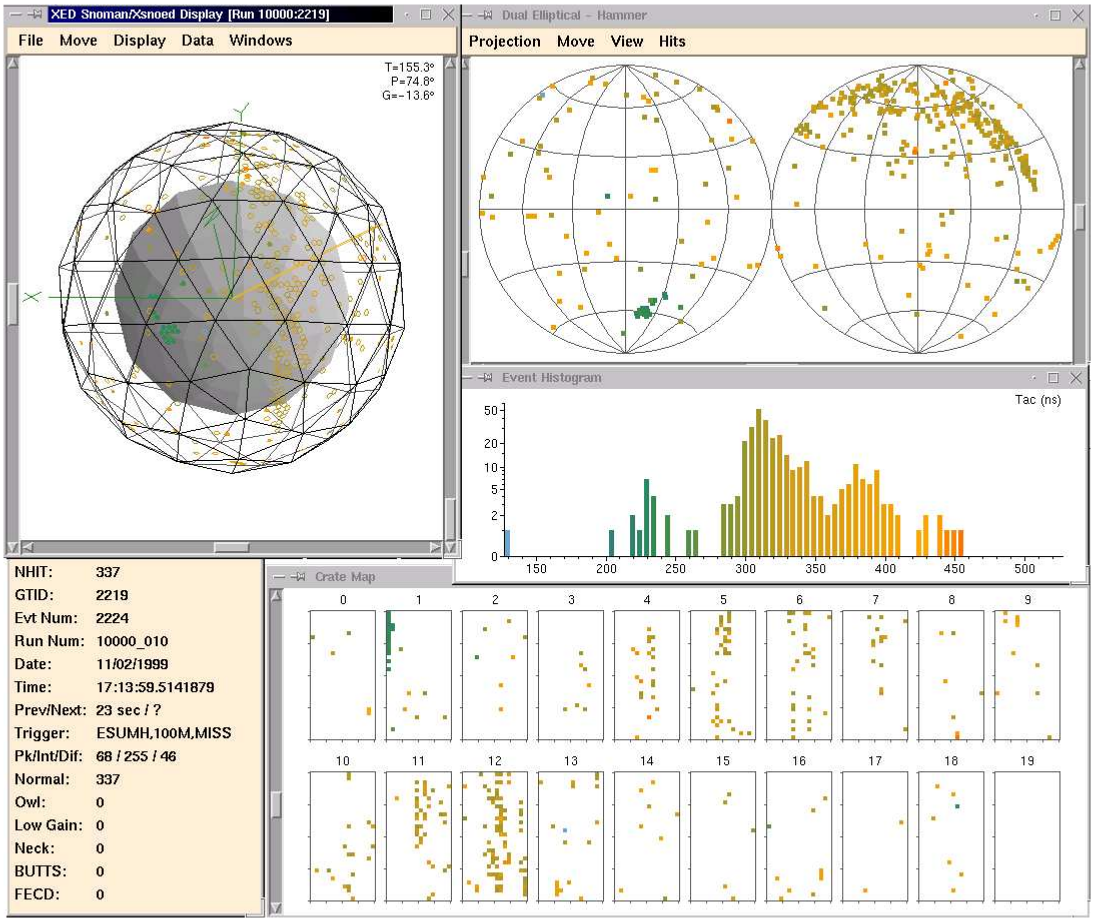

Figure 3.3: A typical 'golden' flasher event from SNO. In this display the ECA calibrated time is shown on the histogram and represented in the colours on the various projections.

ment is also fed into the ESUM signal meaning that flashers typically trigger the detector via the ESUM trigger, and the ESUM pulse is highly abnormal for flasher events.

One interesting characteristic of flashers is that they can be caused by seismic activity in the mine. A seismic event that propagates to SNO usually causes a burst of flashers, which can contain thousands of these events. This feature is very useful as it allows us to collect samples of flashers that are biased in a different way to samples based upon selection of individual events.

The flasher rate in the detector is very high compared to the expected solar neutrino rate. The rate is highly variable but is typically 50 per hour [59]. One interesting feature of the flasher rate is that it shows some conditioning to changes in the high voltage on the tubes. Figure 3.4 shows the flasher rate as a function of time around the time when the high voltage was increased. The conditioning effect seen on the flasher rate implies that flashers are indeed a high voltage phenomenon within the PMTs.

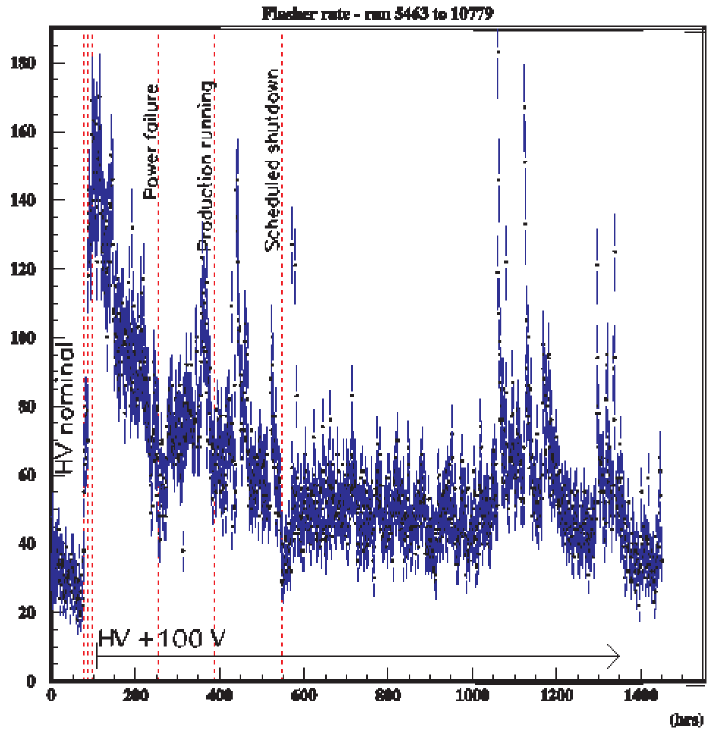

Figure 3.4: The flasher rate in the detector showing the effect of the increase in the high voltage, and of detector shutdowns. The occurrence of various changes to high voltage conditions are shown by the red dashed line. This figure is courtesy of [38].

3.1.3 Pickup

Another class of event seen during the airfill days are external pickup events, known generically as pickup. These are events due to external influence on the electronics and occur when, for example electronic equipment such as a drill, is used near the electronics. They can also occur, to a lesser extent, when there is any activity on the Deck, such as someone walking on to check a valve. This class of pickup event can also occur when the required operating environment is not met, such as the wrong temperature, humidity, or both.

A typical pickup event is shown in figure 3.5. Pickup events are very distinctive in electronics space as most of the tubes on a number of crates will fire. In extreme cases such as the drill mentioned previously most of the channels on every crate will fire. The major feature of all pickup events however is the measured charge on the hit PMTs. The waveforms that cause the pickup tend to be bipolar in nature, and occur at high frequency. This means that when the charge is measured over a long time period, the measured charge integrates to zero. The long integration time is achieved by the use of

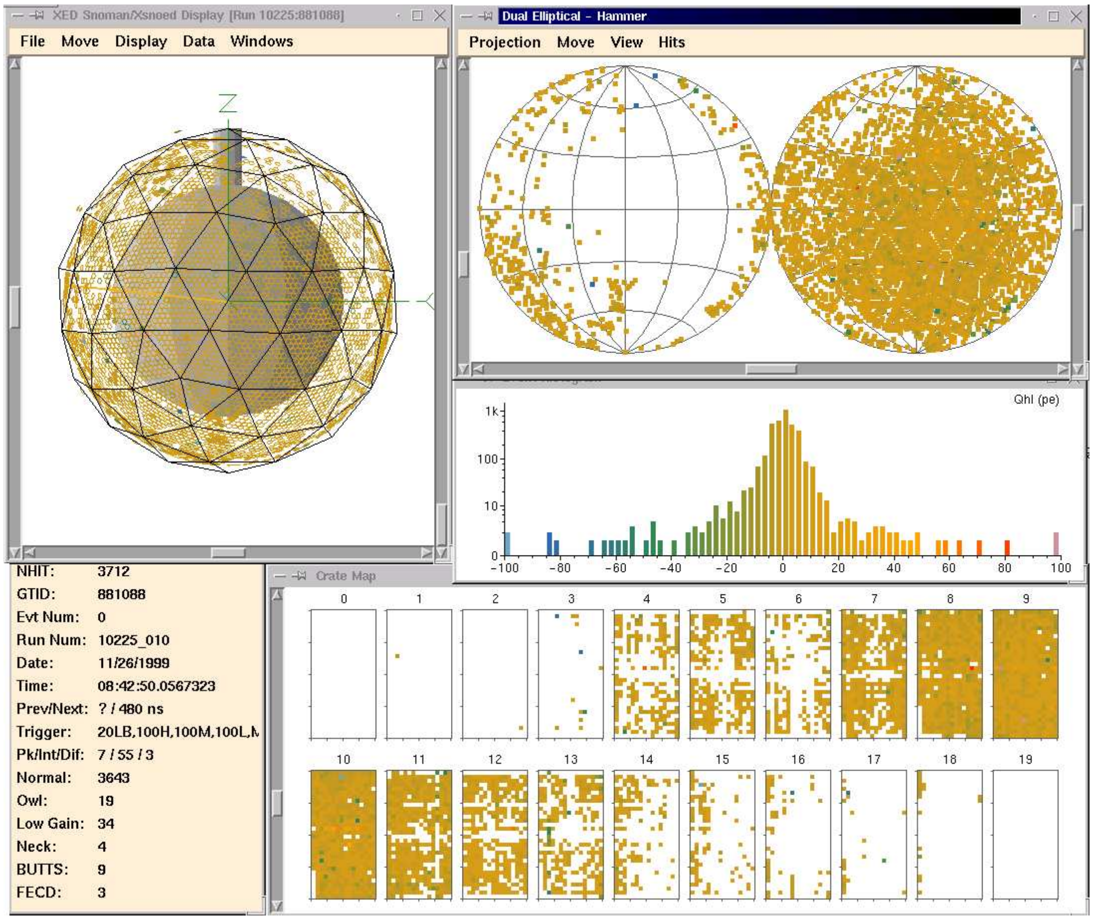

Figure 3.5: A typical pickup event. In this display pedestal subtracted QHL is shown on the histogram and represented in the colours on the various projections.

the QHL charge, which as shown in figure 3.5 is centred about zero counts.

3.1.4 Wet-end Breakdown (WEB)

Wet-end breakdown was first seen during the days of partial fill[4] and at that time there was concern that it could mean the end of SNO. The signal seen was that an entire paddle card (at least) would be railed, there would be pickup in both that and neighbouring crates, with light seen in the rest of the detector. These events occurred at high enough rate to prevent us from taking any useful data, and could in the long term be associated with tube death. Because of this the detector was turned off approximately half way through the partial fill period to prevent further damage to tubes and allow this problem to be investigated. Filling of the detector was completed in this time.

[4]Partial fill is the period of time while the water was being added to the detector.

The story of WEB during partial fill is a long story and could be the subject of a thesis chapter on its own. WEB was finally discovered to be breakdown in the wet-end high voltage connectors. These connectors contain air, which when placed into degassed water (as it was at the time) leeched out of the connector, leaving low-pressure gas under high voltage, an ideal condition for breakdown. As all testing of the connectors had been carried out in gassed water, the problem had been missed before partial fill began. A simple solution of re-gassing the light water with nitrogen was used and this effectively eliminated WEB from the detector.

Throughout the D_2O data set WEB is rarely seen in the detector. When it occurs there is a burst of WEB-like activity before the detector quietens down again. These bursts are associated with tube death. The WEB seen during these bursts appears to have slightly different characteristic to that seen during partial fill, suggesting that it comes from a different source. It is not too surprising that an effect that causes a PMT which is under high voltage to break, would have similar characteristics to WEB. Other than these bursts there does not appear to be a low level background rate of WEB events.

3.1.5 Sharkfins

Sharkfins are so called because of the appearance of the ESUM pulse in these events. Sharkfins are characterised by a large pulse in a single PMT. The large pulse causes a cluster of hits in adjacent channels by crosstalk just as was the case for flashers. Sharkfins actually look just like flashers, but without the associated light. This could mean that Sharkfins are a class of flasher where the light cannot escape the PMT, or they could be a different class of event altogether. They do however show a similar high voltage conditioning behaviour to flashers. Sharkfins are usually well below threshold for the NHIT trigger, but are present in the data stream as they trigger via the ESUM trigger.

Sharkfins occur at a high rate in the detector, triggering the detector at 2.3Hz [59]. Due to their high rate they can cause a sacrifice problem by pileup with Čerenkov events. This pileup is one of the major sources of sacrifice, as sharkfin characteristics are the same as some of those used to cut flasher events.

3.1.6 Bubblers

The remaining background classes were first seen in full D_2O running. The first and most dangerous of these, were the bubbler events. Bubbler events are due to light coming from the bubbler tubes in the acrylic vessel. These tubes were used to monitor the levels of both the H_2O and D_2O to ensure that they remain within safety levels. These are very important quantities for SNO, as a relatively small deviation from the norm can lead

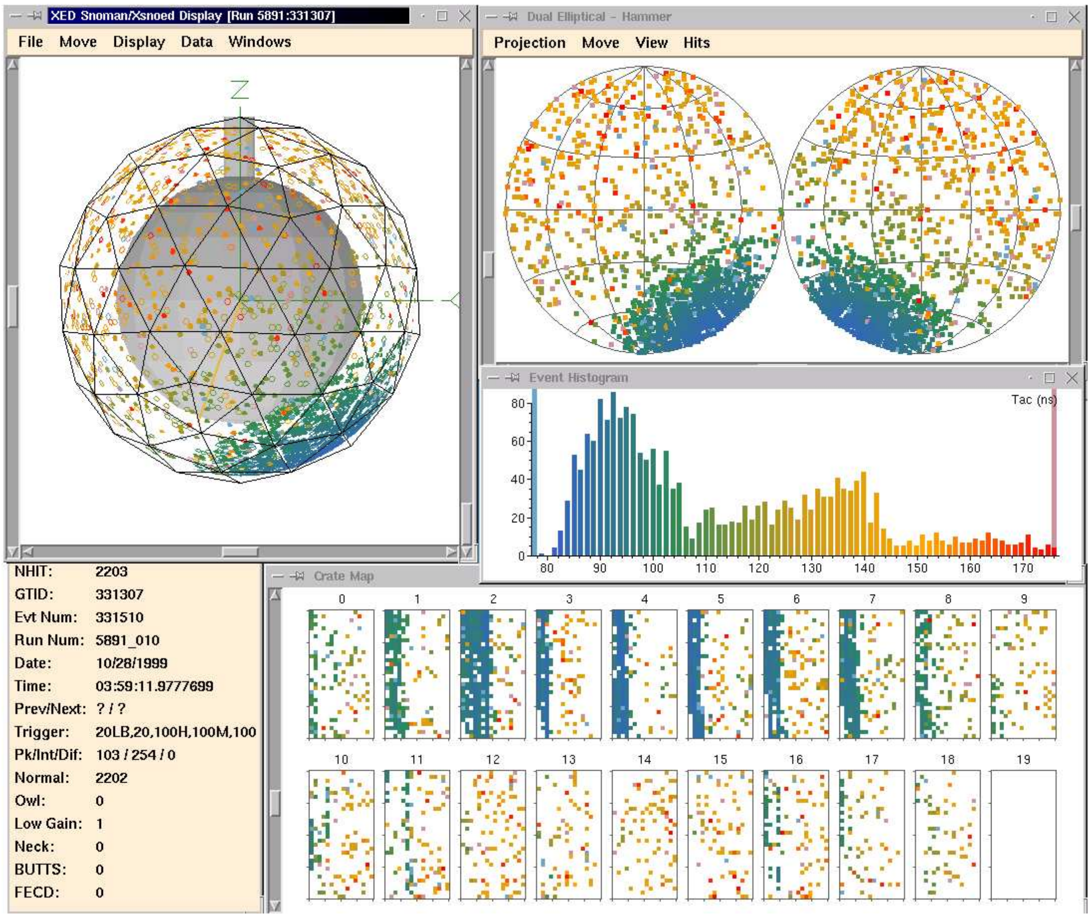

Figure 3.6: A typical bubbler event. In this display fully calibrated times are shown on the histogram and represented in the colours on the various projections.

to a pressure differential across the acrylic vessel that cannot be supported, causing the acrylic vessel to fracture, and a premature end to the experiment. It was found that gas water interface in the tubes would emit light. A typical bubbler event from one of the pipes is shown in figure 3.6. The conclusive evidence that these events come from the gas-water interface in the tubes was seen when the bubbler tubes were flooded. When this was carried out the source of the events was seen to follow the gas-water interface as the flooding progressed.

After this was discovered it was realised that an alternative approach to monitor the water levels was required. During the period from the beginning of the D_2O commissioning phase in May 1999, until the beginning of production running on the 2nd of November 1999, a system to monitor the water levels using ultrasound was developed and tested. With the use of this system the bubbler tubes are flooded throughout the neutrino data set, eliminating bubbler events at the source.

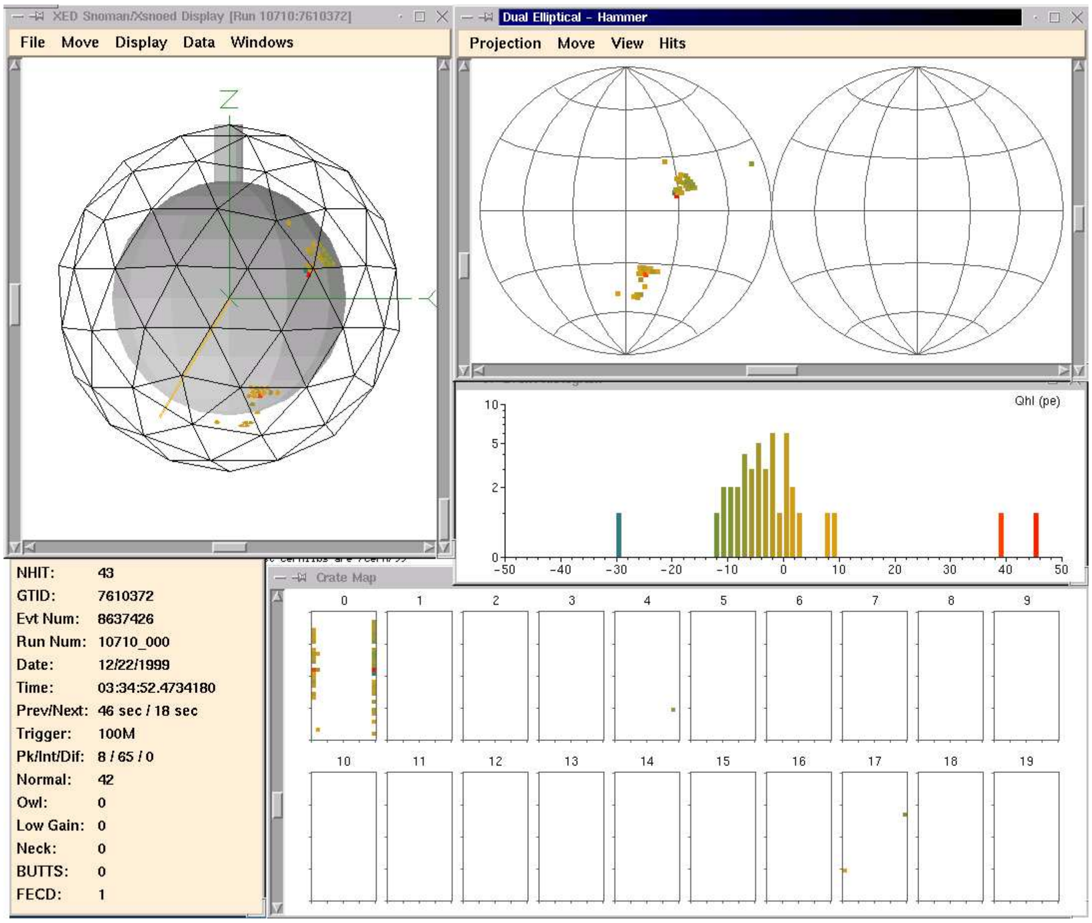

Figure 3.7: A typical ring of fire event. In this display pedestal subtracted QHL charges are shown on the histogram and represented in the colours on the various projections.

3.1.7 Rings and Slots of Fire

Rings of fire (RoF) are so called due to their appearance on an electronics map. A ring of fire is a pickup event internal to a crate, where the channels at the edge of the crate trigger the discriminator. Slots of fire (SoF) are closely related to rings of fire, except that the affected channels are on one or two cards inside a crate.

Figure 3.7 shows a typical ring of fire event from crate 0. This was the crate that had the worst problem with rings of fire. The distinctive hit pattern in electronics space is clearly shown in this event. As with external pickup events the measured QHL charge is compatible with zero.

The rate of these events is highly variable. They can occur in bursts which can require a run to be stopped, or they may not be seen from a crate for months only to start again. At the beginning of the production data set, rings of fire were controlled by the careful adjustment of the thresholds on the affected channels. However as the detector become better understood a permanent fix for these events was found. A backplane mod to the crate was designed and installed. This mod reduces the possibility for internal pickup in a crate and vastly reduces the number of rings of fire that occur. This means that rings of fire are more prevalent at the start of the data set than the end, but they are present throughout.

Rings and slots of fire are a worrying background because, although they are well defined they occur inside the range for study of solar neutrinos, indeed they have a very similar NHIT distribution to neutrons in pure D_2O. Furthermore, as these events are common, coincidences do occur with radioactive backgrounds When this occurs the characteristics used to remove the events are weakened meaning that this type of ring of fire event is the most likely to survive and appear in the data set after instrumental cuts have been applied.

3.1.8 Retriggers

Retriggers are events that follow large events. These events will follow both high NHIT events such as muons and high charge events such as flashers. Typically they follow the retrigger time (500ns) after the initial event, but there is a time distribution for these followers, see figure 6.5.

A typical retrigger following a muon is shown in figure 3.8. There are three distinctive features of retriggers visible in this event. The first is the flat time spectrum for the hits. Secondly most of the PMTs on channel seven are present in the event. Finally the PMTs hit hardest by the muon are preferentially present again in this event.

There are four causes of retriggers all of which can be present in an individual retrigger. These are:

- Large events such as muons contain large numbers of Čerenkov photons. Some of these photons can continue to bounce around the detector for a time longer than GT_valid[5]. If enough of these photons hit a PMT after this time then this will cause the detector to retrigger. Channel 7s are particularly liable to see this effect as they have a shorter reset time than other channels.

- Internal pickup can occur due to the large number of signals passing through the electronics. Channel 7s are again more liable to be affected due to the shorter reset time.

[5]GT_valid is the length of time for an event defined by the electronics. In SNO GT_valid is approximately 400ns.

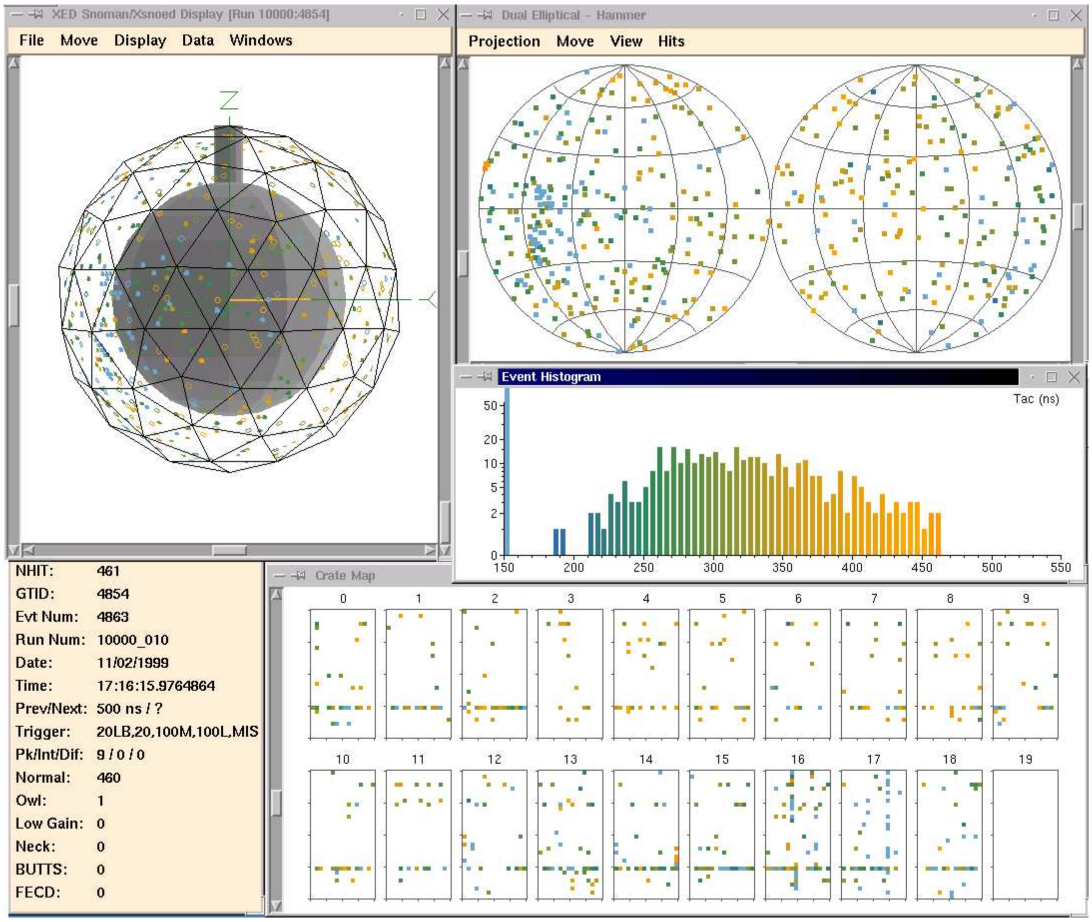

Figure 3.8: A typical retrigger event. In this display ECA times are shown on the histogram and represented in the colours on the various projections.

- PMT after-pulsing can occur. This happens after a high charge hit on a PMT. The increased number of electrons accelerated through the dynode stack, can liberate positive ions in the gas inside the PMT. These ions travel back to the photo-cathode where they can liberate more electrons causing the PMT to fire again. This can occur for many cycles depending upon how hard the PMT was hit initially. As the timescale for after-pulsing is longer than GT_valid, this can cause retriggers.

- Pickup can occur in the trigger system, as is seen to occur on the NHIT20, causing this trigger to fire. This creates events that occur 4.3μs after the initial event, that have the same NHIT spectrum as PulseGT events.

3.1.9 Neck Events

Along with flashers and rings of fire, neck events are amongst the most dangerous background to neutrino events present in the D_2O data. They are the subject of chapter 4.

3.1.10 Flat TAC Events

Yet another PMT based instrumental background exists, the flat TAC events. These events, as the name suggests, are characterised by the time distribution of the event being flat. Due to the trigger system, this is not strictly true and not all flat TAC events have completely flat TAC spectra. However the TAC spectra is always much broader than that expected from Čerenkov events. Flat TAC events occur in bursts, which when summed shows more light on one side of the detector than the other, with OWLs on the opposite side. This spatial distribution is only visible after summing a burst as individual events only show weak spatial correlations. A flat TAC event is shown in figure 3.9.

Flat TAC events occur when a region in a PMT, probably the base, emits d.c. light. Bursts of flat TAC events are always associated with tube death, so it is believed that the d.c. light is emitted by the tube as it dies. This explains the hit pattern seen from a summed flat TAC burst. The OWLs behind the tube see light coming through the translucent base of the PMT.

The NHIT of flat TAC events occurs within the region of interest so while they are not expected to produce good fits cuts have been developed to identify these events. These cuts help provide a good diagnostic for detector health as well as being essential for the data cleaning process.

3.1.11 Leslie Events

Leslie events[6] are named after the first person to investigate these events. These events are present at a much lower level than any of the other instrumental backgrounds, and as such were only discovered in studies of the data after the instrumental cuts were applied.

A Leslie event is shown in figure 3.10. These events all reconstruct to the acrylic vessel and appear to be isotropic from their fit position. They are present throughout the neutrino energy spectrum and at much higher NHIT values. The large range of NHIT values where these events exist, suggest that they are not due to radioactive decay, but are an instrumental background. How light is produced in these events is unknown.

[6] These events are also known as Queen's or Heeger events

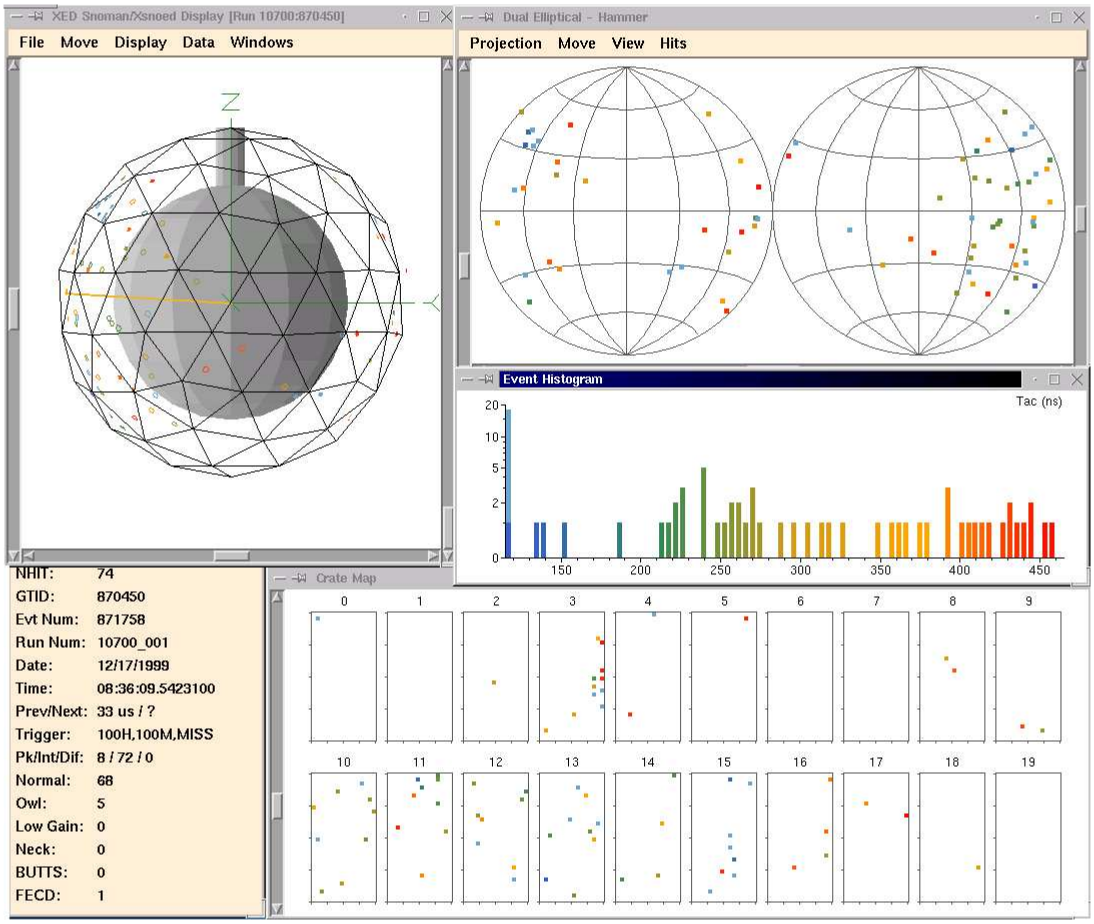

Figure 3.9: A typical flat TAC event. In this display ECA times are shown on the histogram and represented in the colours on the various projections.

3.1.12 Orphans

Orphans are not an instrumental background, but they are included in this list[7], as they are dealt with in the same way. As data comes out of the SNO electronics it is passed to a program called builder by the DAQ system. The builder combines PMT bundles with MTC data (see chapter 2). On rare occasions a PMT bundle will not be built with the rest of the corresponding event. This PMT bundle is called an orphan. Orphans are usually collected together to form orphanages of 80, or 10000 hits and are output to the data stream. On even rarer occasions it is possible for orphans to be built into the wrong event potentially resulting in an event with more than one hit on the same PMT. This only occurs for a particular class of orphans known as 00-orphans. These occur for events where the last eight bits of the GTID are zero, and are caused by a hardware design

[7] In the historically incorrect position

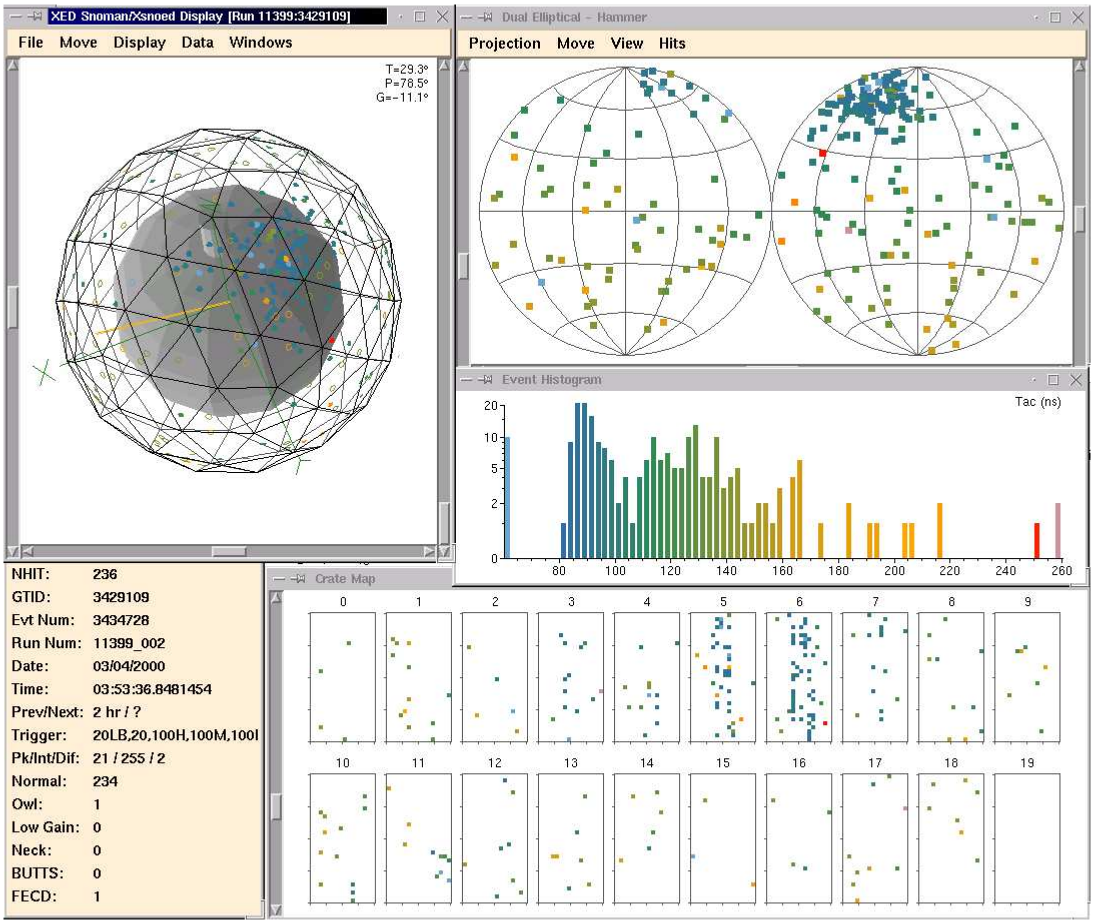

Figure 3.10: A typical Leslie event. In this display fully calibrated times are shown on the histogram and represented in the colours on the various projections.

problem where the CMOS chip is left in an uncertain state upon roll-over. 00-orphans will on occasion be built into the wrong 00-event.

Orphans can also be caused by the hardware faults, where the GTID on a particular channel can become stuck. Perhaps the most serious orphan problem however are orphans caused by periods of extremely high rate in the detector, which can overwhelm the builder. In this case all PMT bundles are output as orphans, the efficiency of the detector is essentially reduced to zero for this period of time. This causes complications for the measurement of livetime.

3.2 Background rejection techniques

There are two principle levels of cuts used in SNO to reject the large number of instrumental background events. These are the instrumental and the high level cuts. The instrumental cuts are the first level of analysis applied to raw data. They consist of a number of simple cuts designed to remove background events. Some of the cuts are targeted at a single class of background, while there are complementary cuts aimed at a broader range of events. Often a cut is designed that is targeted at a single class of background, only to find that it is also effective in the removal of other backgrounds. The instrumental cuts are applied once to the data, and the results are stored in a DAMN bank for use at later times. DAMN banks are described in appendix A.

The high level cuts differ from the instrumental cuts in a number of ways. The principal difference is that these cuts are designed to accept signal, opposed to the instrumental cuts which have been designed to reject background. Each high level cut requires a reconstructed vertex in order to calculate the cut parameter. The one exception to this is the reconstruction algorithm, which can also be considered to be a high level cut.

The high level cuts are used not only to remove instrumental backgrounds missed by the instrumental cuts, they will also remove background events that were not known when the instrumental cuts were designed, such as the Leslie events. The high level cuts also remove badly mis-reconstructed events.

The two techniques are used in combination to provide a data set that is very clean. Events must pass both levels of cut to be further analysed in the study of the solar neutrino signal.

3.2.1 Standard data sets

In order to construct and investigate the cuts a number of standard data sets were constructed. A number of these are data from calibration sources and the study of the effect of the cuts on this data is described in detail in chapter 5. There are however a number of data sets of instrumental backgrounds constructed as well.

- Flashers.

  There are two different sets of flashers. One is a golden flasher data set with events taken from runs 10000 - 10655. The selection process is to search for events with one cluster of 6 or more continuous hits in channel space. Events with more than one such cluster, or events containing more than 6 smaller (3-5 hit) clusters are rejected. Rings of fire are removed from the sample by requiring 20 or more hits on the crate not containing the cluster. The second flasher sample is a flasher burst from run 10714. The burst is 8.4s long and is unmodified except for the removal of orphans. This means it contains a number of non-flasher events, such as flasher retriggers and a small number of events from the background wall. Both these samples of flashers are biased, but with the use of both data sets the effect of these biases can be reduced.

- Neck events.

  There are two samples of neck events, taken before and after the cabling changes. Both these samples were derived from the hand scanning described in chapter 4.

- Flat TAC events.

  There is a burst of flat TAC events taken from run 10700. It is 2.9s long and like the flasher burst is unmodified except for the removal of orphans.

- Pickup.

  There are also two samples of pickup. The first is a set of external pickup due to activity on the deck. This set is derived from hand scanning a number of runs with known external pickup. The second is a set of internal pickup such as rings and slots of fire. This data set is derived via a simple algorithm requiring 14 or more hits on a card, and that the card comprises more than half the total hits in the event.

- Retriggers

  This is a data set of all events following within 2μs of NHIT>1000 events in run 11371.

These data sets were used in the design of particularly the instrumental cuts. They are not now used to measure the contamination of the signal data by background, as the statistics of the data sets are small, and not all known backgrounds are present in the data sets. A more sophisticated technique is used for this measurement and it is discussed in chapter 5.

3.2.2 Data blindness

In order to avoid statistical bias SNO implemented a data blindness scheme. This scheme is particularly important when designing and testing the cuts. Data taken from the start of the production data set on November 2^{nd} 1999, until the end of run 12168 on July 1^{st} 2000, were all available for analysis, but only 10% of data taken after this time were used.

The process for producing the blind data is known as data splitting. This was achieved by dividing the run into 30 equal pieces based upon the start and stop times of the run. The data sets were numbered in a pseudo random way based upon the run number. Three randomly chosen data sets: 3, 13 and 23 were then opened and added to the open data set.

This procedure left 31% of the data unanalysed while all the analysis tools were developed. Only after the entire analysis had been fixed was the blind data analysed. No statistically significant difference was seen between this data and the open data. While in principle 69% of the data was available for the setting up of the cuts, the instrumental cuts were setup using only a tiny fraction of this data. This is because they had been fixed by March 2000. Even by this point, only a small fraction of the available data had been used in cut design. The only changes made to the cuts after this time were on the basis of results from calibration data[8]. The high level cuts however were setup using the entire open data set.

3.3 The instrumental cuts

Given the large number of instrumental backgrounds present in SNO data, there are correspondingly a large number of instrumental cuts, sixteen in all, that are used to reduce the background content of the data. Each of the cuts has been designed with three criteria in mind. Two of these are of course the background rejection power and the sacrifice of the cut. A target of less than 1% for both contamination and sacrifice was set. The bias in the sacrifice must also be small. This is important as it avoids the need to model sacrifice bias in down stream analyses. The third criteria used when designing the instrumental cuts is that they should be robust. In this context this means that the cut should be stable against badly calibrated or faulty channels. Thus a cut based on a single channel should have some protection against a breakage in a channel causing excessive sacrifice in events where that channel is hit.

The instrumental cuts remove events based upon five different properties of events. The cuts use these properties either singly or in combination to remove background. The five event properties used are:

- Intra-event timing

  The time distribution of the hit PMTs in some background events, for example flat TAC events, is very distinctive and different to that expected from Čerenkov events. This occurs because Čerenkov light emission, from the time scales important in SNO, is essentially an instantaneous process. This is however not the case for some of the classes of instrumental background. The difference is in some cases big enough so that the difference in the time distribution is great enough to be identified without reconstruction of the event.

[8]The only example of this is the change made to QvT on the basis of the ^{8}Li sacrifice analysis.

The time of specific PMTs in a event can also be important for some classes of background. A PMT identified by one of the other criteria used by the instrumental cuts may be expected to have a specific relation in time to the remainder of the PMTs in an event. This is important in the removal of both flashers and neck events, where light is emitted by or near an identifiable PMT and must then travel across the detector to be detected by any other tubes.

- Inter-event timing

  Many of the different classes of instrumental background are known to occur in bursts. This is not the case for solar neutrino events. Thus by setting up a burst cut with the relevant threshold and timescale it is possible to remove large numbers of instrumental background events.

  Not all signals of interest in SNO are time uncorrelated, such as signals from anti-neutrinos, or events from a supernova. For the analysis of events from these phenomena it is important to disable any burst cuts working on the relevant timescale before the analysis is carried out. The burst cuts are discussed in chapter 6 when the livetime of the experiment is discussed.

- Event Charge

  If the time of the PMT hits can be important for the removal of background events, the charge of the hits is even more important. In particular flasher events, which comprise the most numerous instrumental background, deposit large amounts of charge in the PMT which flashed. Unfortunately just measuring the charge on a PMT and searching for PMTs with a very high charge measurement is not enough. This technique falls foul of the requirement that the instrumental cuts are robust. However the same requirement in combination with another event criteria can successfully be used in the removal of these events.

  As well as the measurement of the charge on a single PMT the total charge measured in an event also forms a useful cut. The charge deposited on a flashing tube is much bigger that the total charge deposited in the detector in a typical solar neutrino event of similar NHIT, so this value can be used to remove these events. The same value can be used to remove pickup events where the total deposited charge is much less than that expected from a typical solar neutrino event.

- Event geometry

  Instrumental backgrounds, as well as having unusual signals in the hit PMTs, will also cause unusual combinations of tubes to fire in the same event. The geometry of instrumental background shows correlations that are rarely expected in Čerenkov events. These correlations are particularly apparent in electronics space due to crosstalk. Examining event geometry in electronics space reveals these correlations and allows events to be cut. Geometry cuts do not need to take place in electronics space. Due to the correlations between electronics space and real space the correlations in real space can be used.

- The presence of veto tubes

  Events due to light from outside the PSUP, such as neck events, will also cause light in regions of the detector where light originating inside the PSUP would not be expected to reach. There are veto tubes on the outside of the PSUP, and near the neck. Hits on these tubes can be used to remove events. This hit information can either be used individually, or in combination with other criteria used to remove instrumental backgrounds. The cuts using veto tubes are all aimed at least partially at neck events and are discussed in chapter 4.

A number of the instrumental cuts are discussed in detail elsewhere in this thesis. Descriptions of the neck, OWL and OWLEHi cuts can be found in chapter 4 where neck events are discussed. The three burst cuts used[9] are discussed in chapter 6, along with the impact these cuts have on the livetime of the experiment. The remaining cuts are described in the rest of this section.

3.3.1 The junk cut

The junk cut[60] is designed to be a catch all for artifical events that are introduced into the data stream either by accident or design. The most important action of the junk cut is that it removes orphans and orphanages. However there is more to the junk cut than this simple process. The cut will also search for events that have picked up orphaned PMT bundles. The process for doing this is to remove events where the same PMTs appears to have been hit more than once in the same event. It is important to remove these events from the data stream as soon as possible in the analysis chain because, despite their rarity, these events can cause software problems. The junk cut also removes pedestal events. These events can occur in neutrino data. The events are not useful for analysis, except in specific circumstances and are thus removed by the junk cut.

[9] The retrigger, burst and NHIT burst cuts

3.3.2 The analogue measurement board cut

The analogue measurement board or AMB cut [61],[62], is probably the single most powerful cut in the suite of data cleaning cuts. The analogue measurement board is designed to provide diagnostic measurements of any of the analogue trigger sums. When it is used to measure the ESUMHi signal it can be used to measure the total charge deposited in an event independently of channel properties. This avoids problems with errors in the calibration of individual channels. The ESUMHi signal is a linear sum of all PMT pulses in the detector, shaped to a Gaussian profile, of 120ns width. The AMB provides three 8-bit measurements of the ESUM signal, the peak, integral and derivative in a time window around the trigger crossing. The time window is big enough to include most hits in a normal event, although for some backgrounds this is not always the case.

Only the integral and the peak are used to remove events. Each is a linear function of the number of hits in the event, and the ratio of the measured AMB value to NHIT is used in the cut. The AMB value cannot be used directly as it must be corrected for the pedestal (the value at zero hits), measured by extrapolation of calibration data to zero hits. The deviation of the event from the mean value of the corrected AMB/NHITs ratio is the cut parameter. As the spread of the AMB/NHITs distribution is function of NHITs then this must be corrected for in the cut. The σ of the distribution is measured as a function of NHITs, and events are cut if they are more than 3.7σ from the mean value. The AMB cut requires calibration to determine these parameters, this is carried out using ^{16}N data.

Figure 3.11 shows the corrected AMB/NHITs integral distribution for ^{16}N data, the flasher burst and rings of fire. With the choice of the cut at 3.7σ the sacrifice of the cut is small, but is still the largest sacrifice of any of the instrumental cuts. The cut is very good at removing flashers from the data set, and has some power to remove pickup events as well. The fraction of flashers removed by the cut is large as they tend to rail (i.e. record the maximum possible value of) the AMB measurement. This appears in the plot as the smooth curve followed by the majority of the flasher events. An asymmetric cut could improve the rejection of pickup, while controlling the sacrifice.

3.3.3 The QCluster cut

As flashers (and a number of other event classes) contain extremely high, usually railed charges, it is natural to design a cut that looks for the occurrence of this high charge. However it turns out that the charge cut fails the robustness criterion in a quite spectacular fashion, such that when the ^{16}N data was taken as the end of 1999, the sacrifice of the charge cut was 10%. Likewise, due to the clusters seen around the high charge tubes,

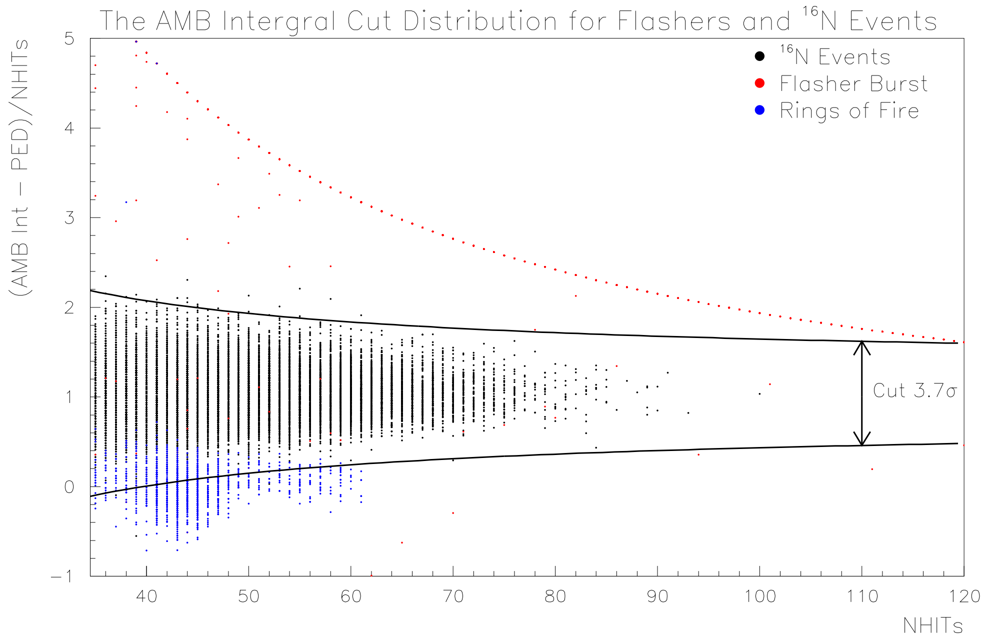

Figure 3.11: The AMB integral distribution and cut values for ^{16}N , flashers and rings of fire. The distribution of the signals is discussed in the text.

it is also natural to use the clusters to remove events [63]. This cut too falls foul of the requirements for data cleaning. The sacrifice bias in both radius and energy are too large to effectively use a cluster cut.

It is from the ashes of these two cuts that the QCluster cut [64] was born. The QCluster cut searches for events with a cluster containing a high charge tube. This combination gives robustness to the charge cut, and reduces the sacrifice bias of the cluster cut. The cut succeeds in this to such an extent that it has been made more aggressive in both charge and cluster requirements than the charge or cluster cuts alone.

The cluster cut was designed to search for clusters on a single paddle-card. As cross-talk became better understood, it was realised that this requirement would reduce the efficiency of any cluster finding algorithm. For this reason the cluster search algorithm was updated for the QCluster cut. A sliding window is now used that scans though electronics space searching for clusters of adjacent tubes irrespective of boundaries in the electronics, as pickup can occur in the cables. The QCluster cut searches in this way for clusters of four or more hits in a window five channels wide. This allows a single lost hit in a cluster to be ignored, irrespective of why the hit was lost.

Assuming one or more clusters have been found, the cluster still must contain a hit with an anomalously high charge. The charge distribution of hits from ^{16}N data taken

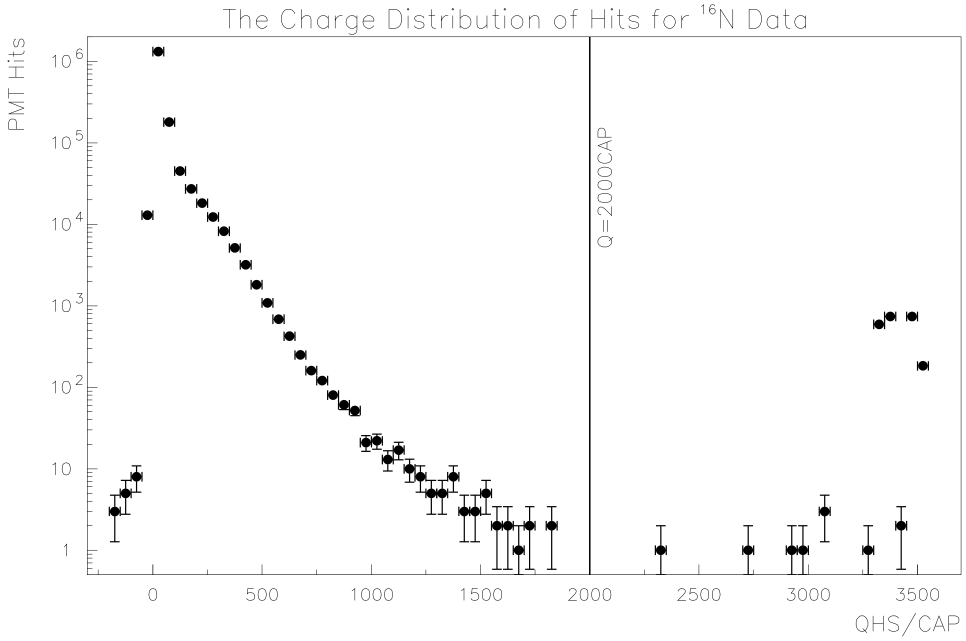

Figure 3.12: The charge distribution using QHS for ^{16}N data taken in the centre of the detector, in counts above pedestal (CAP).

in the centre of the detector is shown in figure 3.12. The hits at extremely high charge are due to the bad ADC present in the ^{16}N data. This is the problem that caused the charge cut sacrifice to be so high. The combination with the cluster requirement vastly reduces the sensitivity of the QCluster cut to these effects. In general the plot shows that the charge distribution does not extend significantly above 2000 counts above pedestal (CAP), with the exception of the hits on the bad channels. The choice for the charge cut is therefore made at 2000 CAP. For the low gain charge this requirement is lowered to 300 CAP, to allow for the reduced gain. Any hits where the charge is on the negative rail are also considered bad.

Using the requirement that a bad charge must have an associated cluster the QCluster cut allows the instrumental backgrounds, to be removed in a much more robust way. The QCluster cut also removes sharkfin events, as they also fit the criterion of the QCluster cut. These events by their high rate actually account for a significant part of the sacrifice of the QCluster cut, by coincidence with real events. This fact must be considered when monitoring the cut.

3.3.4 The flasher geometry cut

The QCluster cut demonstrated a method to recover the use of the cluster cut, avoiding the problem of the large bias in the sacrifice using that cut. However the QCluster cut by the use of the extra criterion of the charge, is not as effective as the cluster cut alone at removing flashers. This is because it is not unknown for the flashing tube, the one with the high charge, to be missing from the event. This can occur as the tube can be orphaned, or cut off from the start of the event. The rest of the cluster can however be processed normally and appear in the event. The flasher geometry or FGC cut [65] provides another modification to the cluster cut to reduce the sacrifice, while doing a good job at removing flashers where there is a problem in signal processing of the flashing tube.

The cut works, by searching for a cluster of 4 or more hits, either in real or electronics space, in the same way as for the QCluster cut. The real space cut defines a cluster as more than 4 hits within a 1m radius. For each cluster in an event the average distance in real space to the remainder of the tubes in the event is calculated. For flashers where the light originates in the cluster and then propagates to the other side of the detector, this is expected to be large. In the case of Čerenkov events, on the other hand, this distance would be expected to be small as clusters should only appear when the hit density is high. By placing a cut on this distance the FGC cut can remove flashers while maintaining a low sacrifice.

The FGC distance parameter is shown in figure 3.13. In this figure the number of events in each data set are equal before the application of the cluster finding algorithm. The figure shows that with the cut set at 1200cm the sacrifice even at 5m is small, showing that the radial bias has been reduced. The sacrifice at 740cm is higher partly due to the increased cluster probability and partly due to the concentrator cut-off causing clusters that are separated from the other hits in the event. The FGC cut is so good at reducing the sacrifice of the cluster cut that for data within the acrylic vessel the principal cause of sacrifice is the occurrence of a separate cluster, such as a sharkfin by coincidence, that is well separated from the rest of the hits within the event. The FGC cut is very good at removing flasher events, as shown from the distance distribution of both flasher data sets.

3.3.5 The QvT cut

Just as the FGC cut rescues the cluster cut in a charge independent way, the QvT cut [66],[67] does the same for the charge cut in a cluster independent way. In a flasher event the light is emitted by a single tube that then records a very high charge measurement. The light then propagates across the detector and is detected on the far side of the detec-

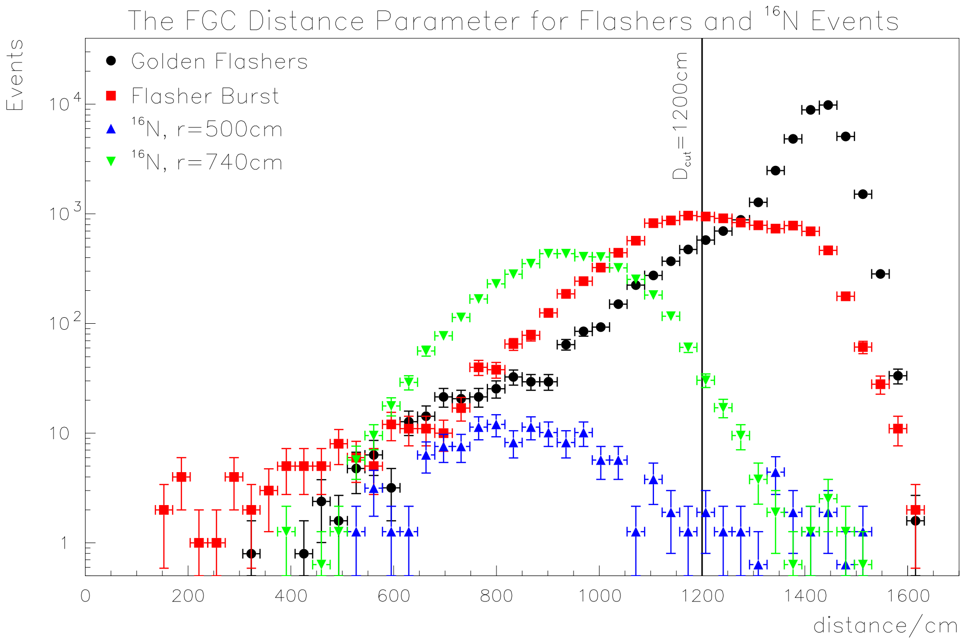

Figure 3.13: The FGC distance parameter for ^{16}N and flasher event. The plot is normalised such that there are the same number of events before the cluster search. Two ^{16}N source positions are used, these are 500 and 740cm from the centre of the detector respectively.

tor. This means that the high charge tube must occur early in the event, the requirement of the QvT cut.

The QvT cut operates by creating a time and charge box into which the highest charge tube in an event must fall for the event to fail the cut. In order to gain some robustness against bad calibrations only the highest charge tube is considered. To give the cut stability against the event moving within the time window, for example due to event triggering, the tube time is compared to the median hit time in the event. The charge cut is carried out by comparing the mean charge in the event. This means that events with naturally high charges, such as muons are less likely to fail as the average tube charge will be higher.

The distribution of the highest charge tubes, after correction for the median time and the average charge, for both ^{16}N and flasher data are shown in figure 3.14. In the plot tubes on the negative rail appear without pedestal subtraction, and are shifted to the high charge area of the plot. The ^{16}N data shown was taken during the period with the bad board, on which all hits appear with railed charges. The plot demonstrates the reduced sacrifice of the QvT cut compared to a charge cut. The bad hit must appear in

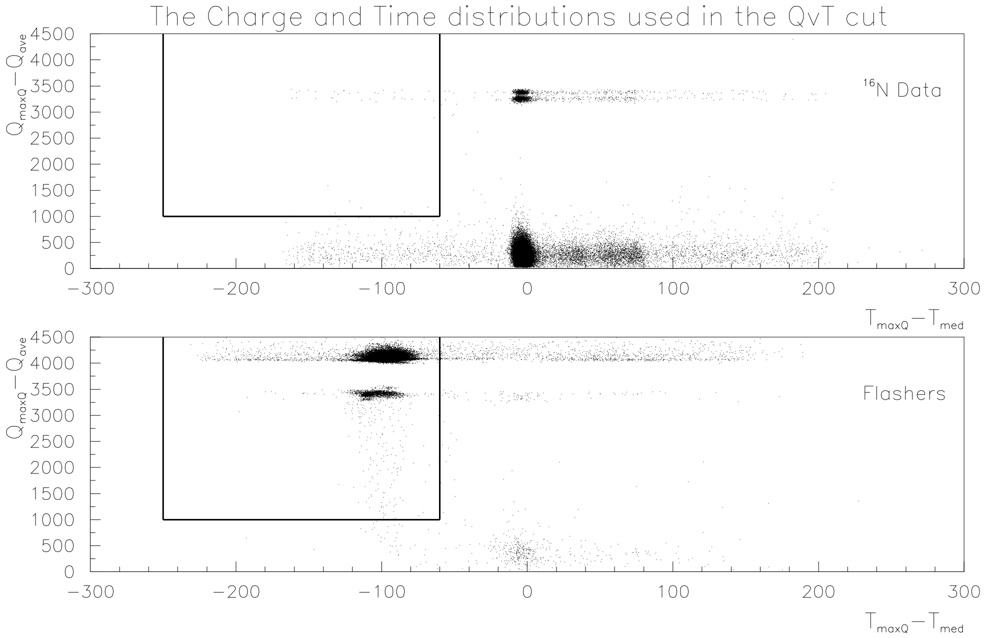

Figure 3.14: The corrected high gain charge and time for the highest charge tube in flashers and ^{16}N data. The event will be cut if the high charge tube falls in the box.

a small time window, and will generally only appear in this window as a noise hit; direct, or reflected light cannot appear at these early times. The majority of flashers are also removed, although this fact is somewhat hidden in the scatter plot by the clustering at early times on the negative rail. Also demonstrated is the extra rejection power available in the cut, made possible by the reduced charge threshold allowed by the time cut.

The QvT cut is the least robust cut amongst the instrumental cuts. This fact was emphasised when ^{8}Li was taken. A single channel would, some fraction of the time, appear railed in both charge and time. This lead to an unacceptably large bias in the NHIT sacrifice. The QvT cut has since been fixed to require a valid time measurement on the high charge tube. This fix worked and the ^{8}Li sacrifice is now acceptable. However since the QvT cut does rely on the calibrated measurements on a single tube, it is the least robust instrumental cut in use. For this reason monitoring the sacrifice of the cut is more important than for the other cuts.

3.3.6 The ESUM cut

Like a number of the cuts in the suite of instrumental cuts, the ESUM cut is backup against the failure of other cuts. The ESUM cut is based upon the way that the event triggers the detector. It will remove events if the ESUMLo trigger has fired (with or without the ESUMHi), but no other trigger is observed in the event. Unfortunately requiring no missed trigger bit vastly reduces the rejection power of the cut. The ESUM cut is however very good at removing sharkfins. For the same reasons that reduce the power of the cut, the sacrifice of the cut is very small, so small that the cut has never been observed to remove any tagged event from a calibration source.

3.3.7 The fitterless timespread cut

So far all the methods described that remove flashers have done so by the use, either directly or indirectly, of the flashing tube. One concern with a background as large as flashers is the question: do these events occur without the the appearance of the flashing tube or the associated pickup in the event (a so called blind flasher)? This could occur either due to some pathological case of flashers, or due to human error while configuring the detector. If the high voltage remains on a PMT while the sequencer[10] is disabled, then the tube has the potential to flash without the flashing tube appearing in the event.

The fitterless timespread, or FTS cut [68],[69] is a cut that removes flashers not based upon the flashing tube, but upon the other hits, thus having the possibility to cut blind flashers as well. The cut is based upon the median time difference between nearby tubes. The algorithm finds all valid pairs of tubes in an event. To obtain a valid pair, the tubes must be within 3m of each other. To reduce the sacrifice and bias of the cut, and to attempt to examine in-time light only, the time difference must be less than 25ns. This avoids the situation where noise and reflected light can bias the distribution to high values. The cut has one final restriction, in that it requires at least 15 valid pairs to calculate the median.

Figure 3.15 shows the distribution of the median value for flashers and ^{16}N events. The plot shows that with the cut placed at 6.8ns, approximately 50% of flashers are rejected. As pickup hits in the cluster will in general have no PCA time, the cut operates solely on the hits across the detector from the flasher. Thus assuming that this distribution is no different for blind flashers then, 50% of these events will be removed as well.

The FTS distribution for ^{16}N events, also shown in the plot, shows the most serious problem with the FTS cut. While sacrifice close to the centre of the detector is tiny, the bias on the sacrifice of the FTS cut is rather large. The bias is the reason that the FTS cut cannot be made more aggressive than it is. The bias will be discussed in more detail in chapter 5.

[10] The sequencer enables electronic readout of its associated PMT.

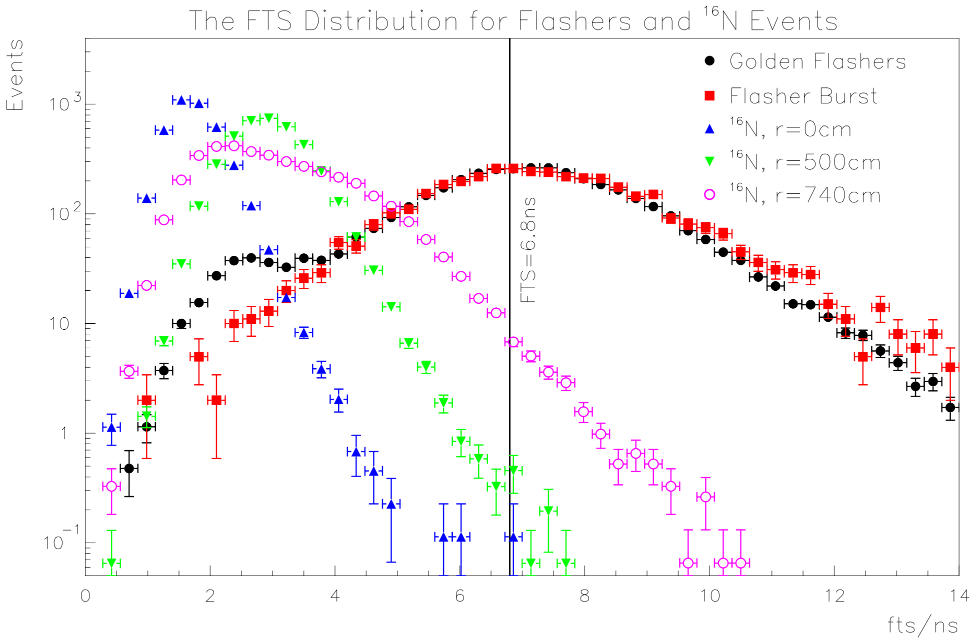

Figure 3.15: The FTS distribution for ^{16}N and flasher events.

3.3.8 The q/NHIT cut

The q/NHIT cut[70] is a cut aimed at removing the various forms of pickup seen in SNO data. These include ring/slot of fire events and external pickup events. Retrigger events which contain a component due to pickup are also targeted by the q/NHIT cut.

As discussed above, the charge measured over the long integrate time should be consistent with zero for pickup events. Whenever a Čerenkov event occurs in SNO nearly all PMTs that fire the discriminator are due to single photo-electrons (p.e.). The charge measured on the PMTs should therefore be consistent with this fact. Thus the mean value of the QHL charge should show good discrimination between Čerenkov and pickup events.

At the time of the construction of the cut, the gain of each PMT was not available. However the average gain of each PMT had been calculated. In order, therefore, to use more physically meaningful charges the pedestal subtracted charge is converted to a p.e. measurement using the average gain. A further correction is applied for those tubes with blown terminator resistors[11].

Robustness of the cuts is one of the requirements for the instrumental cuts. In order

[11] Terminator resistors were installed to remove reflections from the PMT cables at the expense of reduced PMT gain. Some of these resistors were blown during the WEB activity seen during partial fill.

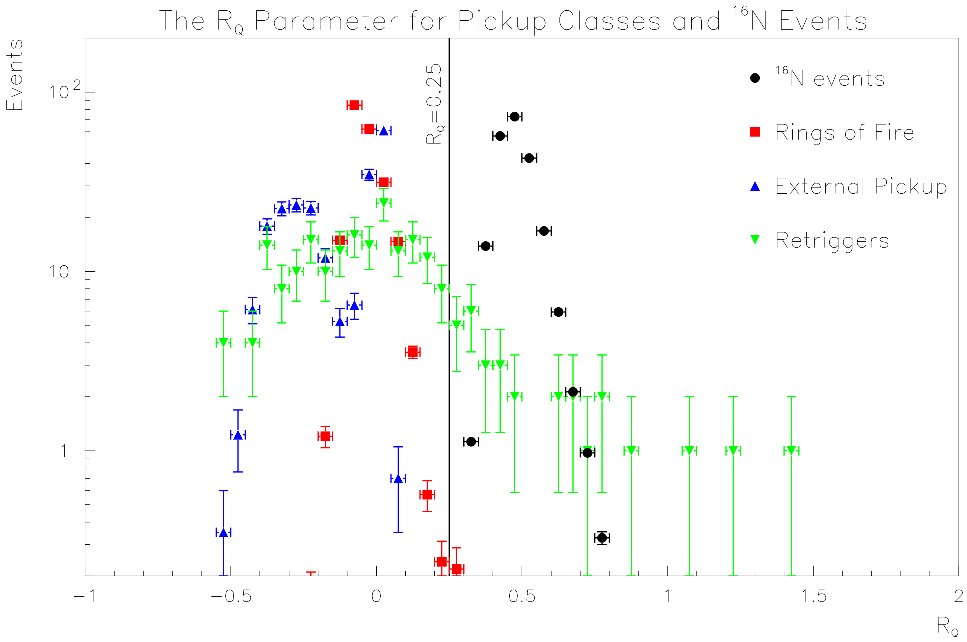

Figure 3.16: The R_Q distribution for ^{16}N events and various classes of pickup.

to avoid spoiling of the average charge measurement by bad ADCs or bad ECA constants, the mean charge is not used to form the cut parameter. Instead the cut parameter, R_Q, a modified version of the mean is used. Firstly non-physical charges, where the measured value is more than 100 counts below pedestal, are removed. Then the 10% of the tubes with the highest charge are also removed from the calculation. This protects the cut against channels that have anomalously high charge due to ADC or ECA problems. It also helps make the cut more effective at removing events where the pickup occurs in coincidence with a low energy Čerenkov event, as these hits are preferentially removed from the calculation of R_Q. Thus R_Q is the mean value of the QHL charge after corrections for tube gain and the removal of invalid channels.

The R_Q distribution for the various classes of pickup and ^{16}N events is shown in figure 3.16. As shown in this figure, with the cut value chosen at $R_Q = 0.25$ the sacrifice of the cut is minimal. The cut has excellent discrimination against external pickup events, which is important as this is the primary cut for these events. The discrimination against ring of fire events is also extremely good. Finally the cut is also better at removing retriggers than would have been expected. This is because the removal of the highest charge events selects the hits that are due to pickup in the determination of R_Q.

Overall the q/NHIT cut does a good job at removing all types of pickup event. Given the extremely low sacrifice of the cut, at its current setting, the rejection power could be easily increased by making the cut more aggressive. However how much more aggressive is not certain due to the the steeply rising R_Q spectrum for Čerenkov events just above the cut threshold. The cut could also be improved by the use of individual tube gains that were unfortunately not available at the time this cut was written.

3.3.9 The crate isotropy cut

The crate isotropy cut is a cut designed to remove ring (and slot) of fire events. This is achieved by the use of an algorithm that finds events with the hit pattern observed in electronics space. Rings of fire occur due to pickup on a single crate of the electronics. For this reason the hits due to the ring of fire occur only on that crate. In general there will be other hits in the event, due to noise or coincidence with a radio-active decay. For this reason the algorithm finds the crate with the largest fraction of hits and calculates the fraction of the total hits on that crate.

Once the crate with the highest fraction has been identified, a search for the hit pattern expected for rings and slots of fire can be carried out. In general a slot of fire event will have the majority of the hits on a single card, or two adjacent ones. A ring of fire looks slightly different as the majority of the hits are confined to cards 0 and 15, the edges of the crate. Thus in order to find these two classes of events the maximum card fraction is found. The card fraction is the ratio of the number of hits on two adjacent cards to the number of hits on the crate. In order to preserve sensitivity to ring of fire events, cards 0 and 15 are considered to be adjacent cards.

In order to construct the cut the crate and card fractions are considered in combination. The distribution of crate and card fractions for ring/slot of fire events and ^{16}N events from three different positions are shown in figure 3.17. The structure seen in these scatter plots is due to both fractions arising from the division of two relatively small integers. The rings/slots of fire only occupy a small region of the full parameter space. For ^{16}N events the card fraction is relatively evenly distributed, but the crate fraction distribution for events is a strong function of the position of the source. Only as the source gets close to the PSUP are crate fractions observed comparable to those seen in ring/slot of fire events.

The cut is made in a simple way using the combination of crate and card fractions. A box is constructed in crate, card fraction space and events inside this box are removed as ring of fire events. The box is shown in figure 3.17, and corresponds to a cut at crate fraction greater than 0.7 and card fraction greater than 0.8. A more effective cut could be constructed by the use of a more complicated contour in the crate, card fraction plane.

Because of the way that events must start close to the PSUP in order to have a crate fraction that is big enough for the event to be removed, the crate isotropy cut does not

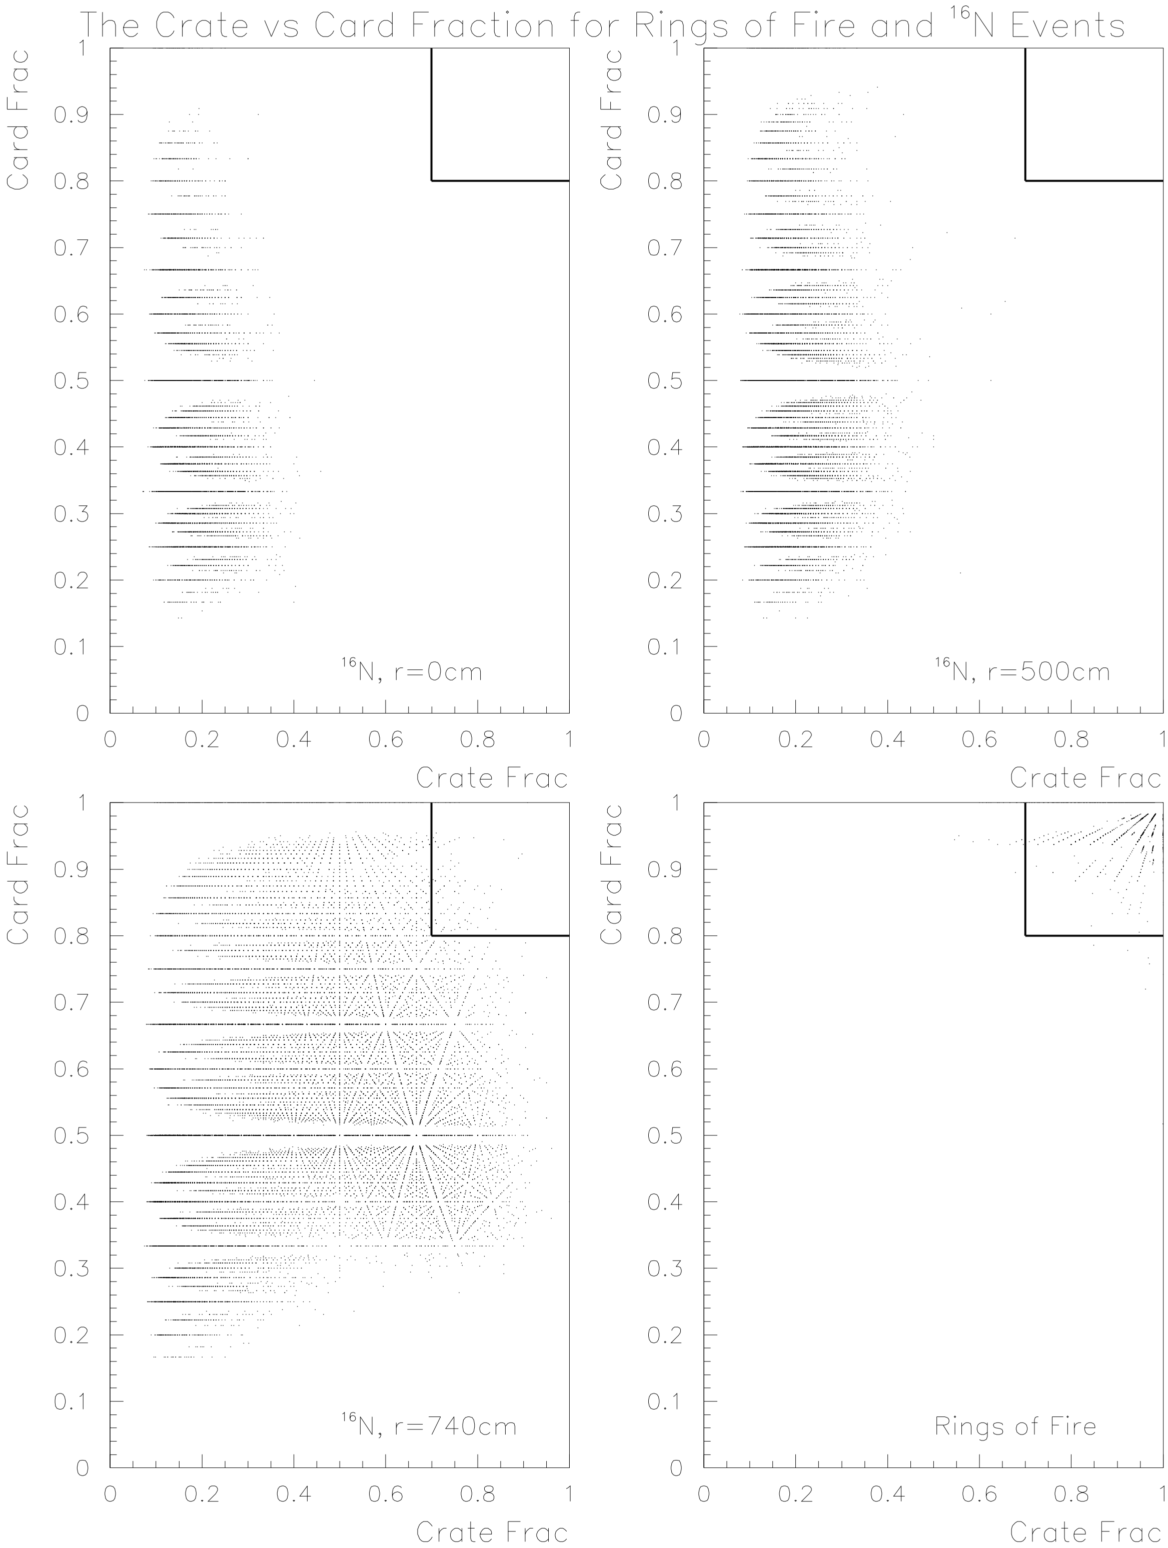

Figure 3.17: The distribution of crate fraction compared to card fraction for ring/slot of fire and ^{16}N events, showing the position of the cut box.

have a significant sacrifice for events inside the D_2O. As events get closer the the PSUP the crate isotropy sacrifice grows, in a complicated position dependent way, due to the mapping between real space and crate space.

The weakness of the crate isotropy cut for the removal of ring/slot of fire events, is also due to the crate fraction cut. When these events occur in coincidence with a low energy radioactive decay, they can pass the cut due to the hits from the decay lowering the crate fraction below the cut value. This is an unfortunate correlation with the way that the same events can fail the q/NHIT cut. These events are not expected to survive the high level cuts which examine event geometry due to the unusual hit pattern on the PSUP of ring/slot of fire.

3.3.10 The in-time channel cut

The in-time channel (ITC) cut [71] is aimed at the removal of flat TAC events, independent of the burst cuts. This is required as, while flat TAC events are only observed in bursts, they could also occur at a low level in data outside of bursts.

For Čerenkov events, even including the transit time of light across detector, the majority of the hits in an event occur in a relatively small fraction of the available time (GT_valid). This is not the case for flat TAC events, as the hits are much more uniformly distributed throughout GT_valid.

The ITC cut works by finding the maximum number of hits inside a sliding time window, and calculating the ratio of this value to the total number of hits. The size of the time window is chosen to be 93ns, the same as the NHIT trigger coincidence time. It is found that the rejection power of the cut is not a strong function of the size of this time window, and that 93ns sits comfortably within the stable region.

With this choice for the size of the time window, the ITC ratio is shown in figure 3.18. The non-statistical structure seen in the plot is because the ratio comes from the division of two relatively small integers. The ITC ratio is not a strong function of the position of the ^{16}N source which, suggests that the cut will not produce a strong radial bias to the sacrifice. With the choice of the cut at ITC=0.6 the sacrifice is even less of a problem. The choice of the cut in this position still however removes a large number of flat TAC events.

The ITC cut, as well as removing events where the hits are distributed in a uniform way throughout GT_valid, the ITC cut will also remove events with bimodal hit distributions. This means that the ITC cut will also cut events which are coincidental occurrences of two or more events within GT_valid. This turns out to be important when attempting to measure sacrifice at high NHIT using the ^{16}N source, as events at very high NHIT are often in this regime due to the pileup of different events.

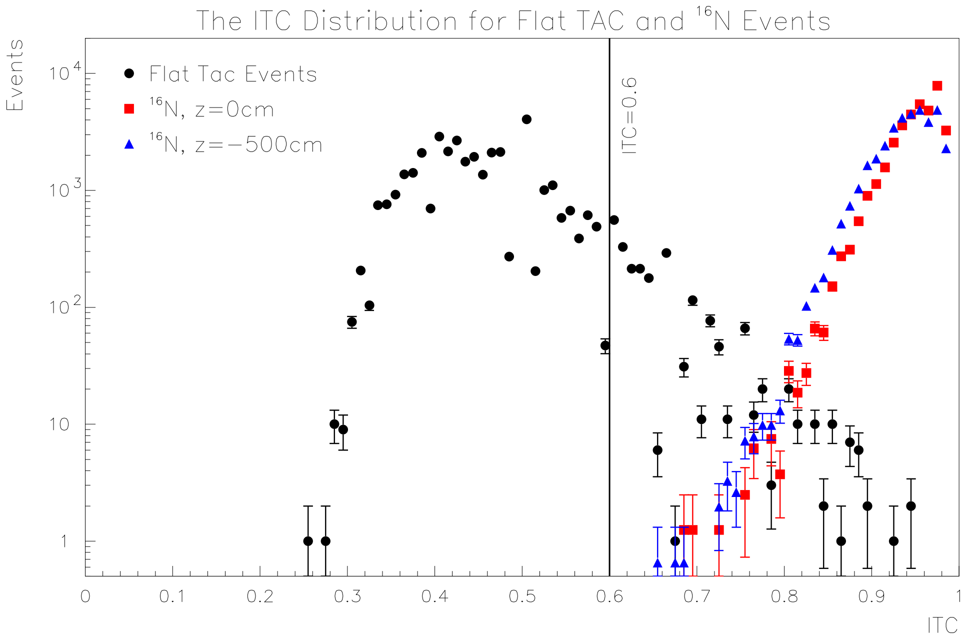

Figure 3.18: The ITC distribution for ^{16}N events flat TAC burst.

3.4 The high level cuts

The high level cuts are a set of further cuts that examine events in a more detailed way than the instrumental cuts. Apart from reconstruction itself, the high level cuts all use the result of the reconstruction process. There are two figure of merit cuts based upon the result of the fit. There are then two further cuts designed to select neutrino events, not reject background, also dependent on the reconstructed vertex. There is one final cut that is used as a high level cut, the muon follower cut, discussed in chapter 6.

3.4.1 Reconstruction

Reconstruction is the determination of the position, direction and possibly energy of the electron, or γ-ray that caused the observed PMT hits. This is a complicated subject and it is not possible to do justice to the full complexity of the subject here. The package that carries out this procedure in SNO is called a fitter.

The fitter of choice for neutrino analysis is the path fitter [38]. This fitter is a maximum liklihood fitter that uses probability density functions (PDFs) in both time residuals and angle. These PDFs are based upon 5MeV electrons from the Monte Carlo and timing information from the laserball. The fitter has a vertex resolution of 16cm and an angular resolution of 26.7^o in the D_2O [57]. The most important property of the fitter,

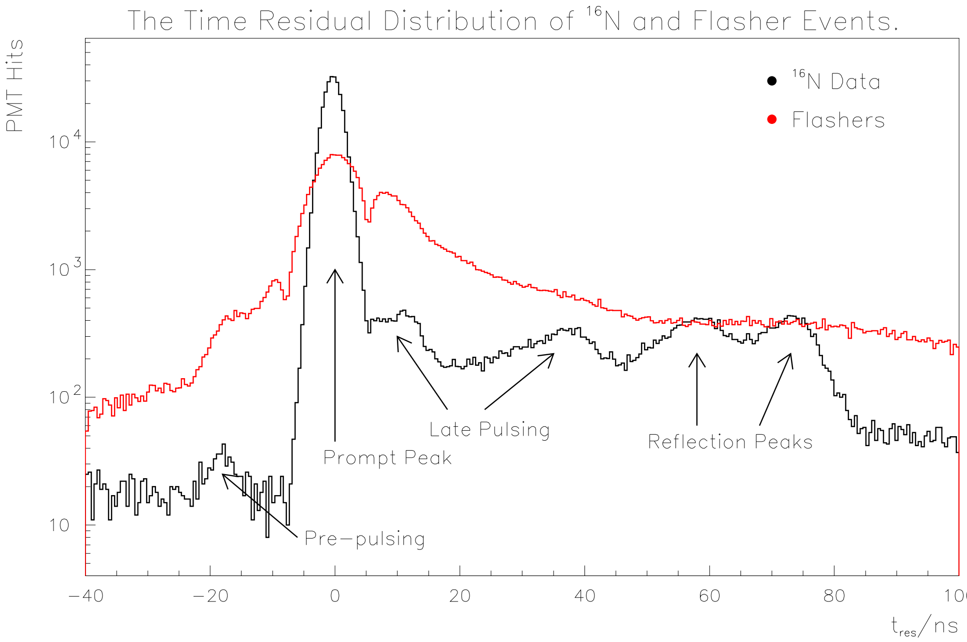

Figure 3.19: The time residual distribution for ^{16}N events at the centre of the detector and flasher events.

the vertex shift: a systematic mis-reconstruction of events is bound by a 1% radial shift.

From the point of view of a cut the fitter is important in two ways. The first of these is fitter sacrifice, mainly due to the way the path fitter is seeded. This is described in section 5.3.3. The second way is in the use of the goodness of fit, or figure of merit cuts, calculated with the vertex. These cuts are used for two reasons. Firstly they check that the fit hypothesis used is compatible with a single Čerenkov electron, not an event such as a flasher or relativistic proton. Secondly they help remove badly fit events from the data set. There are two figure of merit cuts, the first of which is a KS test against a flat azimuthal distribution about the fit direction. The second is a 2D-KS test that compares the event to the angular PDF used in the fit.

3.4.2 The in-time ratio

The in-time ratio or ITR cut is based upon the time residual distribution of events. The idea behind the cut is that the time residual distribution for Čerenkov events that have been correctly reconstructed is very sharp, whereas that for badly fit events, or instrumental backgrounds it is not. The time residual distribution for ^{16}N events in the centre of the detector, and the golden flasher sample, are shown in figure 3.19.

Visible in the ^{16}N peak are several features that appear because the ^{16}N source is in the centre of the detector. These include the pre and late pulsing peaks, as well as two reflection peaks from the PSUP[12], but the time residual distribution is dominated by the prompt light peak. The reflection features become more spread out as the source moves off centre, but the main feature, the prompt peak remains unchanged. The time residual distribution for flashers looks quite different. The prompt peak is much broader and contains a smaller fraction of the light than is the case for ^{16}N events.

The ITR cut works by comparing the light in the prompt peak to the total amount of light. The number of hits within a time residual window of -2.5ns $\rightarrow$ 5.0ns is divided by the number of calibrated hits, to construct the in-time ratio (ITR). The use of channels with valid calibrations avoids problems that could arise due to variation in the number of channels with valid calibrations. The choice of this time window is very aggressive, and this does mean that the ITR cut is sensitive to the quality of PMT time calibrations as a function of time. This has implications for the stability of the cut.

By the use of the fit vertex to constrain the time information provided by the PMTs the ITR cut is much more effective at removing backgrounds that reconstruct than the time based instrumental cuts. This can be seen by comparing the ITR distributions of flashers and neck events to the distribution from calibrations sources. The ITR distributions for these events are shown in figure 3.20. The neutrino data shown in the figure, agrees with the distribution from ^{16}N data, except at low ITR values. These events are badly fit events, and a small amount of residual contamination that has been passed by the instrumental cuts. Even with the choice of the cut at ITR=0.55, at the very tail of the Čerenkov distribution a large fraction of flashers and an even larger fraction of neck events are removed by the ITR cut. The cut parameter is chosen at 0.55 to attempt to compensate for the very aggressive choice of the time window. This allows some leeway for bad time calibrations in a typical event before the sacrifice of the cut becomes a serious problem.

3.4.3 The θ_{ij} cut

The θ_{ij} cut is based upon the θ_{ij} distribution, a distribution designed to select Čerenkov events based solely upon the hit pattern observed from the fit vertex. The value of θ_{ij} is the mean angle between all possible hit PMT combinations measured from the fit vertex[72]. For a typical Čerenkov event the value of θ_{ij} is 1.1 radians, but the distribution is quite broad. The θ_{ij} distribution for CC electrons and γ-rays from NC events in pure D_2O is shown in figure 3.21. The figure shows that there is a distinct difference between γ-rays

[12]Pre and late pulsing peaks are due to the PMTs, and the two PSUP reflection peaks are due to the reflections at 38^o and 180^o. These reflections are principally due to the concentrators and PMTs. Reflections are also seen from the acrylic vessel, but this peak is small and hard to see on the plot.

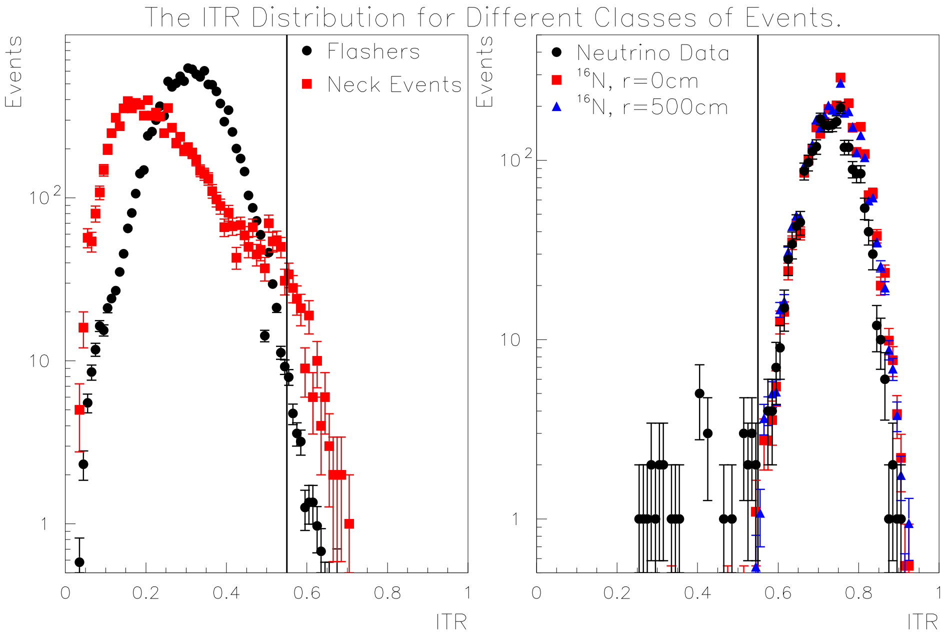

Figure 3.20: The ITR distributions for backgrounds, calibration data and neutrino data after the application of instrumental cuts.

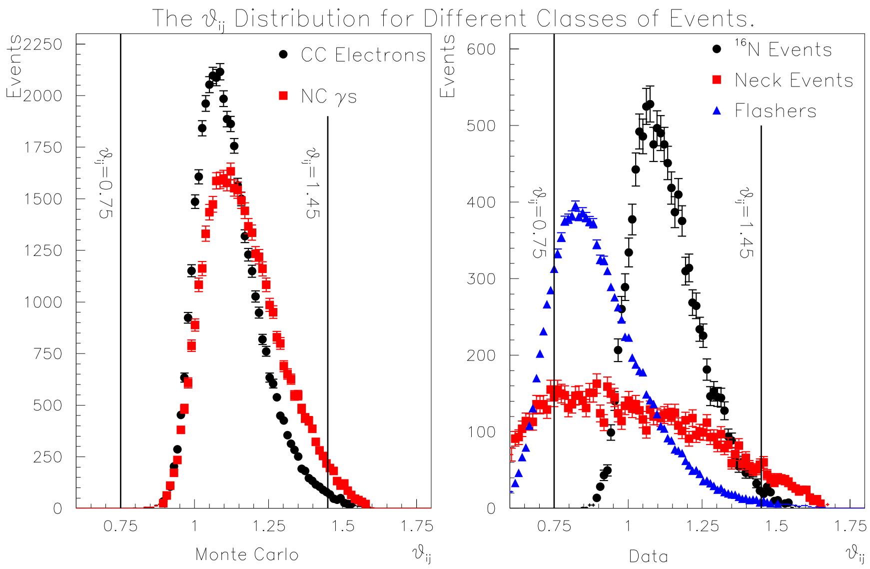

Figure 3.21: The left plot shows the θ_{ij} distributions from the Monte Carlo of both CC electrons and γ-rays from neutral current interactions in pure D_2O. The right plot shows the θ_{ij} distributions for flashers, neck events and events from the ^{16}N source. The distributions are normalised such that each contains the same number of events.

and electrons, and this causes problems when measuring the sacrifice of the θ_{ij} cut. Also shown in the plot is the θ_{ij} distribution for flashers and neck events. The θ_{ij} cut is not that good at removing these events with the cut parameters set where they are.

However the cut parameters have not been set to remove flashers and neck events, but instead a different criteria for cut selection has been used. The low end cut is very effective at removing badly fit events, in particular events from the H_2O that have mis-reconstructed inside the D_2O. These events have a much narrower hit distribution from the fit vertex, and this appears as a low value of θ_{ij}. The high end cut is even more powerful. Isotropic events have an expected value of $\theta_{ij} = \frac{\pi}{2}$, separated from the bulk of the Čerenkov region. Leslie events which are isotropic from their reconstructed vertex have no instrumental cut aimed at them. Thus the θ_{ij} cut is important as the high end cut is the only cut, other than a fiducial volume cut that explicitly removes these events. It is for this reason that the θ_{ij} cut is used despite the changing sacrifice for electrons and γ-rays . This problem must be dealt with in the sacrifice analysis.

3.5 Conclusions

The large number of instrumental backgrounds inside the SNO detector leads to the need for a large number of cuts to remove these events. These cuts fall into two distinct categories: the instrumental cuts and the high level cuts. These two sets of cuts differ in the philosophy behind them, the instrumental cuts remove background, the high level cuts accept signal.

This chapter has dealt with the cuts themselves, not the results of the cuts when applied to neutrino data. This will be dealt with in chapter 5 where the sacrifice of the cuts and the contamination will be measured. However it is satisfying to observe the change in the raw NHIT spectrum shown in figure 3.1. The effect is shown by splitting the instrumental cuts into groups based upon their target categories and showing the effect of each. A final reduction in the NHIT spectrum based upon the high level cuts and a fiducial volume cut at 550cm is also included. The reduction of the NHIT spectrum is shown in figure 3.22. The resulting NHIT spectrum has the shape expected for observation of solar neutrinos, however further statements on the effect of the cuts can only be made after the sacrifice and contamination analyses are complete.

The one instrumental background and associated cuts that were not discussed in this chapter were the neck events. These are the subject of the next chapter.

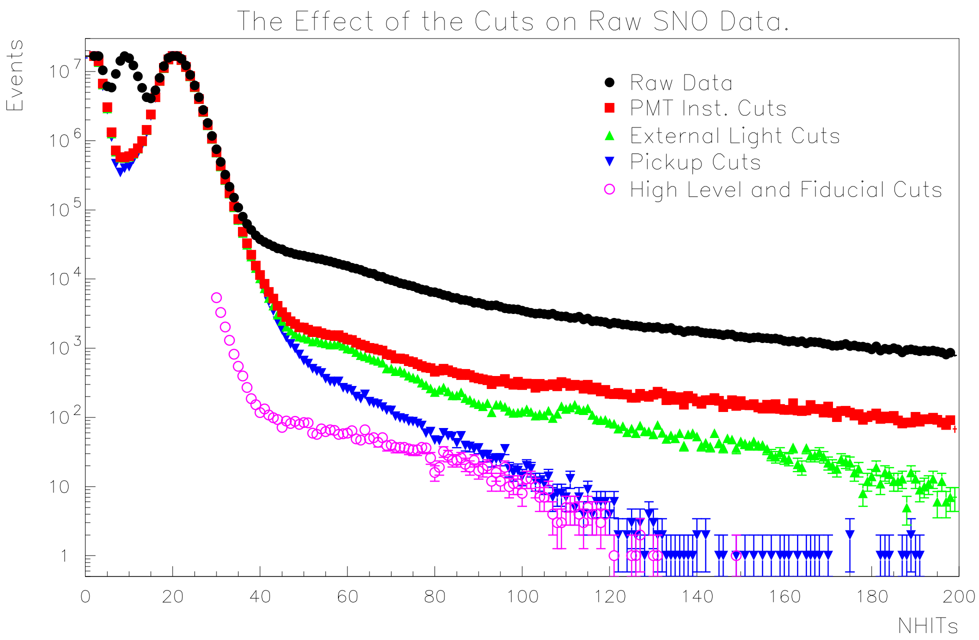

Figure 3.22: The reduction of the NHIT spectrum when various levels of cuts are applied.

Chapter 4

Neck Events

Neck events are one of the most serious instrumental backgrounds in the SNO detector. Not only do many of them fall into the NHIT regime of interest, but their reconstruction properties make them difficult to remove effectively by the use of a fiducial volume cut. Due to the range of the neck event NHIT spectrum these events also result in a large background to the through-going muon analysis in SNO.

Neck events are removed by a number of different instrumental cuts. The most effective of these rely on use of the PMTs outside the PSUP, to veto events. However neck events can also be removed by taking advantage of some of the other characteristics of the events, such as the broad pulse time and the bursty nature of the events.

The distinctive hit pattern, which is not used to attempt to remove the events, is used to estimate the residual number of neck events left in neutrino data after the application of the instrumental cuts. The number of neck events found by this method shows that the residual contamination of these events in the final data set is very small.

4.1 A pain in the neck

Historically neck events were first noticed as a class of bubbler event, during the commissioning phase of the detector. Bubbler events are described in section 3.1.6. It was found that a class of bubbler event did not disappear when the bubbler tubes were flooded, and cannot therefore be associated with such a tube. The events look like bubbler events coming from the neck of the acrylic vessel. These are the neck events, an example of which is shown in figure 4.1.

In order to do something about neck events the detector had to be upgraded. The upgrade was to install PMTs at the top of the detector. These allow neck events to be vetoed during the analysis. Two tubes were installed in the neck of the detector, and are known as the neck tubes. Two more were installed in the pipe-box and are known as the

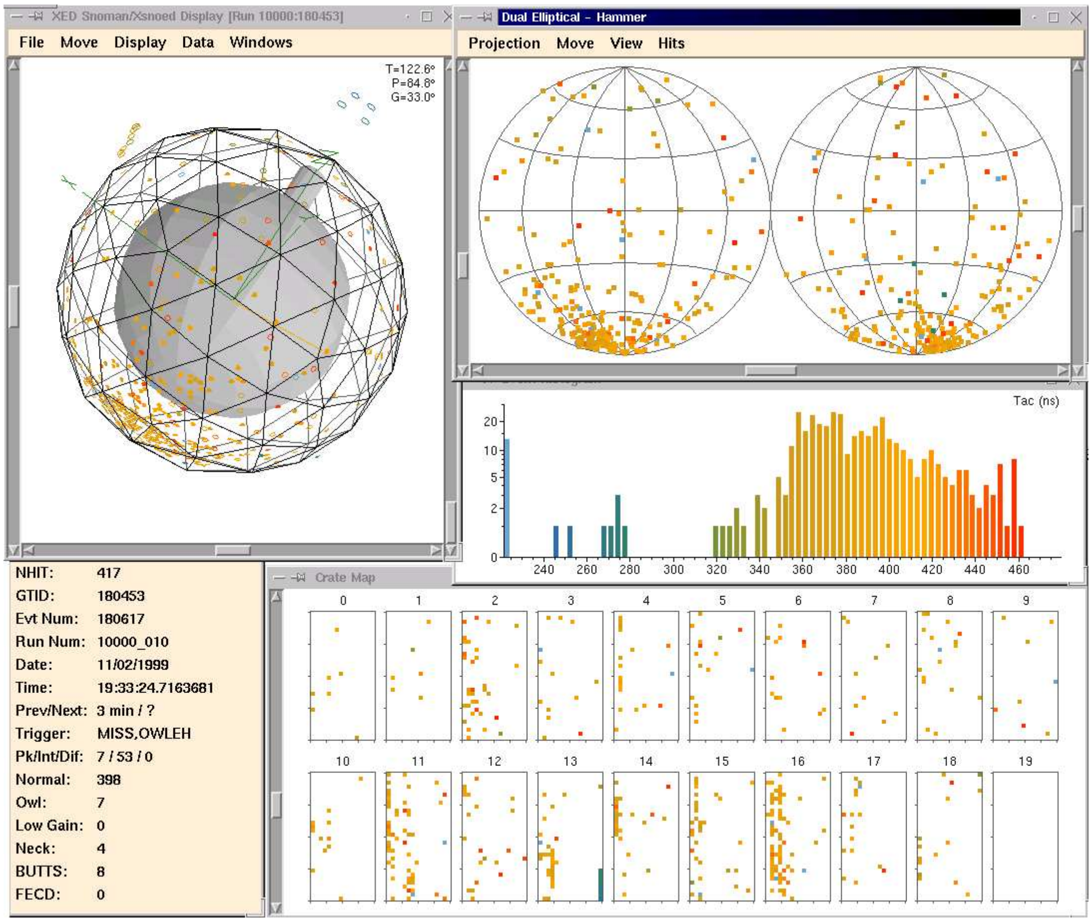

Figure 4.1: A neck event, showing the ECA times of the tubes.

pipe-box tubes. The neck tubes were plugged into spare OWL channels on crate 13, card 15. The two neck tubes are connected to channels 1 and 2 and the pipe-box tubes to channels 0 and 3. Due to this configuration these tubes are subject to crosstalk from one another. As they are connected to OWL channels the neck tubes are also plugged into the OWLESUM trigger system.

It can clearly be seen in the event shown in figure 4.1 that the tubes in the neck and those in the pipe box all fire. The tagging on these tubes is crucial for removing neck events. Figure 4.1 shows the neck event where the ECA times of the PMTs are displayed. The use of ECA time allows comparisons to be made between the neck tubes and the tubes inside the PSUP[1]. The event clearly shows that the neck tubes have fired early in the event. Due to the early firing, the neck tubes often appear very early in the TAC window. This led to the fear that many neck events would not have neck tubes

[1]There is no PCA calibration for the neck tubes.

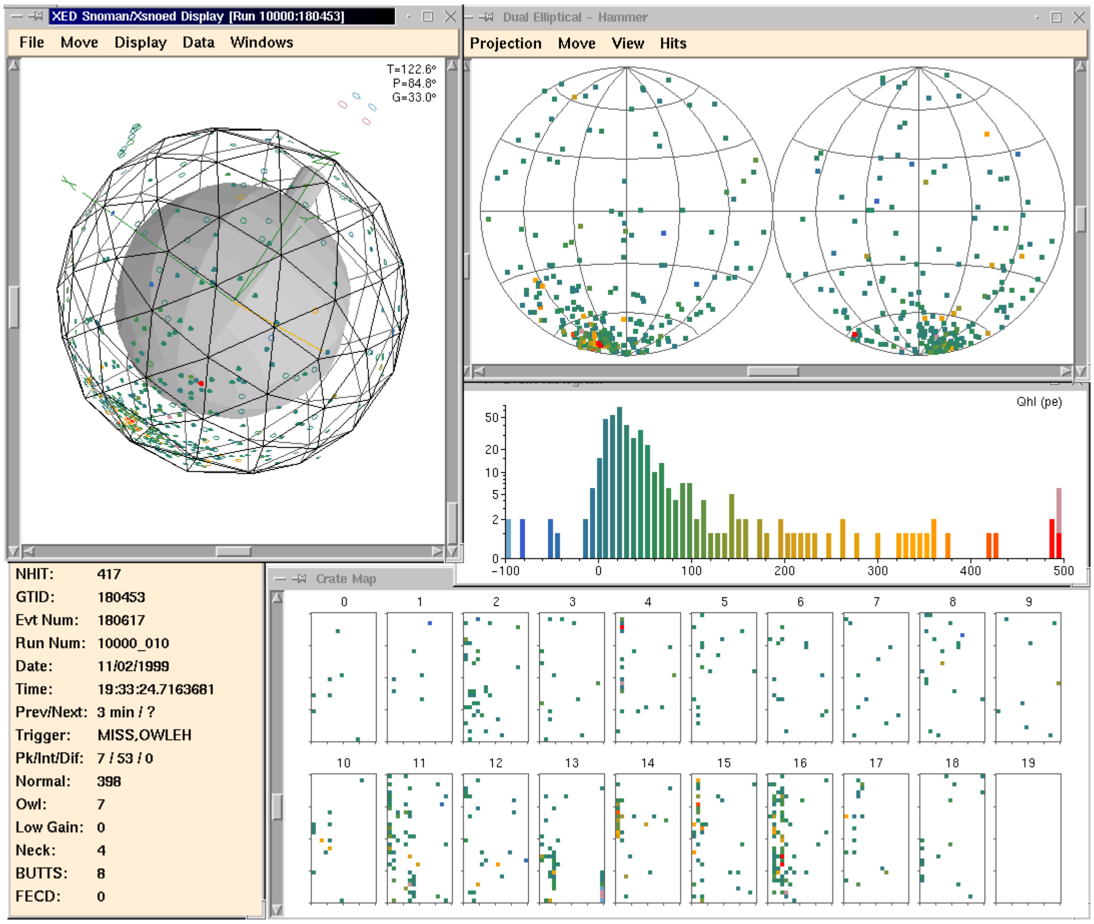

Figure 4.2: A neck event, showing the pedestal subtracted QHL charge.

in them, due to the tubes being so early that they get cut off at the start of the event. Despite the lack of firm evidence that this was the case, extra cable was added between the neck tubes and the electronics. This extra cable delay moves the neck tubes closer to the middle of the TAC window, ensuring that they will not be lost in this manner. The extra cable delay was added on the 16$^{\text{th}}$ of December 1999, shortly after the start of production running.

It is also instructive to examine the event in charge as well and this is done using the long integrate charge in figure 4.2. The event shows that the two neck tubes have fired with a high charge, while the charge on the pipe-box tubes is consistent with crosstalk from the neck tubes. Another property of neck events clearly shown in the figures is that the BUTT tubes are a good indicator of neck events, as some of the tubes are pointing at the neck region of the detector. Finally the neck event shown fired the detector via the OWLEHi trigger. The high charge on the neck tubes causes this trigger to fire. The

NHIT trigger is then collected as a missed trigger[2].

4.1.1 Light Production in Neck Events

The light production mechanism in neck events is not known. A number of theories exist the most compelling of which is that they are caused by static discharge. The neck is a complex region of the detector, containing a number of different components that could contribute to light from the neck. This part of the detector contains the interface between five different regions of the detector: the acrylic, the D_2O, the H_2O and the cover gas above both media. In particular the acrylic is a very good insulator, and can therefore build up static charge. The water level is not completely stable but fluctuates slightly. As it fluctuates different parts of the acrylic of the neck become exposed. These changes in the water level can cause the charged acrylic to discharge emitting light into the detector. The acrylic could also discharge if the amount of charge grows too high. The humid cover gas can facilitate this.

However the light is emitted into the neck region of the detector it is then easy to see how it can travel into the detector. The neck is one of only a small number of openings into the PSUP, and it is the biggest of these, so light produced here has no problem making it into the detector. Light produced here also has a potential guide from the neck inside the PSUP. The acrylic vessel could trap light inside by total internal reflection, and then transport it down inside the PSUP. The light could then escape from the acrylic where the neck joins the spherical part of the vessel. This can facilitate the transport of light produced in the neck into the PSUP. The acrylic that makes up the neck is not of the same optical quality of that in the rest of the vessel to attempt to stop the neck acting as a light pipe in this way. However the effect could still be going on at some level.

4.2 Properties of Neck Events

Neck events have a number of properties that make them a very dangerous background to the analysis of SNO data, but also allow them to be removed before any more complicated analysis is carried out.

In order to study the properties of neck events a sample is required. For the following plots the selected neutrino dataset (see appendix C) was scanned using the neck cut described in section 4.3.1. This cut was used in conjunction with the retrigger (section 6.2.3) and q/NHIT (section 3.3.8) cuts to give a large relatively pure sample of neck

[2]The missed trigger bit is set when another trigger occurs after the 10ns time window for triggers has elapsed following a previous trigger, but still during the same event.

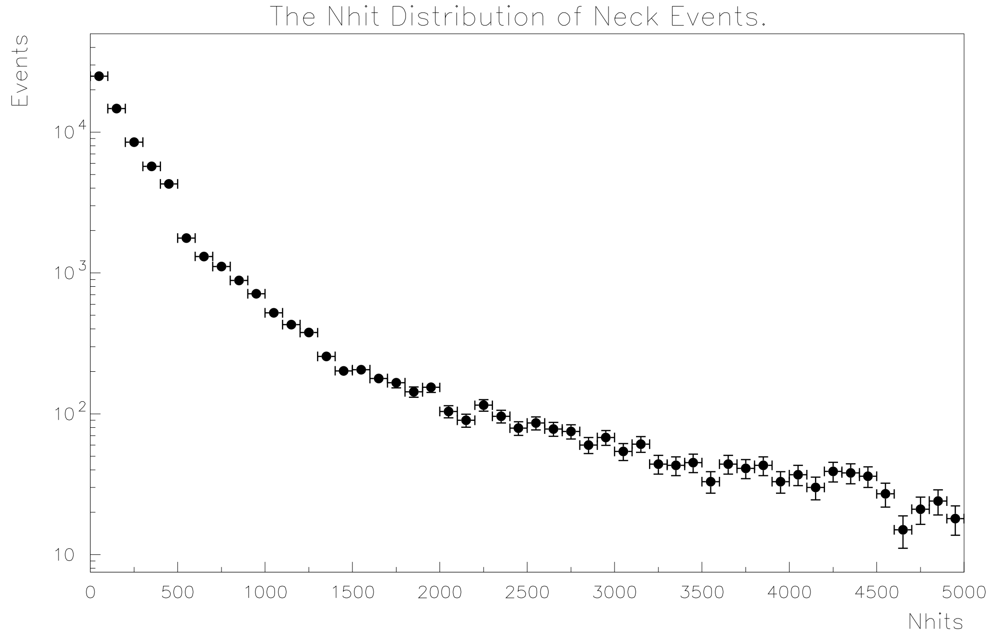

Figure 4.3: The NHIT distribution of neck events from the selected neutrino dataset.

events. These cuts are required as the neck cut does not discriminate well between pickup events, retriggers and the neck events desired.

The data set obtained contains neck events in the range 30-5000 hits. Events below 500 hits are then reconstructed using the path fitter. While this data set contains a relatively pure sample of neck events there are some other classes of event remaining. The biggest of these are low energy Čerenkov events coincident with activity in the neck tubes. For the moment these events will be ignored by placing a cut at 40 hits. However they will be considered later when the sacrifice of the neck cut is discussed (see section 4.5).

The first distribution to be examined is of course the NHIT distribution, shown in figure 4.3. The distribution shows that while there are events at very high NHIT, which is how neck events were first seen in SNO data, the vast majority of neck events lie at much lower in NHIT values, inside the analysis window for study of solar neutrinos.

Using the sample of neck events the reconstruction properties of the events can also be examined, shown in figure 4.4. The radial distribution shows a very disturbing property of neck events. Despite the true vertex in the neck of the detector, these events tend to reconstruct inside the acrylic vessel. This means that reconstruction does not provide a good defence against neck events, making it even more important that the cuts designed to work against them do an effective job. The neck events of course come from only one place in the detector and this is reflected in their reconstruction. It is found that the neck

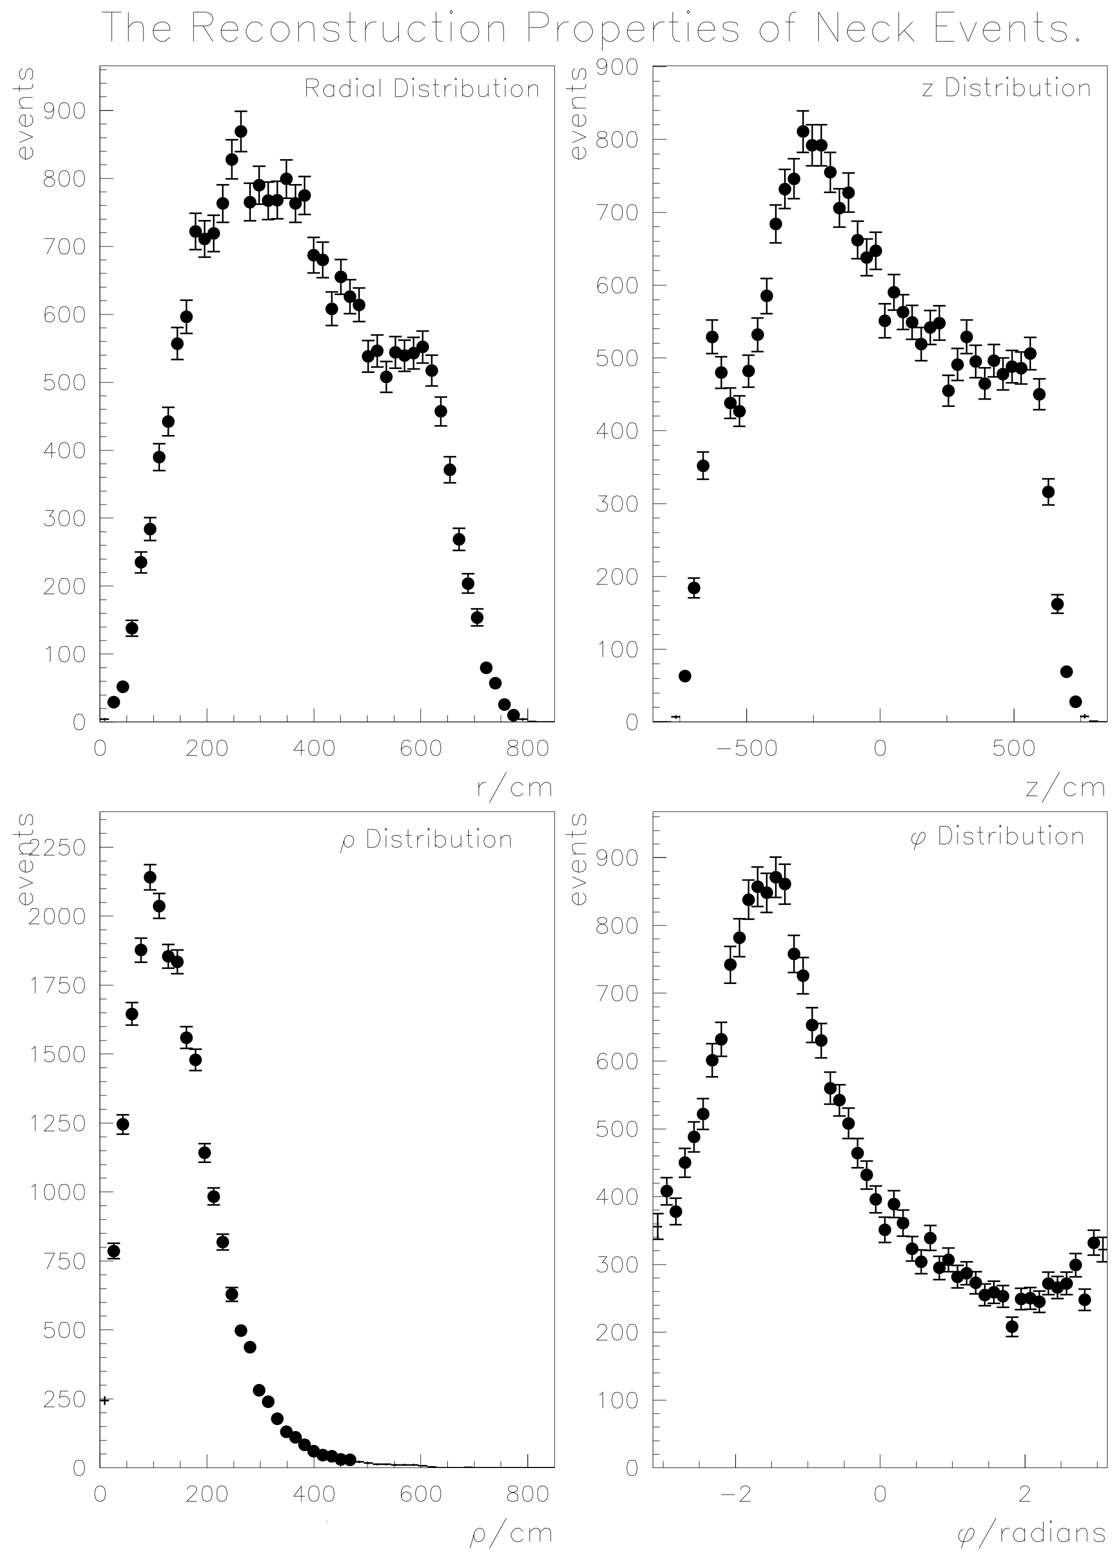

Figure 4.4: The reconstruction properties of neck events shown in reconstructed radius and in cylindrical co-ordinates (z, ρ, ϕ).

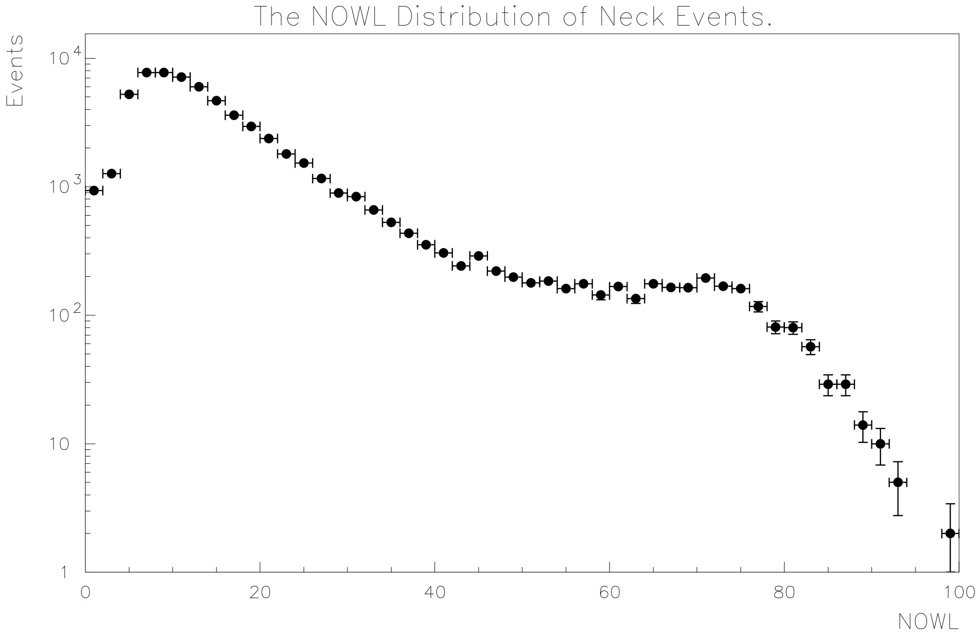

Figure 4.5: The number of OWL and BUTT tubes firing in neck events.

events reconstruct very close to the z-axis of the detector. This can be demonstrated by examining the events in cylindrical co-ordinates, also shown in figure 4.4. These plots show that the events reconstruct close to a line down the centre of the detector. This line is actually slightly offset along the y-axis perhaps giving us a clue to the origin of these events.

There are three basic properties of neck events, however, that do distinguish them well from Čerenkov events. The first is the presence of various veto tubes (see section 2.4) in neck events. As well as the neck tubes, the use of which is described in detail in section 4.3.1, a large number of outward looking tubes also fire in neck events. The number of outward looking tubes firing in neck events is shown in figure 4.5. As in a typical contained Čerenkov event there will typically be no outward looking tubes firing at all, these tubes provide a powerful discriminant against neck events.

The other important property of neck events is the width of the pulse that produces neck events is very broad in comparison to that from Čerenkov events. This broad pulse is responsible for the bad reconstruction properties of neck events, but it allows cuts based on the time distribution of hits to do a good job of removing these events.

Finally the last property of neck events that helps in their removal is that they occur in bursts. This can be seen by examination of the time difference, δt between neck events, shown in figure 4.6. The δt distribution is not an exponential distribution, as would be

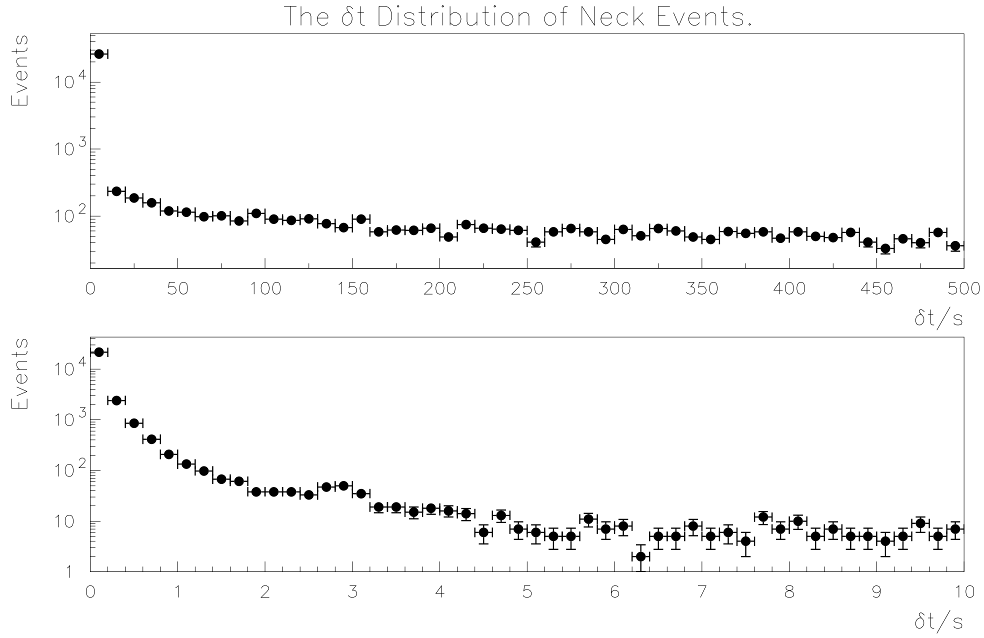

Figure 4.6: The δt distribution for neck events.

expected for events distributed in a random way. Indeed the occurrence of the majority of the events in the first bin tells us that neck events occur in bursts. Looking at this first bin, as shown in the lower plot, the typical δt for a neck burst is a few tenths of a second.

4.3 The Neck Event Cuts

While most of the cuts in the suite of instrumental cuts will remove neck events to a greater or lesser extent, there are a number of cuts which specifically target neck events. Some of these cuts are designed solely to remove and in one case identify neck events. Others are designed with neck events in mind, but neck events are not the sole targets of the cut. While neck events do have a distinctive hit pattern in the detector, this information is not used in the neck event cuts, because as with all of the instrumental cuts, these cuts are designed to be as simple as possible. The use of the hit pattern to measure neck event contamination is discussed in section 4.4.

Run	Run length/h	Total Events	Selected Events	Neck Events
10031	8	905628	241	177
10040	35	2246913	550	382
10196	11	700981	427	388
Total	54	3853522	1218	947

Table 4.1: The runs used to construct the neck cut data sample.

4.3.1 The Neck cut

The most important of the neck event cuts is the neck cut. This cut uses the neck tubes installed in the detector to attempt to uniquely tag and then remove neck events. This means that the goal of the neck cut is slightly different from that of the majority of the instrumental cuts as it is specifically targeted at a single class of events and attempts to leave alone some of the other classes. For example, when flashers are considered from the perspective of the neck tubes, they do not look any different from Čerenkov events. This means that the performance of the cut on data samples that include flashers can be used to judge how well the cut will behave given a sample of solar neutrinos.

There are two classes of instrumental backgrounds that are ignored however. The first of these are retrigger events. Retriggers from both neck events and other events can contain one or more neck tubes with random charge and time. To attempt to avoid these events using the neck cut would degrade the performance of the cut sufficiently that it was decided to ignore retriggers when formulating the cut. Similarly pickup events where pickup occurs on the neck tubes can also confuse the cut. For this reason these events are also ignored in the construction of the cut.

Armed with this knowledge the specificity of the cut can be used to construct samples of neck events when required. The retriggers and pickup events can be removed by the use of the retrigger and q/NHIT cuts respectively. This method was employed to construct all neck event samples discussed except for the one used in the construction of this cut.

The data sample used to construct the cut

Before the cut can be fully constructed, a data sample containing both neck events and potentially good events needs to be developed. Three runs were chosen for this job, all before the changes to the cabling of the neck tubes were made. These runs are shown in table 4.1

The timing of the neck tubes will play an important part of the neck cut. For this reason a second data set containing events after the changes were made to the neck tube cables was also constructed. This sample was only used to measure the changes to the

Run	Run length/h	Total Events	Selected Events	Neck Events
10779	21	1272072	707	643
10801	24.5	1641583	140	24
Total	45.5	2913655	847	667

Table 4.2: The runs used to make any adjustments to the neck cut based upon the changes to the neck tube cabling.

neck tube timing, and played no other role in the development of the neck cut other than compensating for these changes. The two chosen runs are shown in table 4.2

The neck cut will only cut neck events containing neck tubes so the first step in reducing the number of events is to require that at least one neck tube fires. The next step is to remove all events lying outside the NHIT validity range of the cut (10-5000 hits). An additional NHIT cut of more than 24 hits is then applied to reduce the amount of hand scanning required. At this point the retrigger and q/NHIT cuts are applied to remove retriggers and pickup from the sample. Finally any remaining PulseGT and orphans events are removed.

This sample will contain neck events, events where light reaches the neck tubes from within the PSUP, or events containing noise on the neck tubes. The latter two categories are combined together and classed as accidental events. The neck cut with some best guess parameters was then used to mark bits in a DARN bank (see appendix A) to provide a first guess for the neck events. This procedure was used simply to reduce the number of clicks that would have to be made in the following hand scanning. All the events tagged in this method as neck events were hand scanned and mis-tagged events were unmarked. The data set that was not tagged originally was also hand scanned and any missed neck events were tagged. This results in a data set of events with a DARN bank that contains a bit separating neck events from accidentals, to the limit of hand scanning. The neck events found in this way were also used to construct the neck event standard data sets, discussed in chapter 3.

The NHIT spectrum of the neck events and accidentals is shown in fig 4.7. The NHIT distribution of neck events is as expected with many events at high NHIT, as well as smaller NHIT values. The accidentals are dominated by events in the lowest NHIT bin. These are low energy radio-activity events. There are also a small number of events at higher NHIT.

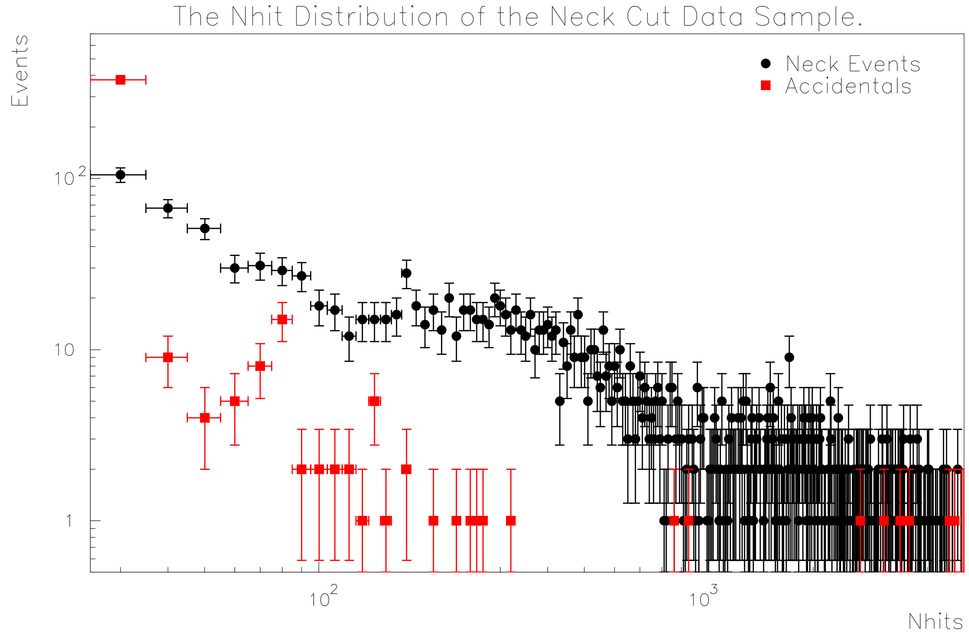

Figure 4.7: NHIT spectrum of the events classed as neck events and those classed as due to accidental firing of the neck tubes.

The neck cut algorithm

The neck cut has two modes of operation. These modes are independent of each other with the exception that they both use the neck tubes but not the pipebox tubes. The first mode is the simplest: if both of the neck tubes fire in the event, then event is considered to be a neck event and is tagged by the cut. This technique is known as the neck mode and it removes a large number of neck events.

The other technique is more complicated and relies upon more knowledge of the neck events. It will cut neck events that only fire one neck tube with a relatively small charge. This cut uses the fact that the neck tube charge should still be above the single photo-electron (p.e.) peak and that the neck tube should be early. For this reason this mode is known as the QandT mode. The first step is to apply a charge cut to the neck tube. If this tube is above the single p.e. peak, then the timing of the tubes is taken into account. This is done by first calculating the mean time for the PMTs in the lower hemisphere of the detector, and then comparing this to the time of the neck tube. If the neck tube is earlier by a set amount of time then the event is considered to be a neck event. The QandT mode uses ECA + HCA calibrated charges and times by default. It will however work given raw ADC values if required, as for example within the supernova trigger. Given PCA times, for the inward looking PMTs the code will re-calculate the required level of calibration from the raw ADC values. This step is required as the neck tubes do not have PCA constants and it is best in this case to compare like with like.

The cut also has an NHIT validity range, with an upper and lower limit. The lower limit is used to avoid numerical difficulties with low NHIT events and is set to ten hits. The upper limit is set to five thousand, above this threshold the cut is overwhelmed by muons.

The parameters of the neck cut

Of the used modes of operation of the neck cut only the QandT mode has adjustable parameters. However this mode of the neck cut is the primary mechanism in the entire suite of instrumental cuts to remove neck events that do not contain large amounts of light and do not occur in bursts. These are the class of neck events that are the most serious concern, as they fall directly inside the analysis window for solar neutrinos. For this reason it is important that the parameters of this cut be as aggressive as possible, without causing too much sacrifice.

As discussed in section 4.3.1 data samples containing both neck events and events with neck tubes are collected and tagged. These are used to determine the cut parameters. Figure 4.8 shows the distribution of the charge and time difference for the data sample taken before the cable changes[3].

Figure 4.8 also shows that the QandT cut will operate best using the and of the charge and time cuts instead of the or of the two. The use of the and allows both cuts to be made as aggressive a possible, without compromising the control of the sacrifice of the cut. The charge and time cuts used are to first order independent of each other so the parameters of each cut are determined independently.

The charge cut

The SNO detector measures the charge deposited on a each PMT in three different ways. The neck cut applies the charge cut to all three, however the parameters are defined on the basis of the high gain long integrate charge (QHL) only. This measurement of the charge provides the best discrimination for neck events, as the long integrate time allows the collection of as many photons as possible from the neck event. The high gain is used to to help separate the multi-photon signal from the single p.e. peak as much as possible. The other two charges show very similar distributions to QHL, and the use of

[3] In this and all subsequent figures in the case where both neck tubes fire, only the tube with the highest charge (taking into account the negative rail) will be shown. This helps prevent the addition of noise to the charge spectra due to crosstalk from one neck tube to the other.

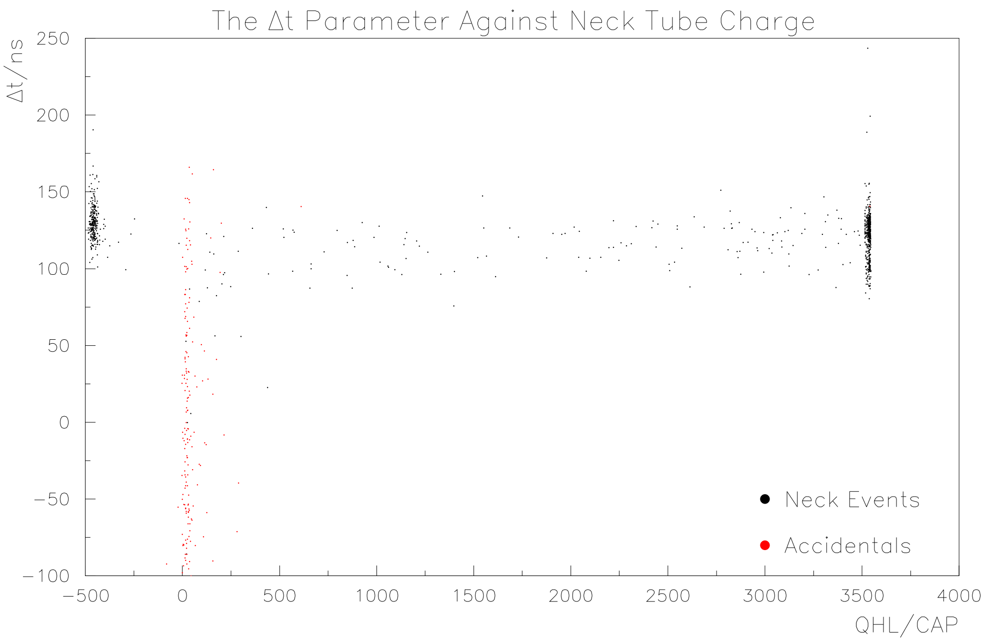

Figure 4.8: QHL vs. the time difference parameter for neck events and events with neck tubes.

these charges alongside QHL provides a redundancy against possible problems with the QHL measurement.

Figure 4.9 shows the pedestal subtracted QHL on the neck tubes for neck events and other events with neck tubes. The left hand plot shows the entire charge spectrum of the neck tube with the highest charge in each event. There are three distinct peaks in the charge spectrum. The small charge peak close to zero for the events classed as accidentals is the single p.e. peak. The cut is designed to keep events with charges in this region. The other two peaks are the positive and negative rails seen for the neck events. The charge measurement is railed due to the large amount of light created during a neck event. Away from these peaks the events are dominated by the neck event class. In order to accurately set the charge cut values a closer look at the region around the single p.e. peak is required. This is shown in the right hand side of figure 4.9 along with the final cut values. The cut is most sensitive just above the peak and this is also where the cut is most aggressive. The lower cut to remove the negatively railed charges is much less aggressive as there is much better separation at this end of the spectrum. If the charge lies outside the window defined in figure 4.9 in any of the three charges the events is deemed to be a possible neck event and the time cut is then applied.

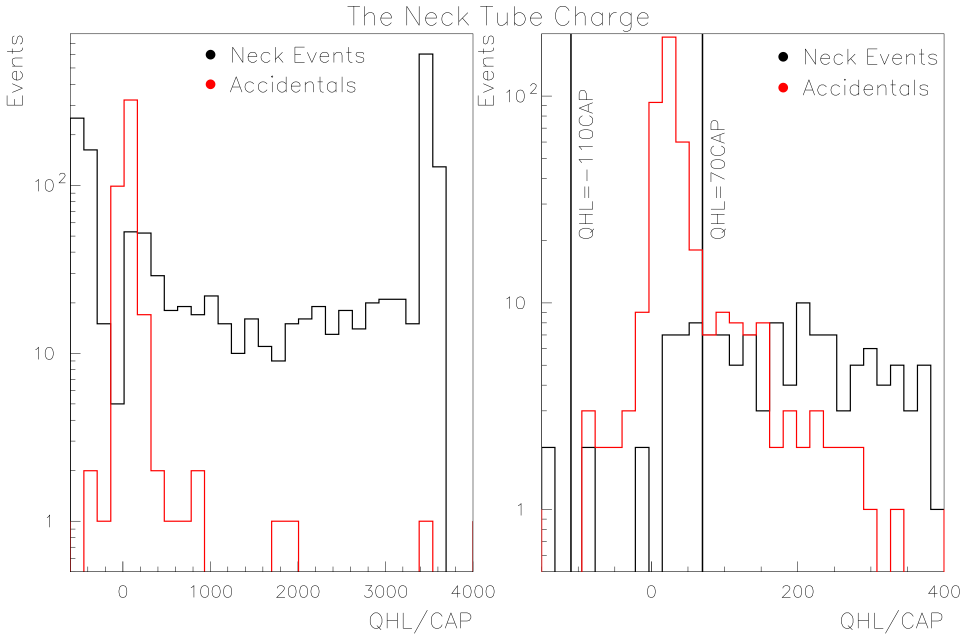

Figure 4.9: Pedestal subtracted QHL for the neck tubes.

The time cut

The time cut requires that the neck tubes be early as would be expected for light from the neck. The actual algorithm used is to require that the neck tube be a set number of nanoseconds earlier than that average time of the PMTs in the lower hemisphere of the detector. This average time is only calculated if a neck tube has already failed the charge part of the neck cut, improving the speed of the algorithm. Fig 4.10 shows the time difference between the neck tubes and the mean time for PMT hits in the lower hemisphere of the detector (Δt). This parameter is shown for both neck events and accidentals. The plot clearly shows that the neck tubes are early in most of the neck events, but are more uniformly distributed for the other classes of events. The right hand plot shows the same data but more closely centred around the cut region. The actual cut value is chosen to attempt to maximise the rejection power of the cut, while minimising the sacrifice. This is a slightly different approach to that taken with the charge cut which was made as aggressive as possible. This combination was found to be the most efficient at removing neck events, while keeping sacrifice minimised due to the distribution of the non-neck event signal for the two parameters. The cut value of 70ns shown on this plot is applicable to the data taken before the changes to the neck tube cabling. The time cut is obviously sensitive to the changes to the cabling, and the cut parameter has to be

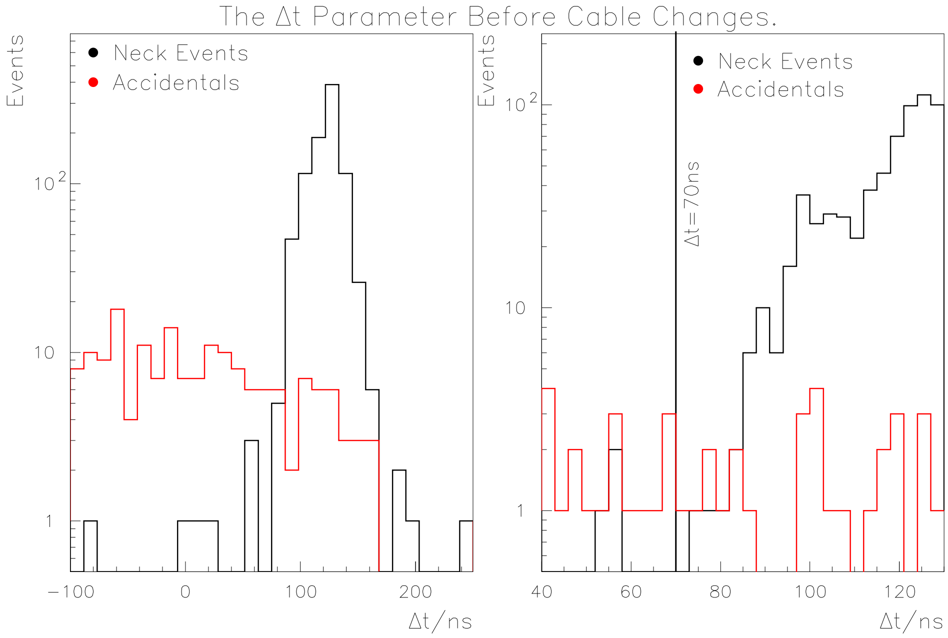

Figure 4.10: Time difference between the neck tubes and the mean PMT time for hits in the lower hemisphere, before the changes to the neck tube cabling.

recalculated for data taken after this time. The time distribution after these changes is figure 4.11. This figure suggests that the value for the time cut should be 15ns for data after the cable changes, 55ns earlier than the previous value. This number agrees with the expected change calculated from the length of cable added in the cable delay.

The time cut used relies on the ECA calibration of the neck tubes and this requirement can lead to extra contamination from the cut due to situations where the neck tubes fails ECA time calibration. This can occur for two reasons, the most common of which is that the neck tube could fall into the TAC curl. As the neck tubes are expected to fire early in neck events this is a particular problem. Neck tubes can also fail ECA, as can any channel due to problems with the TAC slope. These features are further discussed in [50]. In order to protect the neck cut against TAC failures on the neck tubes the neck cut will default to a cut based on raw TAC if the neck tube has failed ECA time calibration. The determination of the TAC cut before cable changes is shown in figure 4.12. As the raw TAC in SNO has early times at high TAC values the direction of this cut is in the opposite direction to the calibrated cut. For data after the installation of the cable delays the TAC cut has been re-calibrated in a similar fashion to the calibrated cut.

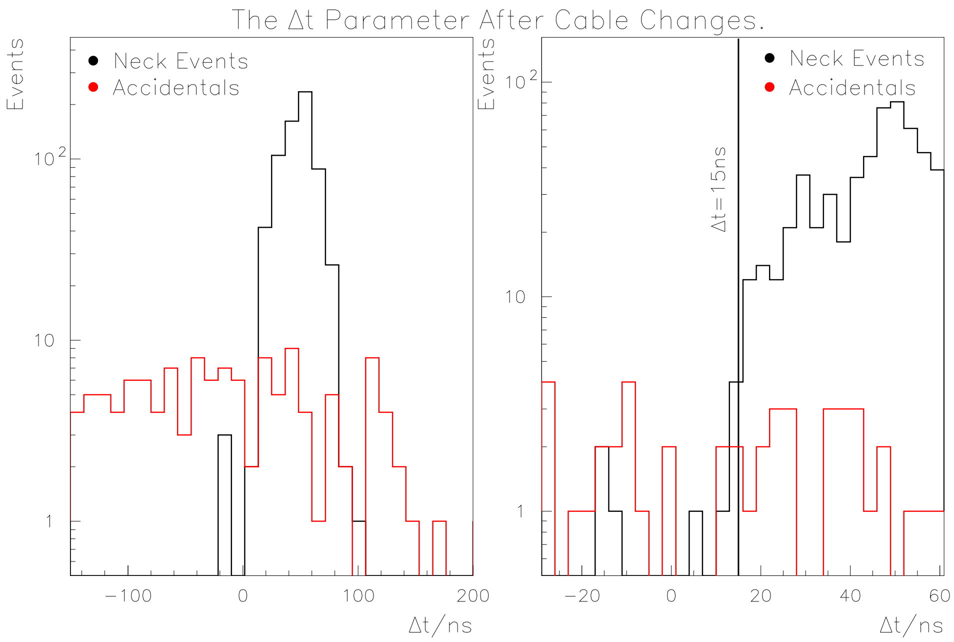

Figure 4.11: Time difference between the neck tubes and the mean PMT time for hits in the lower hemisphere, after the changes to the neck tube cabling.

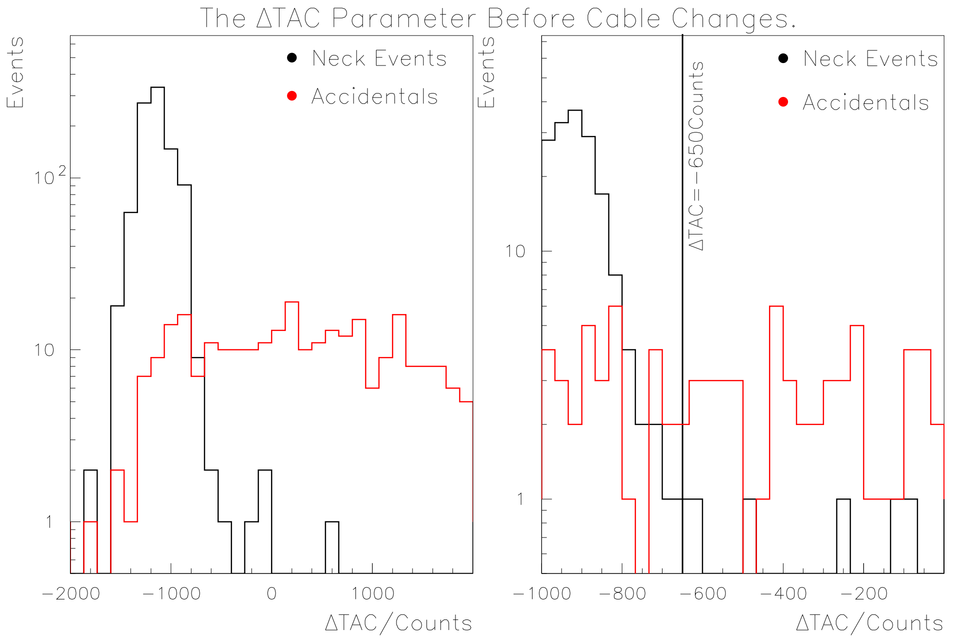

Figure 4.12: Raw TAC difference between the neck tube and the mean TAC of the PMTs at the bottom of the detector, before cable delays.

	Neck Events		*Accidentals*	
	QandT Mode	Neck Mode	QandT Mode	Neck Mode
QandT Mode	891/900	898/900	6/242	6/242
Neck Mode	898/900	893/900	6/242	1/242

Table 4.3: The correlation matrix for the two neck cut modes for identified neck events and accidentals.

Testing the QandT parameters

To test these derived parameters, the QandT mode of the neck cut was also run with a variety of different values for each parameter. Scans around the chosen value were made while keeping the other parameters constant, using both data sets and combining the results.

The minimum charge cut is stable in both sacrifice and contamination to large changes in the cut value, as is expected from the negative rail. The other parameters are not as stable to changes in the cut values, but this is a relative measure. The largest effect is seen in lowering the charge threshold where the sacrifice is increased. This is not a surprising effect as the cut is designed to be very aggressive. Overall the neck cut is not very sensitive to small changes in the cut parameters which is a requirement for a robust cut.

The interplay between the two modes

Each mode of the neck cut has its own particular strengths and weaknesses. These are best seen by examination of the correlation matrix for neck and accidental events, shown in table 4.3. This table shows that the two neck cut modes are very complementary with each other, with each mode removing some neck events that had leaked through the other. This demonstrates that the best way to use the two modes is as an “or” of the two cuts.

4.3.2 The OWL cut

The OWL cut is also a very effective cut at removing neck events. However unlike the neck cut it is not targeted specifically at this class of event. Like most of the instrumental cuts, the OWL cut is targeted at a number of background classes, including the non instrumental background of through going and partially contained muons.

The OWL cut uses the outward looking tubes in the SNO detector as a veto for events. SNO has two classes of these outward looking tubes and uses both categories. The OWL tubes comprise 91 PMTs looking outwards from the PSUP. These tubes were included in the original design for SNO as a veto for through going muons [32]. The

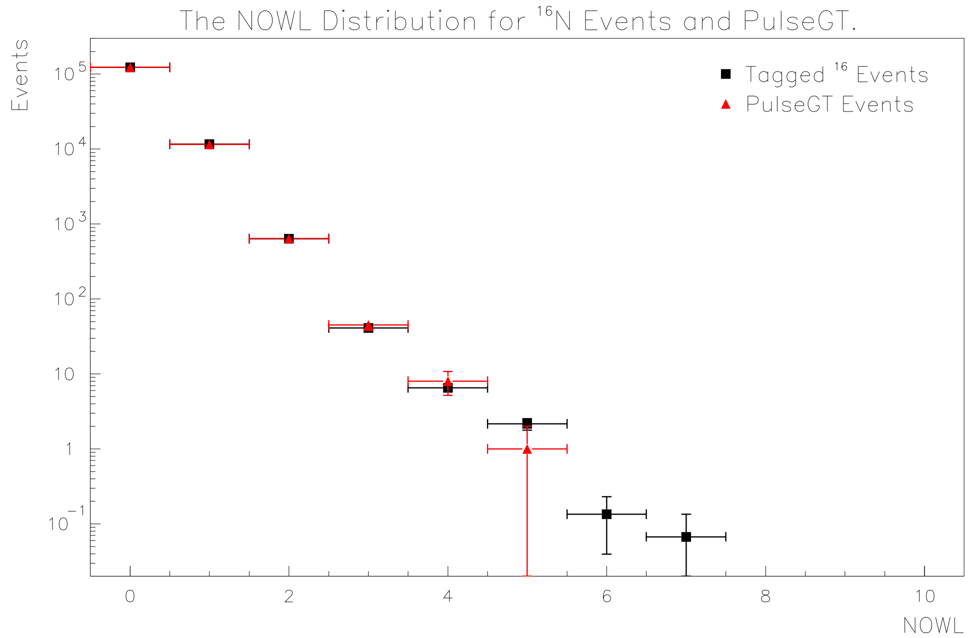

Figure 4.13: A comparison of the OWL+BUTT distribution for tagged ^{16}N and PulseGT events from the ^{16}N data cleaning runs.

BUTTs can also see outside the PSUP and in particular as some are pointing at the neck of the detector, these tubes turn out to be useful for removing instrumental backgrounds including neck events. The BUTTs are extra OWL tubes so the two sets are combined together when considering outward looking tubes. The total number of hit outward looking tubes in the event including the OWLs and the BUTTs is called NOWL.

The OWL cut uses a simple algorithm which removes events if NOWL exceeds or equals a threshold number. This threshold is not compared to the number of these tubes online, but this is an obvious improvement that could be made, although in practice given the final setting of the cut value it would have little impact. The OWL cut is unique amongst the instrumental cuts as the sacrifice of the cut can be monitored and measured exactly using PulseGT data. This is possible because the sacrifice from this cut is dominated by coincidences between noise hits on these tubes. This fact can be demonstrated by comparing the NOWL distribution for ^{16}N events and PulseGT events from the same set of runs as seen in figure 4.13. The ^{16}N runs used are the data cleaning data set discussed in chapter 5, and shown in table 5.1.

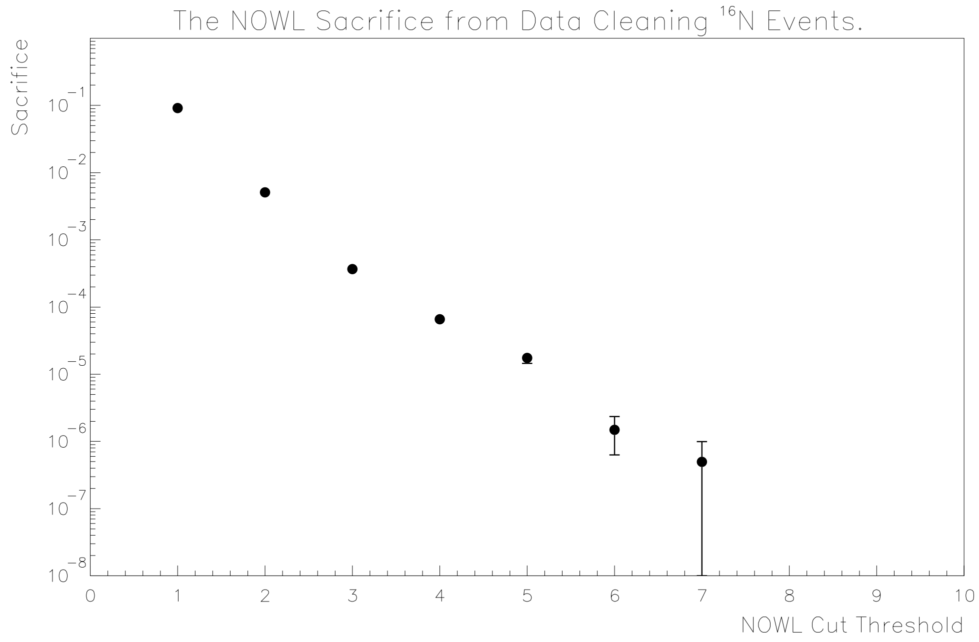

Figure 4.14: The imposed sacrifice from the OWL cut for various thresholds.

The OWL cut threshold

The OWL cut threshold is chosen with reference entirely to the sacrifice that would be caused by the cut value at the given threshold. This is possible because by lowering the threshold the rejection power of the OWL cut is always increased. Therefore the threshold of the OWL cut is set so that the sacrifice is as high as can be tolerated. The runs used to measure the sacrifice of the cut are the ^{16}N data cleaning runs. The sacrifice that would be imposed for a given NOWL threshold is shown in figure 4.14. The target for sacrifice for the entire suit of data cleaning cuts is less than 1%. For this reason the OWL threshold is chosen at three or more outward looking tubes, for which the sacrifice measured using the ^{16}N data is $(0.037 \pm 0.001)\%$.

The stability of this measurement is obtained via the PulseGT data taken with every run. The results of this study are shown in figure 4.15. The OWL sacrifice is not completely stable over time, with the data showing some quite large variations. These biggest deviation, seen after approximately 100 days of production running can be traced to the reduced number of OWL tubes operating at the time. The best fit to the OWL sacrifice using this data is 0.034%, slightly lower than the value measured by the ^{16}N. That data was taken approximately 30 days after the start of production running, so it agrees with the value measured at the time. The spread of the sacrifice around the best

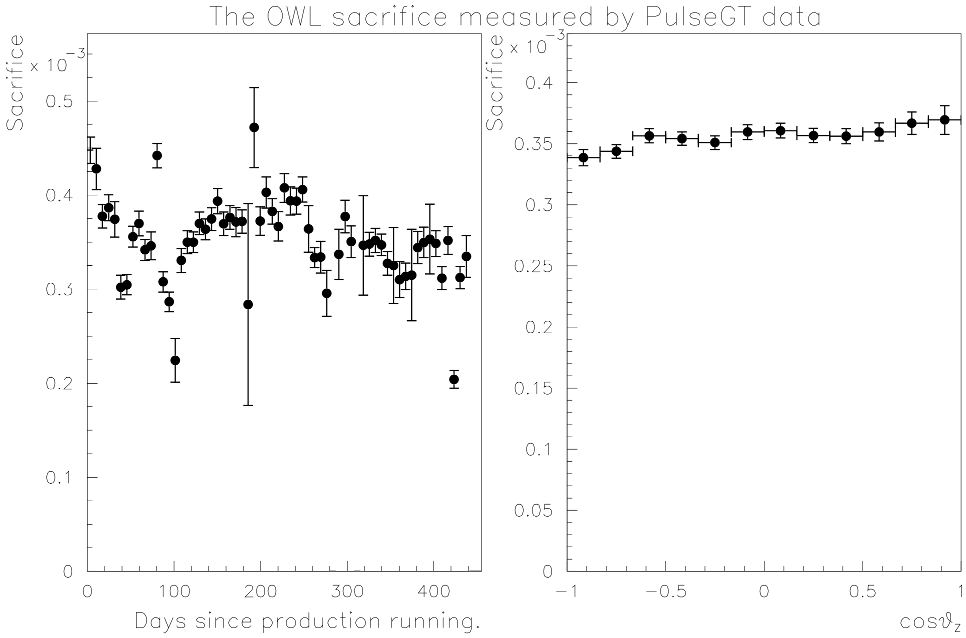

Figure 4.15: The stability of the OWL sacrifice as a function of time and of the direction to the sun.

	Neck Events	Accidentals
Passed	92	432
Failed	1522	19

Table 4.4: The performance of the OWL cut on the neck tube event sample used to develop the neck cut.

fit gives one contribution to the final measurement of the error on total sacrifice of the cuts. The median deviation from the fit is 0.003%, giving us the contribution of the OWL cut stability to the error in the total sacrifice.

The OWL cut also shows a weak dependence on the position of the sun, with there being slightly more sacrifice during the day than the night. This is as might be expected due to lab activity during the day. However the sacrifice difference between day and night is only $(9 \pm 4) \times 10^{-4}$%. This means that this difference cannot make a strong contribution to any measured day-night asymmetry.

The choice of this very aggressive threshold for the OWL cut means that the rejection power of the cut is very good. This can be seen by examining the effect of the cut on the neck cut data set shown in table 4.4. The OWL cut provides an effective way of removing neck events independently of the neck cut.

	Neck Events	Accidental Events
Passed	365	443
Failed	1249	8

Table 4.5: The performance of the OWLEHi cut on the neck tube event sample used to develop the neck cut.

4.3.3 The OWLEHi cut

The OWLEHi cut is a cut that removes all events that fire the OWLESUMHi trigger. This is the simplest of the cuts targeted at neck events. The OWL, BUTT and NECK tubes are all connected to the OWL trigger system, including the ESUM triggers. The OWLESUMHi trigger is always masked in and because of the high amount of charge deposited on the neck tubes, neck events often fire the detector via this trigger.

The neck events that fail the OWLEHi cut are expected always to fail the neck cut as well. However this cut provides some redundancy for that cut, protecting it against potential problems in the signal processing of the neck tubes, such as orphaning of the tubes, ADC failure, or the neck tubes becoming cut off at the start of the event. This was a particular worry before the changes were made to the neck tube cables.

Using the data sample used to designed the neck cut, the performance of the OWLEHi cut can be measured. Table 4.5 shows the performance of the cut on events classed as both neck events and other events.

4.3.4 The Burst Cuts

The instrumental cuts contain three separate burst cuts. The first of these is the retrigger cut, which operated as a tag for re-trigger events and is not relevant for consideration with neck events. There are then two more burst cuts that operate over more relevant time scales: the burst cut, and the NHIT burst cut. The burst cut will remove bursts of events where there are three or more events in 1ms. This cut is found to be particularly effective at removing bursts of flashers, but due to the time scales involved it is not expected to be effective against neck events. The NHIT burst cut operates over a much longer time scale and has been designed with neck events (amongst others) in mind. This cut removes bursts of 6 events in 4s after the application of a cut at 40 hits. Both burst cuts will not count events tagged as retriggers in the rate measurement. Burst cuts are discussed in more detail in chapter 6.

The ability of the burst cuts to remove neck events can be measured using the sample of neck events used to develop the neck cut. Events above 40 hits are considered separately from events below 40 as these events behave differently under the NHIT burst

	Burst Cut	NHIT Burst Cut
NHITs<40	28/135	86/135
NHITs≥40	43/1479	955/1479

Table 4.6: The performance burst cuts on the neck event sample used to develop the neck cut.

	Neck Cut	OWL Cut	OWLEHi Cut	NHIT Burst Cut
Neck Cut	1603	1520	1249	1039
OWL Cut	1520	1522	1249	989
OWLEHi Cut	1249	1249	1249	746
NHIT Burst Cut	1039	989	746	1041

Table 4.7: The performance of the various cuts designed to remove neck events on the neck sample used to determine neck cut parameters. There are 1614 neck events used in the creation of the table.

cut. Events below 40 hits can however be removed by the NHIT burst cut, if they occur inside a burst with higher NHIT events. The results of the burst cuts on the data in the neck cut sample is shown in table 4.6. This demonstrate clearly that the NHIT burst cut is a much more effective cut for removal of neck events than the burst cut. This also demonstrates that the NHIT burst cut is good at removing neck events, even below 40 hits. The NHIT burst cut provides good security against the occurrence of a problem with the neck tubes.

4.3.5 Interaction between the cuts

Each of the cuts discussed above is good at removing neck events. However it is the combination of these cuts that is ultimately of importance in this process. The correlation matrix of the cut fractions, when applied to a sample of neck events tells us how these different cuts interact, shown in table 4.7.

The correlation matrix shows that the overlap between the cuts on the neck sample is large. This is of no surprise as each cut removes by design a large fraction of neck events. The correlation matrix does however show that the neck, OWL and NHIT burst cuts do remove neck events that are missed by the other cuts. This shows that the cuts are working in a complementary way increasing the total fraction of neck events removed. The correlation matrix also confirms as already discussed, that the OWLEHi cut does not remove neck events that are missed by the neck or the OWL cuts. This does not however remove its usefulness as a backup to these more powerful cuts.

Data Cleaning cut	Passed	Failed
Retrigger Cut	1614	0
Burst Cut	1543	71
QvT Cut	1591	23
Q/NHIT Cut	1614	0
Crate Isotropy Cut	1614	0
AMB Cut	342	1272
FTS Cut	418	1196
OWL Cut	92	1522
Junk Cut	1614	0
Neck Cut	11	1603
ESUM Cut	1614	0
QCluster Cut	1416	198
ITC Cut	1601	13
FGC Cut	1293	321
NHIT Burst Cut	573	1041
OWLEHi Cut	365	1249

Table 4.8: The performance of the various instrumental cuts on the neck event sample.

4.3.6 The effect of the other instrumental cuts

As well as the cuts aimed at neck events the other instrumental cuts will also remove neck events. The effect of each cut in the suite of cuts is shown in table 4.8 As well as the cuts already discussed there are two cuts that were not designed with neck events in mind that do a good job in their removal. The first of these is the AMB cut. However this turns out not to be a powerful extra method to remove neck events. This is because the AMB cut only removes events at high NHIT or where the event has triggered the detector via the OWLEHi trigger. When this happens the ESUM signal measured by the AMB board does not contain the real PMT hits which happen later in the event. This means that the expected signal for the NHIT of the event does not match expectations which results in the event being cut.

The other cut that removes significant numbers of neck events is the FTS cut. This cut does provide extra security against neck events as it cuts them in a different way to the cut previously discussed, using the relative timing of the tubes in the event. An examination of the FTS cut parameter shows why the cut is so effective. The FTS distribution compared to that from ^{16}N is shown in figure 4.16. This distribution shows, that like the flashers, for which the FTS was designed, neck events have much broader pulse shapes than Čerenkov light.

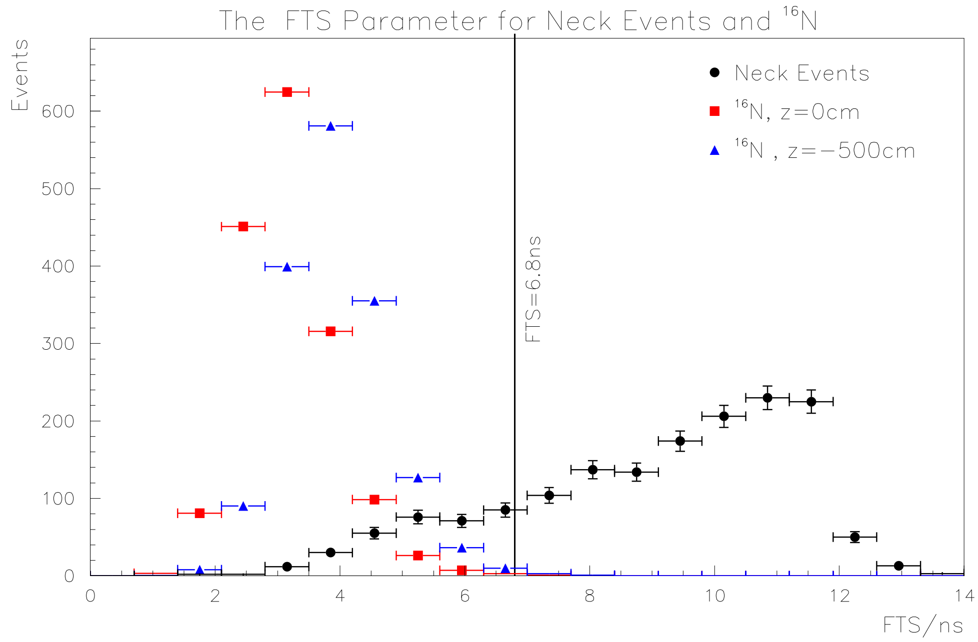

Figure 4.16: The FTS distribution for neck events and ^{16}N data.

4.4 Neck Contamination

We have now produced a set of cuts that remove neck events from the data set. However it is unlikely that the number of these events will be reduced to zero. This means that it is necessary to attempt to measure the residual content of neck events in the data after the application of instrumental cuts, even if only to show that it is compatible with zero.

The one property of neck events that has not been used in order to remove them is their distinctive hit pattern. As the light comes from the neck the tubes at the bottom fire preferentially compared to those elsewhere. Using this information we can use pattern recognition to identify candidate neck events based on their hit pattern. The method chosen in this study is to use a likelihood ratio technique. We use the hit pattern from a large number of neck events to construct a likelihood function for a neck event. A likelihood function can also be constructed from a number of PCA runs. Using these two likelihood functions a likelihood ratio can be constructed to attempt to separate Čerenkov events from neck events.

4.4.1 The Likelihood Ratio

The Neck Event Dataset

The first step required is to create a set of neck events. The procedure for doing this using the neck cut as described in section 4.2. However, in the construction of the likelihood function the entire dataset is not used. Data from the selected neutrino run list is used, with the list truncated at run 11446. This means that it will be possible to use this function to measure neck contamination on data that has not been used to construct the PDFs. As the selection method uses the neck tubes to select the neck events, the selection method is independent of the proposed likelihood method which will rely on the hit pattern.

Finally in the creation of this data set an NHIT cut is also applied. Initially this is set above 35 hits. This runs in conjunction with a maximum NHIT cut, set to 500 hits. This creates a maximal neck event data set that can be reduced depending upon the requirements of the task.

The Likelihood Function

The likelihood function used in this method is constructed by using the hit probability per tube in each event. This hit probability must first be tabulated using the neck event sample described above. This is done using a SNOMAN USR processor [41], [42].

The number of hits in each PMT is recorded along with the number of events where that PMT was enabled, checked using the DQxx banks. The tube-by-tube probabilities are then renormalised so that the sum over the detector equals one. These probabilities are then written out to the titles bank TNEC 1. To create the normalising likelihood function the same procedure is used, except that instead of neck events PCA runs are used instead. The probabilities are here written to a TNEC 2 bank.

The choice of PCA runs for the normalising likelihood may, at first, seem unusual as ^{16}N calibration data would at first glance be a more logical choice. However in order to avoid position dependence, due to the limited number of source locations available, a run at the centre is the only available choice. With the run confined to the centre, the constructed likelihood should be isotropic after the summing of events. The only factors that can change this assumption are optical effects in the detector. The use of a single ^{16}N run, would not accurately represent the state of the detector as a function of time. This is not possible even with a combination of ^{16}N runs, as they do not span the data taking time very well, particularly for the early data used in the construction of the likelihood ratio. PCA runs on the other hand do span the span the data set well. They should also contain the deviations expected due to the optics, with the possible exception

Run Number	Date Taken
10623	12/09/99
10863	01/18/00
10918	01/25/00
10988	02/07/00
11326	02/21/00
11456	03/23/00
11663	04-26-00
11795	05-18-00
12119	06-26-00
12251	07-27-00
12477	08-18-00

Table 4.9: The PCA runs used to construct the normalising probabilities.

of wavelength dependent effects. Thus to create the normalising probabilities the first subrun of the available PCA runs were used, these runs are shown in table 4.9

The likelihood of an event using this method is given by:

$$\mathcal{L}^{\mathrm{neck}} = \prod_{i=1}^{\mathrm{NHITs}} p_i^{\mathrm{neck}} \tag{4.1}$$

This function as with all likelihoods becomes easier to handle by taking the logarithm of both sides.

$$\log \mathcal{L}^{\mathrm{neck}} = \sum_{i=1}^{\mathrm{NHITs}} \log p_i^{\mathrm{neck}} \tag{4.2}$$

Likewise the logarithm of the normalising likelihood function is given in the same way. The logarithm of the likelihood ratio is then given by the difference of the two functions. The distribution of this function is shown in fig 4.17. The profile histogram shown in fig 4.18, shows the likelihood ratio as a function of NHITs. This shows that the likelihood ratio is proportional to NHITs when the likelihood ratio is greater than one, i.e. the logarithm is greater than zero. However this is not unexpected as the log likelihoods involve a sum over hit tubes. The structure at small values of the likelihood ratio is due to contamination of the neck event data set by low energy Čerenkov events sacrificed by the neck cut. The dependence on NHITs can be removed by dividing the logarithm of the likelihood ratio by NHITs. This function called R_{neck} (equation 4.3) is shown in fig 4.19 and can then be used to help separate neck and Čerenkov events.

$$R_{\mathrm{neck}} = \frac{\log \mathcal{L}^{\mathrm{neck}} - \log \mathcal{L}^{\mathrm{norm}}}{\mathrm{NHITs}} \tag{4.3}$$

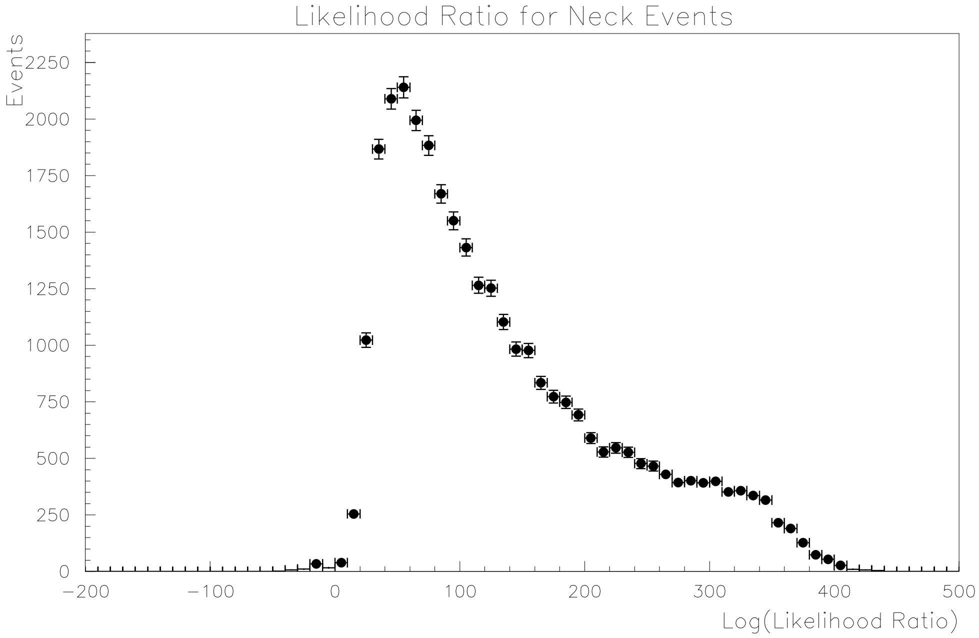

Figure 4.17: The logarithm of the likelihood ratio, shown for neck events.

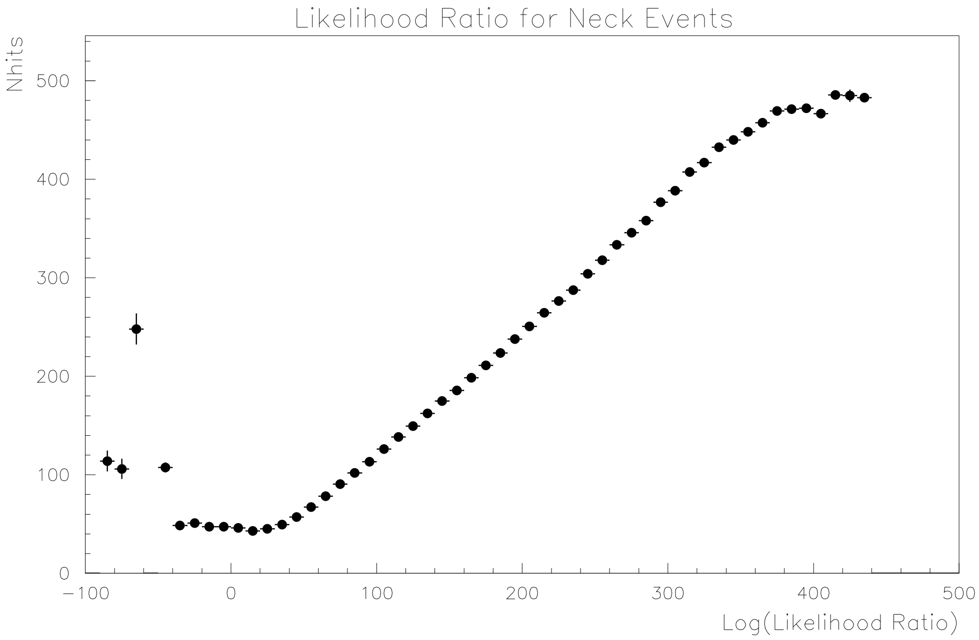

Figure 4.18: A profile histogram showing the dependence of the logarithm of the likelihood ratio on NHITs.

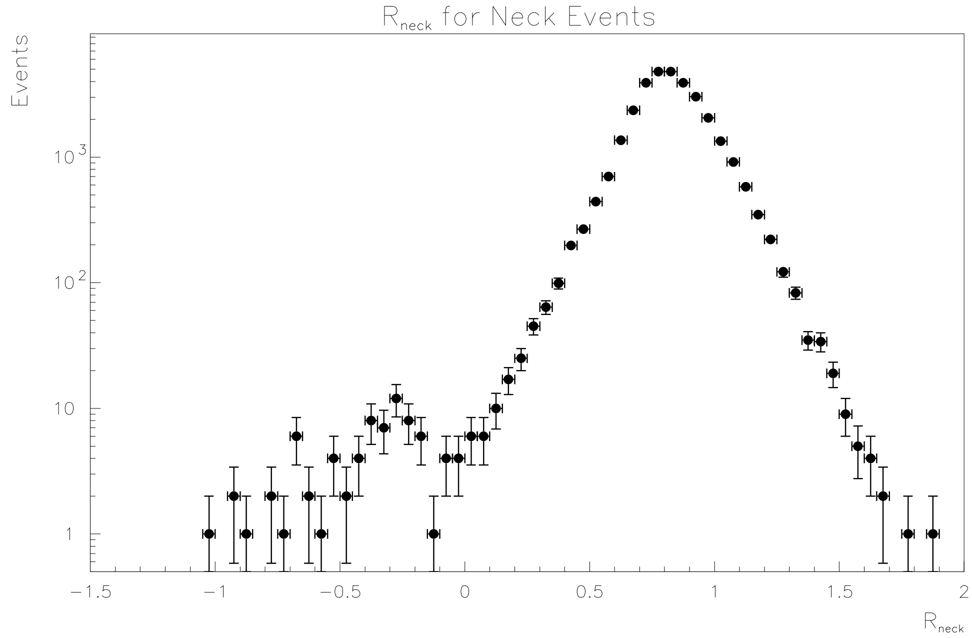

Figure 4.19: R_{neck} for neck events.

These three figures all show that there are some events in the region $R_{\text{neck}} < 0$. These events according to this algorithm are not neck events, and can therefore be used to estimate the sacrifice of the neck cut, see section 4.5.

4.4.2 The Likelihood Function and Čerenkov Events

^{16}N Data

Due to the way that hit patterns are used it is expected that there will be some overlap when the event occurs at the bottom of the detector. The extent to which this is true can be tested using ^{16}N data. Figure 4.20 shows the R_{neck} distribution for ^{16}N events from three different source positions along the z axis. This shows that there is considerable overlap between the neck event and ^{16}N regions particularly when the source is at the bottom of the detector. This is not unexpected. However there are regions where the overlap either disappears or is small. This will therefore allow R_{neck} to be used to estimate contamination due to neck events.

To better measure the overlap of the neck event and ^{16}N regions a data set of ^{16}N was created using equal numbers of events from different source positions. As many source positions as possible were used, these are shown in table 4.10. The ^{16}N data set was then constructed using 20000 events from each of these runs, with the resulting R_{neck}

Run	*Source Position*			*Run*	*Source Position*		
	x/cm	y/cm	z/cm		x/cm	y/cm	z/cm
10295	0	0	-530	10298	516	0	0
10301	300	0	0	10304	400	0	0
10307	200	0	0	10311	100	0	0
10314	-300	0	0	10317	-500	0	0
10320	-200	0	0	10323	-400	0	0
10326	-100	0	0	10328	220	0	220
10331	140	0	140	10333	70	0	70
10335	-70	0	-70	10337	-140	0	-140
10339	-220	0	-220	10341	-310	0	320
10353	0	0	100	10355	0	0	200
10357	0	0	300	10359	0	0	400
10361	0	0	500	10374	-220	0	220
10376	-140	0	140	10378	-70	0	70
10380	70	0	-70	10382	140	0	-140
10384	220	0	-220	10386	280	0	-280
10388	325	0	-330	10405	0	-300	-300
10407	0	-220	-220	10409	0	-140	-140
10411	0	-70	-70	10413	0	70	70
10415	0	140	140	10417	0	180	180
10453	0	220	-220	10455	0	140	-140
10457	0	70	-70	10459	0	-70	70
10461	0	-140	140	10463	0	-180	180
10480	0	0	-100	10482	0	0	-200
10484	0	0	-300	10486	0	0	-400
10488	0	0	-500	10490	0	0	0
10517	0	0	600	10564	-586	208	-100
10566	-586	208	-200	10568	-586	208	-300
10571	-586	208	-400	10573	-586	208	-500
10575	-586	208	0	10577	-586	208	100
10579	-586	208	200	10581	-586	208	300
10583	-586	208	400	10585	-586	208	500

Table 4.10: The ^{16}N runs used to construct the data set.

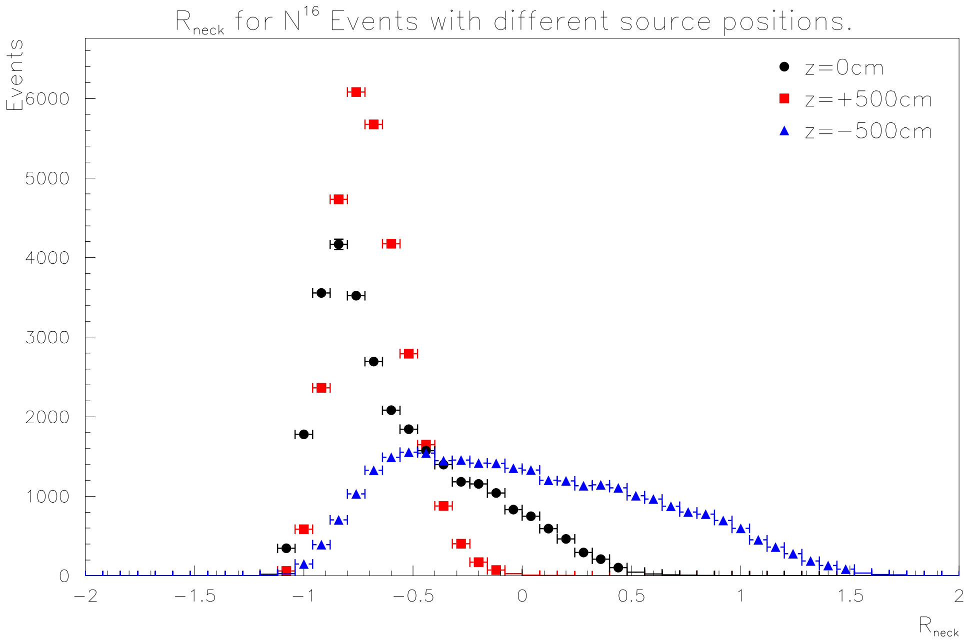

Figure 4.20: R_{neck} for three different ^{16}N runs.

distribution is shown in figure 4.21.

Using the Monte Carlo

Even using the constructed ^{16}N data set does not tell us the full story however, as these events are still not distributed isotropically about the detector. It also does not tell us about any energy effects in this method, although these would be expected to be small at these energies. Therefore to apply some further checks the Monte Carlo can be used.

To do this mono-energetic electrons are generated throughout the D_2O, and R_{neck} is calculated. This has been done for three different electron energies, 5, 10 and 15 MeV. The resulting R_{neck} distributions are shown in figure 4.22. This plot shows that the R_{neck} distribution is independent of energy in the region of interest for studies of solar neutrinos.

A comparison of the R_{neck} distributions between Monte Carlo and the ^{16}N data set can also be made. This allows us to examine the effect of the non-isotropic distribution of events in this data sample. This comparison is shown in figure 4.21. There is good agreement between the two distributions, suggesting that the non-isotropic nature of the ^{16}N sample does not have a large impact on the measured R_{neck} distribution, and that R_{neck} is predicted well by the Monte Carlo.

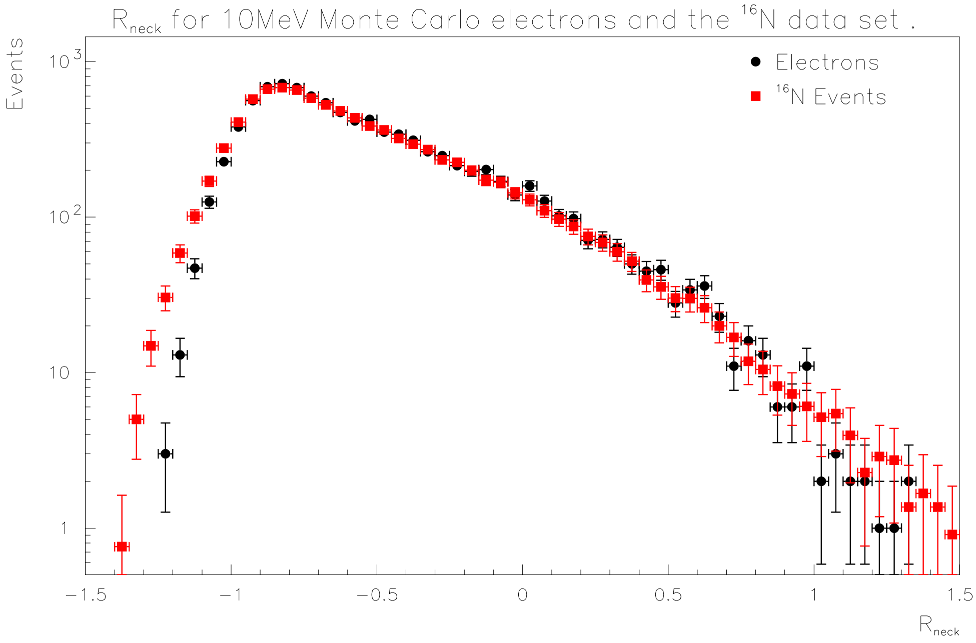

Figure 4.21: R_{neck} for Monte Carlo 10MeV electrons and the ^{16}N data set.

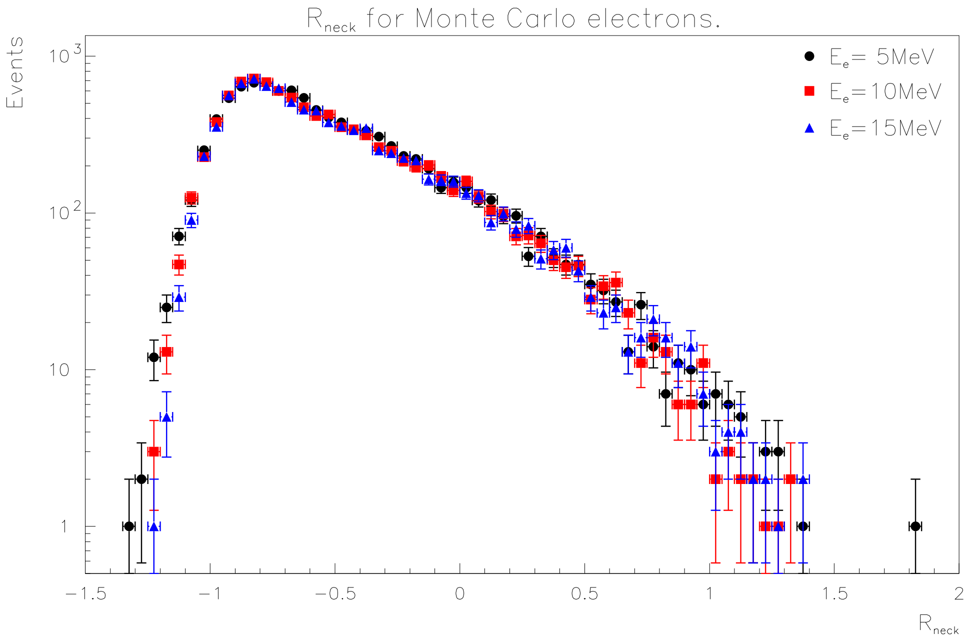

Figure 4.22: R_{neck} for Monte Carlo electrons at different energies.

4.4.3 Estimating Neck Contamination Using R_{neck}

The search for neck events after the application of the cuts, is a rare process experiment, provided that the cuts have done a good job. In order to estimate neck event contamination, we use a technique that is used in other rare process experiments. We construct a neck event signal box in which to look for events. To carry out the contamination measurement it is necessary to know the number of Čerenkov events that will leak inside the box, and the number of neck events that will leak out. The number of events inside the box can then be combined with these probabilities to measure the contamination.

Because there is some overlap between the Čerenkov and neck event regions, R_{neck} cannot be used to define the signal box alone. For this reason the other differences between neck events and signal events must be used to enhance the separation of the two event classes. The most important feature of neck events that can be used for this enhancement are the reconstruction properties (see section 4.2). The conclusion of this discussion is that neck events tend to reconstruct close to the z-axis, but with little structure along that axis. These properties can be used to define an analysis cylinder inside the detector in which to search for neck events. We then need to select the radius and lower boundary in z for this region. There is no need to select an upper boundary as it has been shown that the overlap region in R_{neck} occurs in the lower half of the detector.

The probability for signal events to pass the cut inside the cylinder can be found from the Monte Carlo. ^{16}N data cannot be used as following the use of a reconstruction cut it is important that the distribution of the signal be accurately modelled. To achieve this a sample of 5×10^5 5MeV electrons are generated throughout the D_2O and H_2O. This models the distribution of events present after instrumental cuts more accurately than the ^{16}N data set, particularly as this data is dominated by H_2O $\beta\gamma$s. The probability for a neck event to pass the cuts can be determined from a sample of neck events. This sample of neck events must pass all instrumental cuts that rely on the central PMTs, with the exception of events that fail AMB and OWLEHi (see section 4.3.6).

The choice of the size of the cylinder becomes a trade-off between the number of events expected to pass the cuts, and the sensitivity to neck events. Most of the sensitivity of the two probabilities lies in the choice of the low edge of the cylinder. For this reason the choice of R_{neck} cut and cylinder radius are made beforehand, and the survival probabilities are studied as a function of this low edge. Using figure 4.19 the R_{neck} cut is chosen at $R_{\text{neck}} > 0.5$, and using figure 4.4 the choice $\rho < 250cm$ is made. The survival probabilities of the two event classes after these cuts are shown in table 4.11. The events remaining after these cuts can then be used to determine the optimal position of the bottom of the cylinder, z_{cut}, as shown in figure 4.23.

The data after instrumental cuts contains 14203 events inside the NHIT range.

Cut	Neck Events	Monte Carlo
45≤NHITs≤200	1.00	1.00
Event Reconstructs.	0.83 ± 0.01	0.9515 ± 0.0005
R_{neck} >0.5	0.70 ± 0.01	0.0339 ± 0.0004
$\rho < 250cm$	0.71 ± 0.01	0.167 ± 0.001
Combined	0.58 ± 0.01	0.0145 ± 0.0003

Table 4.11: The Fraction of surviving events after different cuts on neck events and Monte Carlo electrons.

With this number of events a survival probability for Čerenkov events of less than 10^{-4} is desirable. Given the statistics of the Monte Carlo this target is at the limit of possible measurements, but it can be made. For this reason the z_{cut} is chosen at -125cm.

With the choice of z_{cut} made, the final survival probabilities can be measured. The neck selection efficiency is: 0.357 ± 0.006. The probability of Čerenkov events surviving is $(2.0 \pm 0.5) \times 10^{-5}$, which results in an expected number of events inside the neck event box of 0.29 ± 0.07 events. The cuts are then applied to the neutrino data after instrumental cuts, and eight events are seen inside the neck event box. After correction for the expected number of events and the neck selection efficiency there are 21.6 ± 7.9 neck events remaining after instrumental cuts. The error is dominated by the statistical error on the number of events seen in the neck event box. This number can be converted to an estimate of neck contamination of $(0.15 \pm 0.06)\%$, or a three σ upper limit of 0.33%.

The high level and figure of merit cuts are also very effective at removing neck events. This can be seen by the reduction of the neck event data sample to three events after these cuts are applied in combination with a fiducial volume cut. This reduction makes it difficult to repeat the measurement of neck contamination after these cuts are applied. However it is possible to show that the potential neck events missed by the instrumental cuts are not also missed by the high level cuts. This can be achieved by examination of the neck event box after the application of the high level cuts. Once these are applied there are no remaining events inside this region. This observation shows the additional removal power of the high level cuts against neck events, and that the contamination of neck events in the final data sample is very small.

4.5 Estimating the sacrifice of the neck cut

The neck cut relies on measurements in the neck tubes. In general the sacrifice of a cut is measured using calibration data, see chapter 5. However in the case of the neck cut this measurement could be suspect due to shielding effects from the source, or increased light

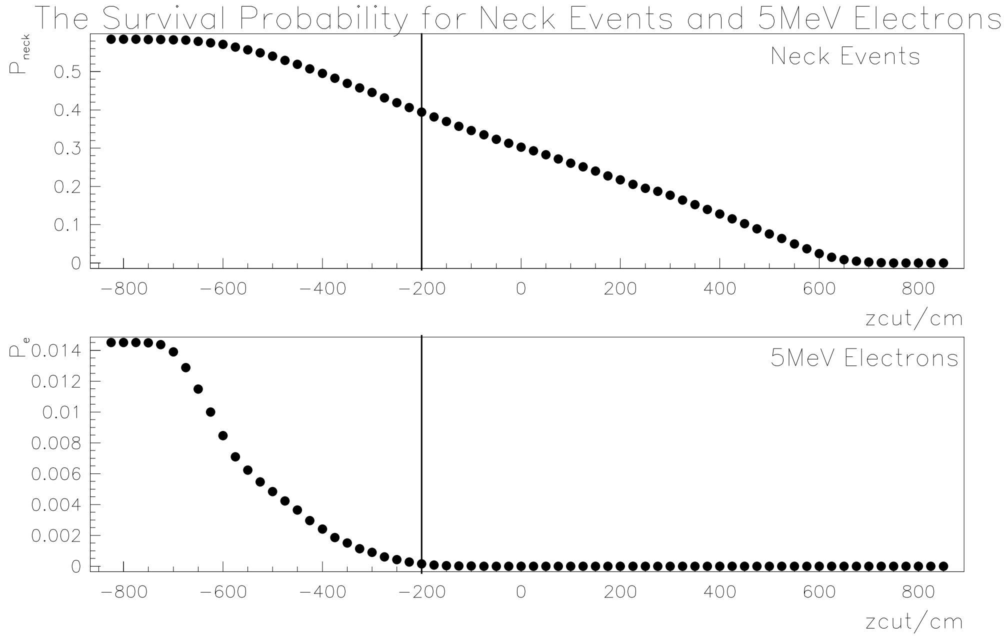

Figure 4.23: The survival probabilities for neck events and Monte Carlo electrons for a variety of z_{cut} values.

levels in the neck due to the source deployment mechanism. It is therefore desirable to confirm that the sacrifice measurement made using calibrations sources is accurate.

As well as using the $R_{\rm neck}$ parameter to measure contamination, it can also be used to estimate the sacrifice of the neck cut. When measuring the contamination the $R_{\rm neck}$ parameter was only considered above 35 hits. Examination of the $R_{\rm neck}$ distribution below this cut reveals a large number of events that do not look like neck events. It is thus possible to use these events to measure the sacrifice of the neck cut directly, to check the measurement made by calibration sources and to monitor the sacrifice of the cut.

The $R_{\rm neck}$ distribution of events cut by the neck cut below 35 hits is considerably different to that above 35 NHITs. It shows a high proportion of the events with $R_{\rm neck}$<0. The events at $R_{\rm neck}$<0 do not have the characteristic event pattern expected of neck events, and can be selected as events that are sacrificed by the neck cut and used to measure the sacrifice of the cut directly. The $R_{\rm neck}$ distribution of events below 35 hits failing only the neck cut is shown in figure 4.24

As shown in figure 4.22 the $R_{\rm neck}$ distribution of Čerenkov events is dominated by events at low $R_{\rm neck}$, so placing a cut at $R_{\rm neck}$<0 will not remove a large fraction of Čerenkov events. By placing this cut on both the sample of events failing the neck cut and on those that do not fail any cut, the neck cut sacrifice can be measured directly from

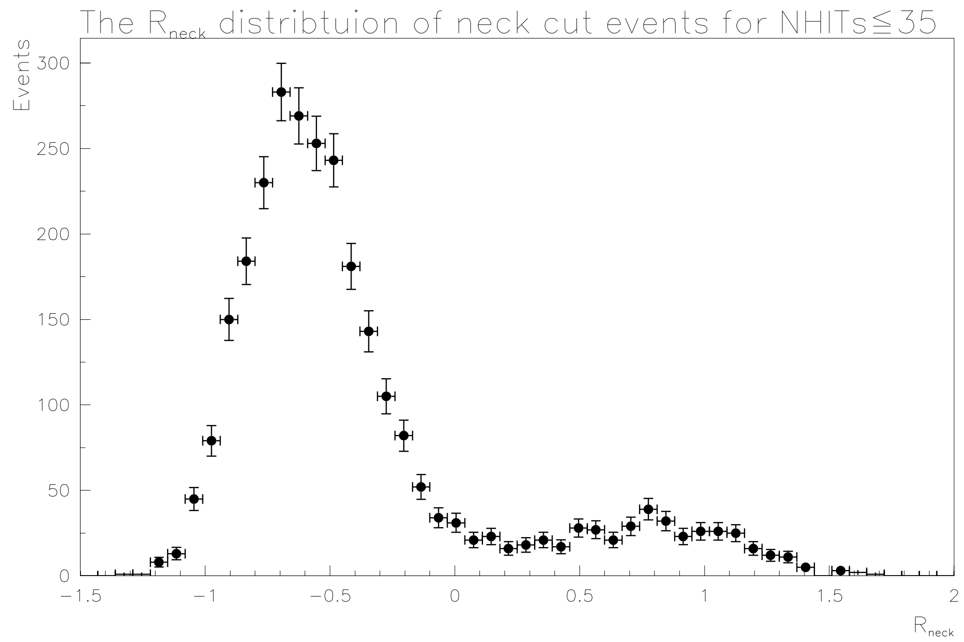

Figure 4.24: The $R_{\rm neck}$ distribution for events failing only the neck cut between 20 and 35 hits.

the data. The selection effect using the $R_{\rm neck}$ cut is removed by measuring the sacrifice in this way, as the effect cancels in the calculation of the ratio.

The sacrifice calculated using $R_{\rm neck}$ and using ^{16}N data in the usual way are compared in figure 4.25. There is good agreement between the measured sacrifice from ^{16}N and $R_{\rm neck}$. This suggests that the source present in the ^{16}N runs is not having a big effect on the measurement of sacrifice of the neck cut. Other sources are not shown as there are insufficient statistics in this low NHIT region with which to make a comparison.

The measured sacrifice is small, at 3×10^{-5}. This is much smaller than the sacrifice of some of the other instrumental cuts, which suggests that neck cut does not contribute significantly to the total sacrifice. It would be nice to confirm the low sacrifice at higher NHIT, but the statistics are too small to use the $R_{\rm neck}$ method. Indeed above 45 hits no events can be identified as neck cut sacrifice, allowing an upper limit of 2×10^{-4} to be set at the 90% confidence level.

The $R_{\rm neck}$ method also allows the sacrifice of the cut to be monitored. It is important to do this both as a function of time and as a function of the position of the sun. These are shown in figure 4.26. Both of these functions show that the neck cut sacrifice is not a strong function of time. The value of the neck cut sacrifice is small enough that the small deviations from the best fit value are much smaller than the other systematics associated

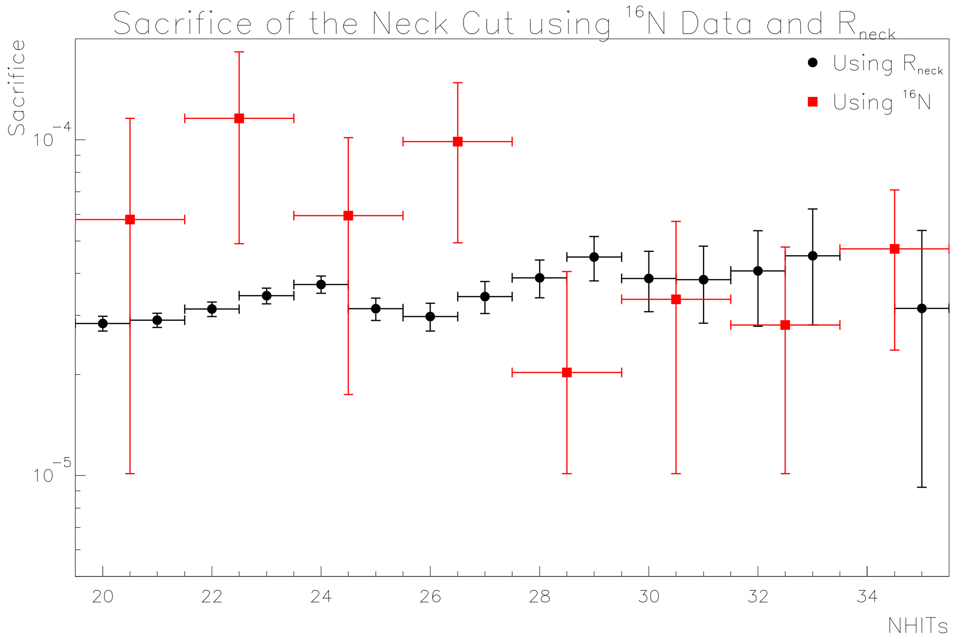

Figure 4.25: The sacrifice of the neck cut calculated using ^{16}N data and the R_{neck} parameter.

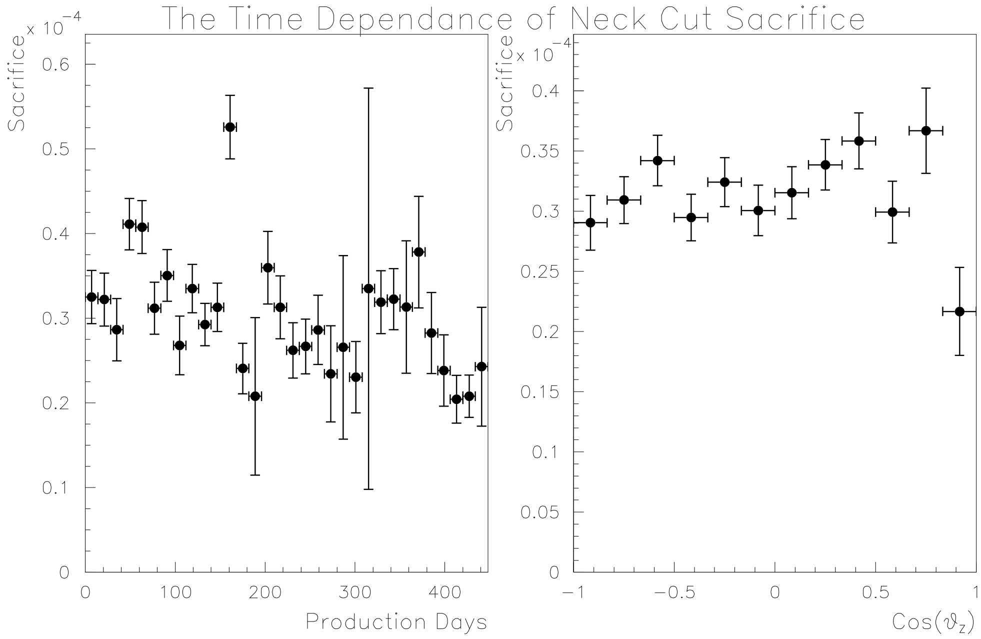

Figure 4.26: The sacrifice of the neck cut as a function of time and as a function of the solar zenith angle.

with the measurement of the total sacrifice. This means that these deviations can be safely ignored in this measurement and its associated error.

4.6 Conclusions

Neck events are one of the most serious instrumental backgrounds seen in SNO. The instrumental cuts used in SNO to remove them are varied, some relying on the light seen outside the PSUP, in these events, but some using the internal properties as well. The most important instrumental cut for removal of these events is the neck cut, which uses the PMTs installed in the neck to tag neck events. This cut can be used to create neck event data sets which can then be used for further study of these events.

Using a sample obtained in this way it is possible to construct a further neck event identifier R_{neck} which identifies neck events by the pattern of light created on the PSUP. The R_{neck} parameter can be used in two important ways to check the results of the instrumental cuts on neck events. The first and most important of these is to measure the contamination of neck events in the neutrino data after the application of the instrumental cuts. This is achieved by combining the R_{neck} parameter with the unique reconstruction properties of neck events to create a neck event signal box, in which to look for these events. This procedure shows that contamination after instrumental cuts is $0.15 \pm 0.06\%$, or less than 0.33% as a three σ upper limit. Once the high level cuts are applied the events in the neck event box are removed suggesting that the contamination is considerable reduced by these cuts.

As the neck cut is designed to identify neck events the R_{neck} parameter can be used to measure and monitor the sacrifice of the cut. Events that are cut by the neck cut, but do not look like neck events by R_{neck} can be identified as neck cut sacrifice, and the sacrifice can be determined. Unfortunately this only works at very low NHIT, however the measurements verify the sacrifice results for this cut from ^{16}N data. This technique can also be used to monitor the neck cut sacrifice and it is found to be stable.

This chapter has dealt with the analysis of neck events. This included determination of the sacrifice of the cuts and the residual contamination of neck events after their application. These two properties must be understood for the data cleaning analysis as a whole, not just for the neck events. This analysis will be presented in the next chapter.

Chapter 5

Sacrifice and Contamination.

The analysis of instrumental backgrounds has three major deliverables at the end of the study. The first is a set of cuts that can be used to obtain a neutrino sample from raw data. The second is the fraction of neutrino events that do not get into that neutrino sample: the sacrifice of the cuts. The final deliverable is the residual content of instrumental background within the neutrino sample: the contamination. Not only are the sacrifice and contamination important for any solar neutrino analysis, they also demonstrate a much more simple fact: they can be used to show that the data cleaning process works. This chapter will discuss the analysis that goes into the measurement of the sacrifice and the contamination.

5.1 Cut sacrifice

The definition of the sacrifice is a simple one:

> *Given a sample of neutrino events, the sacrifice is the fraction of events that fail the cuts applied to that data sample.*

This simple definition hides a complex statistic, because in SNO the sacrifice can be a function of any variable one might wish to use. As SNO carries out signal extraction using the distribution of neutrino data in energy, position, and direction, the sacrifice as a function of these variables must be understood. The neutrino signal is also studied as a function of time, which means sacrifice must also be understood in this dimension. Finally SNO sees a variety of different signals; the events from these signals can have different characteristic distributions requiring an understanding of the sacrifice for each class of event.

In the analysis of SNO data there are two distinct types of cut used, that impose sacrifice on the data in different ways. One type of cut examines each event individually and will remove the event based upon some characteristic of the individual event. The second type of cut works in a different fashion: each event is considered as part of a larger data set. Cuts then search for this data set and remove periods of time that fail some criteria. These cuts do not impose sacrifice in the same way as cuts that use individual events, they instead impose livetime corrections to the data set. The study of the sacrifice of these cuts is discussed in chapter 6, where the livetime of the experiment is considered. The sacrifice of cuts based on individual events forms the main body of this chapter.

These cuts each have their own sacrifice, which can be a function of energy (or NHITs), position and direction. Depending upon the robustness of the cut the sacrifice can also be a function of time. The sacrifice imposed by different cuts can also be correlated. These effects must be understood in order to reach the final goal of the sacrifice measurement: the sacrifice on the flux of solar neutrinos. This is achieved by an integration of the sacrifice over the signal PDFs used in the signal extraction. SNO also attempts to measure the energy spectrum of the solar neutrinos, and therefore the energy dependences of the sacrifice must be understood to reduce the systematic error on any measured spectrum.

5.2 Sources for measurement of sacrifice

5.2.1 Requirements for data used to measure sacrifice.

Unfortunately because the mechanisms of the instrumental background in SNO are poorly understood the detector distributions used to remove them are rarely modelled in the Monte Carlo in a way that allows us to study sacrifice. For this reason it is necessary to study sacrifice using calibration data, as well as the Monte Carlo where applicable.

A calibration source has certain requirements that it should meet before it can be used to study sacrifice. None of the requirements are absolutely essential, but if the requirement is not met the the quality of the measurement of sacrifice using the source will not be as good as it otherwise could have been.

- The source should provide tagged events.

  This is very important. If the events are tagged then each event is known to be a Čerenkov event. This allows a precise measurement of the sacrifice to be made, instead of just allowing an upper limit to be set.

- The source should emit Čerenkov light.

  We wish to measure the sacrifice of Čerenkov events. Many of the developed cuts use the difference between Čerenkov light and instrumental backgrounds. Use of a source that does not emit Čerenkov light introduces an unwanted systematic at best and produces an entirely incorrect result at worst.

- The source should provide sufficient statistics.

  In order to measure the shape of the sacrifice accurately a large number of events is required, much more than the expected number of neutrinos. High statistics usually means high rate. Given the rate effects seen in SNO this is undesirable. These effects can be tested using the same source at low rate, to measure the size of these effects.

- The source should be able to be deployed around the detector.

  Because sacrifice is a function of position and direction being able to move the source around the detector, including the H_2O, allows the shape of the sacrifice to be measured more accurately.

- The source should cover the energy range of interest.

  The energy range of interest here is the solar neutrino energy range, as this is the region that the instrumental cuts have been designed to work over. It is important that the entire range be covered, particularly when looking for shape distortions, or neutrinos from the hep reaction.

- The source should not change the events.

  The ideal source to measure sacrifice would be a containerless one. This of course is not compatible with the requirement for the tagged source, and is not even possible for many sources. Given these problems it is desirable that the container have a minimal effect on the data taken. These problems can include the blocking of backward light effecting geometry-based cuts, the slight change to events when a source is triggered as well as tagged, or the presence of the source deployment mechanism in the neck blocking the neck tubes.

In practice there is no one source that meets all of the above criteria, so the sacrifice has to be measured using a variety of sources. The use of calibration source data has problems associated with it. The most important of these is that the last criteria is never completely met as the source is always present. Source data is only ever taken at long time intervals[1], so it is difficult to use this data to study the time dependence of the effect of the cuts. Variations in the tube noise rate or calibrations can, and do, affect certain cuts. These variations can occur on much shorter timescales than those between source deployments.

[1] One common problem is that the source is usually deployed just after repairs have been made to the electronics rather than before.

In order to attempt to address these problems the Monte Carlo can be used to examine the cuts for which it is applicable. Monitoring PulseGT data as a function of time can also be used to help monitor the performance of some cuts as well. Finally attempts to measure sacrifice from cut data can also be made to verify the measurements of sacrifice. It is however to calibration source data that we turn first in order to begin to understand the effect of the cuts on Čerenkov data.

5.2.2 The ^{16}N source.

By far the most useful source for sacrifice studies is the ^{16}N source[73]. This source provides almost all the required information with the exception of sacrifice at high NHIT, as the source corresponds almost completely to a single energy (6.1 MeV).

The ^{16}N source is the primary calibration source used in the SNO detector. ^{16}N is made outside the detector via the (n:p) reaction on ^{16}O. The activated gas is then transported into cylindrical steel vessel known as the ^{16}N source. The ^{16}N source can be deployed throughout the detector. ^{16}N decays via β-decay into ^{16}O, however only 28% of the decays go directly into the ground state. The γ-rays emitted in the remaining decays are dominated by the a single γ-ray of 6.13MeV (66.2% branch) with a small number of γ-rays at other energies also emitted. The β does not escape the steel, thus the ^{16}N source, provides a sample of γ-rays almost exclusively at 6.13MeV. The β is detected via scintillation light with a PMT inside the source. This means that the ^{16}N source can provide tagged and triggered γ-rays.

The half life of ^{16}N is 7.13s, which combined with the relative ease of ^{16}N production, results in the ability to run the ^{16}N source at a very high rate. The rate at which the source can be run is limited by the ability of the DAQ system to build events at high rate over an extended period of time. For this reason the source can provide a very large number of γ-rays in a large number of positions, making the source a very good source for the study of sacrifice, as it meets most of the required criteria. Questions about the rate effects are checked by running the source at low rate. As the γ-ray travels many centimetres from the source the blocking of backwards light is minimal, which means that geometry based cuts can be tested. Another potential problem is that the ^{16}N source is a γ-ray source not an electron source, this comes into play when the sacrifice of the high level cuts is considered.

The main problem with the ^{16}N source is the range of NHIT values that can be tested successfully. As it is essentially a mono-energetic γ-ray source, the only way that a large number of tubes are hit is by the occurrence of something unusual, such as the coincidence of another event with the γ-ray. However neutrino interactions produce higher NHIT events due to the production of a higher energy electron. This problem means that the

^{16}N source cannot be trusted to measure sacrifice at the extreme high end of the produced NHIT spectrum. Typically the source is used only in the range 30-75 hits. The upper cut off results from the study of cuts that will remove events containing coincidences, such as ITC or ITR. The use of this upper limit results in reliable measurements for sacrifice measured by the ^{16}N source.

The data.

In November and December of 1999 there were a number of special ^{16}N runs taken. These differed from the usual ^{16}N run as the source was not used to trigger the detector, it was instead run in same configuration as for neutrino data. The PMT in the source is still present in the data and this is used to tag the events.

The data taken consists of two vertical scans one in the D_2O and one in the H_2O. This gives source positions covering (most) of the range if interest in detector radius, as well as allowing us to examine the effect of the neck. These runs are known as the data cleaning runs and are used for the majority of the sacrifice analysis. The positions sampled and number of events taken is shown in table 5.1.

The ^{16}N runs were all at high rate. This means that there are more than enough statistics, but that the data does not accurately represent neutrino data, due to the subsequently discovered rate effect in the electronics. To address this fact low rate ^{16}N data taken at a later date is also studied. Also, by a piece of good fortune, at this time there was one card in the electronics that was not working properly. The problem manifests itself as an incorrectly measured charge on 32 channels. The charge was being measured too high, railing the ADC value. By removing this bad board in software before applying the cuts the systematic due to this type of problem can be measured.

In order to monitor the detector for long term changes a series of ^{16}N stability runs were also taken. These are shown in table 5.2. These runs were all taken with the external trigger enabled[2], which is not ideal, but they can still be used to examine sacrifice stability.

5.2.3 The ^{8}Li source

The ^{16}N source does not cover the complete NHIT range of interest for the charged-current spectrum. This means that with just this source the NHIT bias of the sacrifice cannot be properly tested. The ^{8}Li source[74] is used to do this. This source was deployed in December 2000, shortly after the fan upgrade was carried out[3], in four positions along the z-axis, as shown in table 5.3.

[2]The signal from the PMT in the source was used to force a detector trigger.

[3]The fan upgrade was installed to almost eliminate the rate effect seen in the TAC measurement by the electronics, see chapter 2

Run	Media	x/cm	y/cm	z/cm	r/cm	Total Events
10488	D_2O	0	0	-500	500	58823
10486	D_2O	0	0	-400	400	65687
10484	D_2O	0	0	-300	300	62143
10482	D_2O	0	0	-200	200	53741
10480	D_2O	0	0	-100	100	53337
10490	D_2O	0	0	0	0	33665
10353	D_2O	0	0	100	100	68954
10355	D_2O	0	0	200	200	77052
10357	D_2O	0	0	300	300	102727
10359	D_2O	0	0	400	400	147646
10361	D_2O	0	0	500	500	97945
10517	D_2O	0	0	600	600	21442
10573	H_2O	-586	208	-500	797	26879
10571	H_2O	-586	208	-400	739	49896
10568	H_2O	-586	208	-300	690	79921
10566	H_2O	-586	208	-200	653	81928
10564	H_2O	-586	208	-100	630	76993
10575	H_2O	-586	208	0	621	75484
10577	H_2O	-586	208	100	630	75762
10579	H_2O	-586	208	200	653	79526
10581	H_2O	-586	208	300	690	73926
10583	H_2O	-586	208	400	739	68871
10585	H_2O	-586	208	500	797	13090

Table 5.1: The different positions sampled and number of events taken for ^{16}N data used to study sacrifice.

Run	Location	Days since 02/11/99
10278	(0.05,0.17,-0.07)	28
11601	(0.24,0.33,-0.07)	170
12144	(-16.0,23.5,-0.1)	238
12262	(-16.0,23.5,0.03)	269
12498	(-16.0,23.5,-0.15)	291
13381	(-16.0,23.5,-0.28)	330
13467	(-16.0,23.5,-0.1)	350
14273	(-16.0,23.5,-0.3)	405
14501	(0.0,0.0,0.0)	433

Table 5.2: The ^{16}N stability runs.

Run	x/cm	y/cm	z/cm	r/cm	Tagged Events	^{8}Li Events
14329	-16	23.5	0	28	732	592
14331	-16	23.5	0	28	649	525
14333	-16	23.5	0	28	2267	1924
14348	-16	23.5	0	28	26565	23857
14358	-16	23.5	450	450	7266	6473
14360	-16	23.5	450	450	14849	13379
14362	-16	23.5	250	251	1220	1043
14371	-16	23.5	250	251	3965	3514
14373	-16	23.5	550	550	3155	2960

Table 5.3: The ^{8}Li runs taken in December 2000.

^{8}Li decays via β decay to produce an electron, and two α particles. The electron has a similar energy spectrum to that expected from ^{8}B, as both nuclei decay to the same final state. ^{8}Li is produced in a similar fashion to ^{16}N, using the (n:α) reaction on ^{11}B. It is then transported into the detector via a gas system using an aerosol of NaCl in He to carry the ^{8}Li. The source is a thin walled steel ball in which the ^{8}Li decays. The ball is almost spherical, being 4 inches in diameter, with a neck at the top for connection to the manipulator and gas flow systems. The electron can escape the chamber with a typical energy loss of 2MeV. Decays are tagged by the detection of scintillation light from the two α particles using a PMT inside the source.

Due to the much shorter half life (0.838s) and the much smaller yield than ^{16}N, the rate of the source is much less than the ^{16}N source. When the source was deployed the rate of ^{8}Li decays was 0.2Hz. This means that measurements made with ^{8}Li source have much worse statistics than ^{16}N and extensive position scans are not available. The source does however properly cover the NHIT range of interest for study of charge current electrons.

In the production of ^{8}Li a background of ^{16}N events are created. The electron and γ-ray from ^{16}N decay can escape the source, and in some cases give a tag. This occurs when the electron hits the source PMT triggering it. This gives a background of ^{16}N γ-rays to the ^{8}Li signal. In order to provide a pure sample of ^{8}Li events a cut has been developed based upon the charge and time of the source PMT[75].

As well as the problem with the statistics of the source, the ^{8}Li source suffers from one additional problem. As the electrons emit light from the moment they leave the source, a large fraction of the backwards light is blocked. This causes problems particularly with the high level cuts, and for this reason the ^{8}Li is not an ideal electron source.

5.2.4 The laserball

The laserball[76] is the primary source used to measure tube timing[4] and the optical properties of the detector. As the light output is completely adjustable the source has also been used to measure sacrifice.

The laserball is not a Čerenkov source. Instead light from a nitrogen/multi dye laser is channelled into the detector by optical fibres. The fibres feed into an optical diffuser ball to produce light that is close to isotropic. As the source is not Čerenkov light it does not test all the cuts that are in use, such as certain geometry cuts. It can however be used to test the instrumental cuts, both as a constancy check, and for high energy extrapolation.

With these goals in mind a small amount of laser data was taken in December 1999, shortly after the ^{16}N data cleaning runs were taken. Three runs were taken in the centre with the light output adjusted to scan the range of interest. A further two runs were taken at z=+550cm, again with the interesting range of NHIT covered.

5.2.5 The pT source

The pT source [77] is the source used in SNO to bracket the energy range of solar neutrinos. The source creates a γ-ray of 19.8MeV using the $^3\mathrm{H}(\mathrm{p},\gamma)^4\mathrm{He}$ reaction. This is achieved by the acceleration of protons up to 30keV, inside the source, to collide with a target of high purity scandium tritide. As a by-product a large number of neutrons are produced, with approximately 50000 neutrons per γ-ray .These provide a background to the γ-ray signal particularly by multiple neutrons that can cause a problem for sacrifice measurements. However the principle problem when using the pT source to measure sacrifice is that the γ-rays are not tagged. The use of the pT source will be discussed further when the sacrifice at high energy is considered.

5.3 Measurement of the sacrifice

In order to measure the sacrifice every cut that is applied to the data must be studied. The method used to study each cut can differ, depending upon the properties of the cut, or suite of cuts under study. By splitting the cuts into smaller sets, there is the risk of double counting sacrifice. Therefore the correlation of cut sacrifice is also studied by examination of the correlation matrices. Uncorrelated suites of cuts can be found and then the sacrifice examined in more detail.

[4] The PCA constants.

	No cuts	Instrumental Cuts	Fitter	FoM Cuts	θ_{ij}	ITR
No Cuts	619362	0	0	0	0	0
Instrumental Cuts	0	2657	1	13	34	16
Fitter	0	1	258	258	258	258
FoM Cuts	0	13	258	1805	538	271
θ_{ij}	0	34	258	538	9436	464
ITR	0	16	258	271	464	1277

Table 5.4: ^{16}N correlations between different levels of cuts. Only ^{16}N runs within the D_2O are used and only data inside the NHIT window 40-75 are considered.

	No cuts	Instrumental Cuts	Fitter	FoM Cuts	θ_{ij}	ITR
No Cuts	25348	0	0	0	0	0
Instrumental Cuts	0	82	0	2	0	0
Fitter	0	0	1	1	1	1
FoM Cuts	0	2	1	155	5	1
θ_{ij}	0	0	1	5	50	7
ITR	0	0	1	1	7	59

Table 5.5: ^{8}Li correlations between different levels of cuts. An NHIT cut is applied at 40 hits.

5.3.1 Sacrifice correlations

The cuts naturally split into two separate sets, the instrumental cuts and the high level cuts including the fitter and figure of merit cuts. It is informative to examine the correlation matrix of each high level cut compared to the instrumental suite. Examination of the correlation matrices for ^{16}N and ^{8}Li data taken inside the D_2O above a reasonable analysis threshold tells us if the two sets of cuts are sufficiently uncorrelated for us to proceed with an independent approach for each suite of cuts. The Correlation matrices are shown in tables 5.4 and 5.5. These tables show the number of events failing different combinations of cuts for the two source data sets.

These tables show that the correlation between the instrumental cuts and the other cuts is either very small or that they are uncorrelated in the signal region. This allows us to treat the two sets of cuts independently. The fitter is totally correlated with both the high level cuts and the figure of merit cuts, but uncorrelated with the instrumental cuts. This correlation is as expected as cuts requiring a fit vertex fail by default if one cannot be found. By requiring a fit for the measurement of the sacrifice of these cuts, the fitter can be treated as independent and measured separately. The figure of merit cuts and high level cuts are however correlated, and so must be analysed together when calculating the

	AMB	QvT	QCluster	FTS	OWL	Neck	FGC
AMB	930	13	24	0	0	0	11
QvT	13	1020	24	0	0	0	85
QCluster	24	21	225	0	0	0	28
FTS	0	0	0	10	0	0	0
OWL	0	0	0	0	234	0	0
Neck	0	0	0	0	0	29	0
FGC	11	85	28	0	0	0	361

Table 5.6: The correlated sacrifice for instrumental cuts using D_2O ^{16}N above 40 hits. There are 633472 events in the data sample.

total sacrifice.

The study of cut correlations shows us that we have three independent cut sacrifices to measure. We now have to measure the sacrifice from the data cleaning cuts, fitter and the combination of the high level and figure of merit cuts.

As we are attempting to understand the sacrifice of the entire suite of cuts it is also informative to examine the correlation between cuts in the suite of instrumental cuts. ^{16}N data taken inside the D_2O is used, with an NHIT threshold at 40 hits. The correlation matrix for cuts with a non-zero sacrifice is shown in table 5.6. Each of the cuts based upon the charge in some way has correlated sacrifice with each of the others. As each cut is protected against a bad charge measurements in some way, this suggests that in some cases genuine high charges can occur in events and cause them to fail the cuts.

There is also a correlation between FGC and both QCluster and QvT. The correlation between FGC and QCluster is due to to occurrence of a sharkfin in the ^{16}N event. As sharkfins are QClusters, they will fail this cut, and depending upon the tube in which they occur, they can cause FGC to fail if the tube is on the opposite side of the detector to the direction of the event. The overlap between QvT and FGC is due to early clusters with high charges. Many of these events do not fail QCluster due to the smaller charge threshold of the QvT cut.

5.3.2 Sacrifice of the instrumental cuts

The event based instrumental cuts were described in chapter 3 and chapter 4. However until this the point, the sacrifice of the cuts, which is one of the principle cut properties, has not been discussed in detail. Every cut has a sacrifice and this needs to be measured and understood. In some cases the sacrifice of a cut was such that the cut required loosening, or even a change in the algorithm, conversely it was found that some cuts could be tightened due to extremely low sacrifice measurements.

The problem with attempting to measure the sacrifice of the instrumental cuts is that in most cases the cut is based upon a detector property that is not well modelled in the Monte Carlo[5], such as the details of the charge spectrum, crosstalk, the ESUM signal, or veto tubes. For this reason sacrifice of the instrumental cuts is a property that is measured using calibration data.

For the most part the ^{16}N source is used. The data is available at multiple positions, and the presence of the source has minimal impact on the events. For measurements at higher NHIT values that cannot be reliably probed using the ^{16}N source the ^{8}Li source is used. As has been discussed it is desirable to know the sacrifice as a function of energy position and direction. The variables of choice are NHIT, $(\frac{r}{600cm})^3$ and **u.r**. NHIT is chosen because sacrifice can depend upon factors like tube counting statistics, much more than upon the energy of the Čerenkov event. The choice of $(\frac{r}{600cm})^3$ allows things to be volume weighted in a simple way, and normalises r^3 to the detector. In a spherical geometry **u.r** is the only measurement of direction of interest. It is the normalised dot product of the position and direction vectors, thus **u.r** =+1 means the event is radially outwards and **u.r** =-1 the event is radially inwards.

Using the calibration data the sacrifice can be measured inside the analysis region. The choice of fiducial volume and energy threshold means that sacrifice is measured inside r=550cm[6] and above 40 hits[7], however it is of interest to examine the sacrifice outside this region to search for trends. In the following analyses unless otherwise stated a fiducial cut has been applied at 550cm and an NHIT cut at ≥ 40.

The use of ^{16}N data

The main sacrifice measurement is carried out using ^{16}N data. This source was also used to study individual cuts, as well as during their design. The sacrifice and the breakdown into the various cuts is shown in figures 5.1, 5.2 and 5.3. These plots show the combined sacrifice of the instrumental cuts and of the cuts that provide the biggest contribution: the AMB, QvT, QCluster and FGC cuts. The FTS cut is also shown as this cut has the biggest sacrifice bias of any of the cuts. The remaining contribution to the sacrifice is dominated by the OWL cut. Previously the sacrifice of this cut has been shown to be due to coincident noise in the OWL tubes and is therefore unbiased for any of the plotted variables. This cut and the other cuts whose contribution is much smaller have been excluded from the plots to improve their clarity. The total sacrifice shown does however

[5] In the case of OWL or neck tubes for example, the detector property is not modelled at all

[6] This choice of fiducial volume corresponds to $(\frac{r}{600cm})^3 = 0.77$

[7] This choice of NHIT threshold means that all events above the chosen kinetic energy threshold of 6.75MeV are considered, as well as the vast majority of the events above an analysis threshold for neutral current of perhaps 5.5MeV.

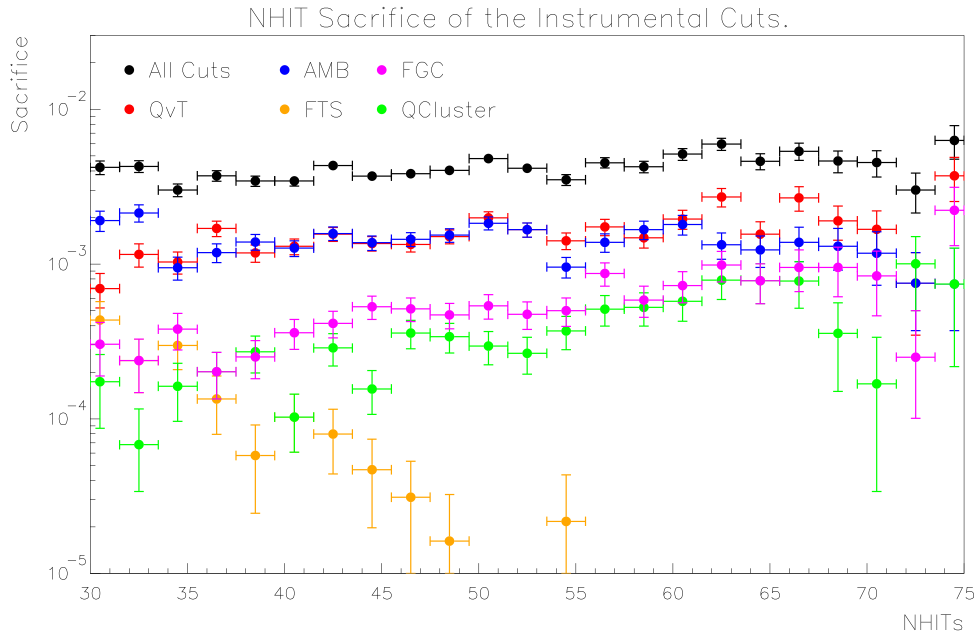

Figure 5.1: The NHIT dependence of the instrumental cuts sacrifice using ^{16}N data, for events within a fiducial volume of 550cm.

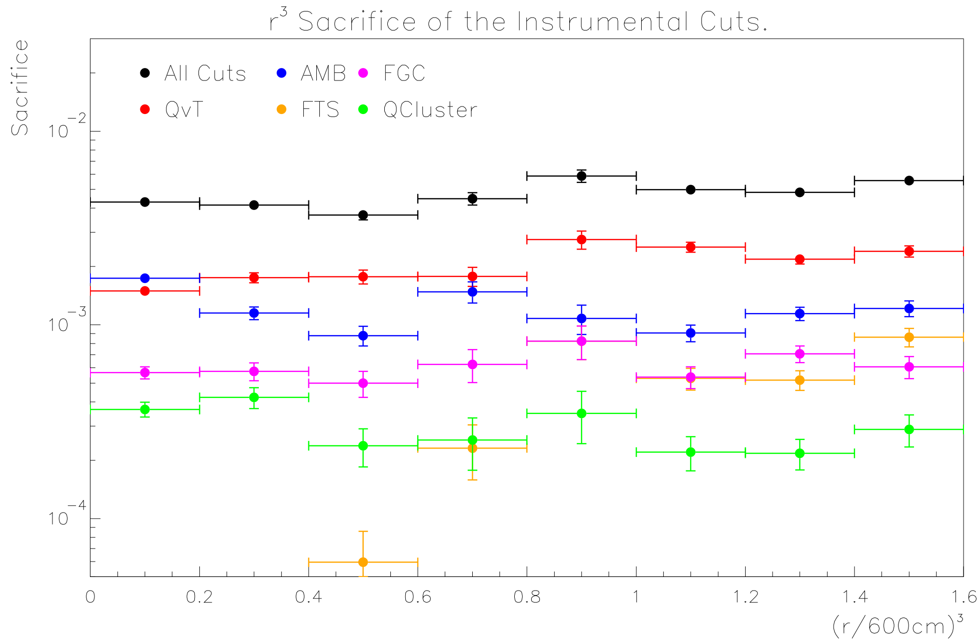

Figure 5.2: The $(\frac{r}{600cm})^3$ dependence of the instrumental cuts sacrifice using ^{16}N data, for events above 40 hits.

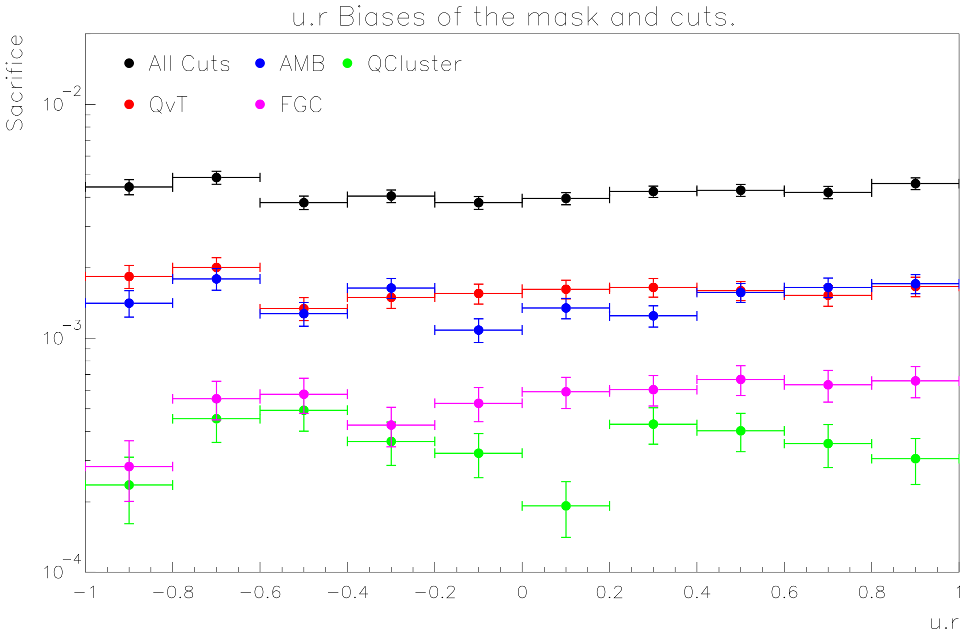

Figure 5.3: The **u.r** dependence of the instrumental cuts sacrifice using ^{16}N data, for events within a fiducial volume of 550cm and above 40 hits.

include these cuts.

The first thing to note about the sacrifice in the signal region is that it is not a strong function of NHITs, position, or direction. This property of the sacrifice is very useful when combining the sacrifice shape with the shape of the neutrino signal. This comes about as none of the four dominant cuts are strong functions of these parameters, at least in the NHIT range covered by the ^{16}N source. Outside of the signal region, beyond the scope of the shown plots, this is no longer the case. At lower NHIT the sacrifice degrades. This is not to surprising as the cuts were designed to work at higher NHIT. Also the statistics in the number of hits causes a number of distributions to broaden causing more sacrifice. The most important effect however is the increase in the sacrifice as the events approach the PSUP. The sacrifice reaches values of several percent by the time the PSUP is reached. This is caused by a number of effects that make the Čerenkov events look more like background. The effects include increased charge deposits on PMTs, and the creation of separated clusters due to the effect of the concentrator cutoff. However as shown in figure 5.2 the sacrifice does meet the target of less than 1% inside 7m.

A feature somewhat hidden in these plots is the shape of the sacrifice of the FTS cut. This cut has an extreme shape to its sacrifice. Events at low NHIT heading inwards at high radius have a much higher sacrifice than events in the centre. This feature lead to the redesign of the FTS cut, to require the maximum time between hits to form a valid pair and to a re-evaluation of the placement of the cut boundary. These changes reduced the problem by reducing the overall sacrifice of the cut. In the signal region the sacrifice of the FTS cut is minimal, rendering the shape of the sacrifice of this cut unimportant. It is worth noting however that increased sacrifice at high radius and low NHIT is still present, but this is hidden by the nature of the cuts applied in the creation of figures 5.1, 5.2, and 5.3.

Of the other cuts not shown in the plots, only the crate isotropy cut shows any deviation from a very low level sacrifice. This cut does have a radial bias, but this only becomes apparent at higher radii than 7m. The shape of the sacrifice is very dependent on the exact location and direction of the event, as it depends strongly upon which channels are hit. Despite this, the sacrifice of the cut inside 7m is very small, even at much lower NHIT values.

Systematic checks.

The first systematic check that needs to be applied is a check on the use of the fitter. In determination of the radial and directional bias reconstruction is essential. It is therefore important to demonstrate that the choice of fitter and reconstruction tails are unimportant.

The best way to do this is to remove badly fit events from the sacrifice analysis by the use of cuts around the source position. The result of this analysis is shown in figure 5.4. This figure clearly shows that cuts around the source have a minimal impact on the sacrifice measurement. The only region where cut becomes apparent near the AV where source locations are sparce, and thus statistics limited. A similar analysis carried out using the various available fitters shows the same result.

The systematic check that has the largest effect using the ^{16}N source is the test of the robustness of the the cuts to bad channels and bad calibrations. As previously discussed there was one bad card present during the ^{16}N data taking. This caused 32 channels to report the charge in every hit as railed. This is not a good situation for data cleaning, as this measurement is what many of the cuts look for. The QvT cut in particular can be badly effected. By removing hits on the bad card in software, before the cuts are applied to the data, the systematic effect on the sacrifice can be measured.

The stability of the sacrifice to bad channels in general can also be tested in this way. As with all runs the ^{16}N data contains badly calibrated channels that are flagged by ECA calibration. The data cleaning process does not use this information as disabling a channel in software is a very good way to allow events to leak through the cuts. However using ECA validation the effect of this choice on the sacrifice measurement can be tested.

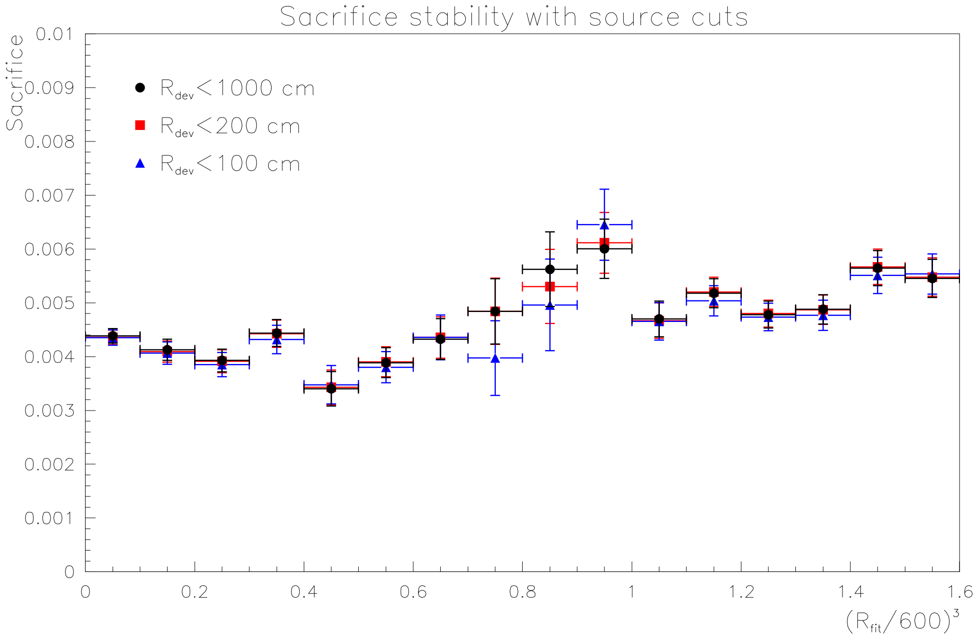

Figure 5.4: The $(\frac{r}{600cm})^3$ dependence of the instrumental cuts sacrifice using ^{16}N data, for events above 40 hits, with various cuts on r_{dev}, the distance from the reconstructed vertex to the source.

This does not change the measurement from the bad card as ECA validation does not flag these channels reliably.

Once the sacrifice has been calculated without the bad card two cuts have their sacrifice reduced. These are QvT and QCluster. The effect on QvT is more dramatic with a sacrifice reduction of $(0.088 \pm 0.007)\%$. The effect on QCluster is smaller with a sacrifice reduction of $(0.030 \pm 0.004)\%$ observed. From this point the sacrifice with the bad card removed will be used as the baseline sacrifice measurement. The effect of the bad card will be returned to when the stability of the sacrifice is considered.

The effect of ECA validation is more evenly distributed across the cuts and the reduction in the sacrifice is much smaller. One interesting feature however is that although the sacrifice in the D_2O is reduced, in the H_2O region the sacrifice is actually increased by the application of ECA validation. The overall change in sacrifice for the three conditions is shown in figure 5.5. The result of the sacrifice measurement with the application of ECA validation measures the systematic error on the baseline level of sacrifice due to extra bad channels.

A third systematic check that has been carried out is to use the low rate ^{16}N source to check for systematics of this type. This study shows that, with the available data, that

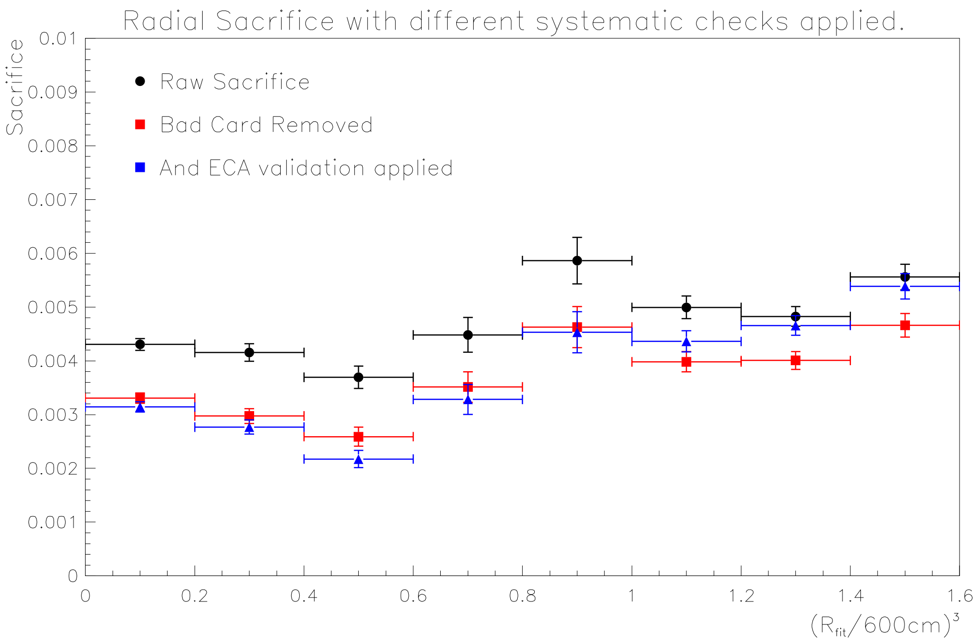

Figure 5.5: The $(\frac{r}{600cm})^3$ dependence of the instrumental cuts sacrifice using ^{16}N data, for events above 40 hits, with various channel property cuts applied before cut application.

is limited both statistically and in position, there is no rate effect observed in the sacrifice measurement[78].

The use of ^{8}Li data

The sacrifice measurement, made using ^{16}N data, is only valid below 75 hits. Thus in order to measure the sacrifice up to the end point of the ^{8}B spectrum, a further source is needed: the ^{8}Li source. When the sacrifice analysis was carried out a distinct bias was seen with NHIT. Further analysis showed that there was a single channel present in the data with railed charges and TAC measurements. Consequently events with hits on this channel always failed the QvT cut, introducing the observed NHIT bias. This was eliminated by the alteration of the QvT cut to avoid the use of channels with invalid times.

The sacrifice from ^{8}Li after this modification is shown in figure 5.6. This shows that once the QvT fix is applied there is no apparent NHIT bias in the ^{8}Li data. This statement is however statistically limited in that the statistics at high NHIT with the ^{8}Li source are not enough to make an accurate measurement of sacrifice. Another key feature of the sacrifice measured by the ^{8}Li source is that in the overlap region there is agreement with the measurement from the ^{16}N source. These observations give us confidence that

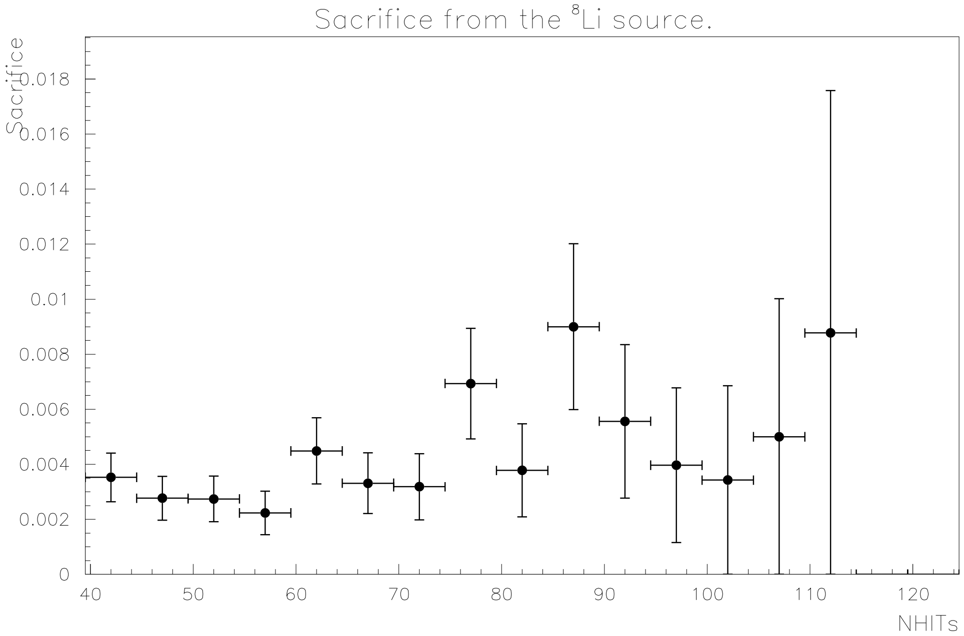

Figure 5.6: The NHIT dependence of the instrumental cuts sacrifice using ^{8}Li data.

measurements made with the ^{16}N source at lower NHIT can be extended to cover the ^8B spectrum.

The laserball

To provide further confidence with the measurements made at high NHIT the laserball is also used to measure sacrifice. The results of this study are shown in figure 5.7. The laser confirms the ^{8}Li measurement at high NHIT values and suggests that the sacrifice is stable to even higher NHIT values. One interesting feature of the laser data is the low NHIT disagreement at high radius. This shows why a Čerenkov source is required to measure the sacrifice, as this feature is due to the laserball. It does however suggest that the response of the cuts to neutrons in salt, where neutron events are much more isotropic, will require further study.

The sacrifice measurement

So far only sacrifice distributions have been shown. To obtain the total sacrifice required for a flux measurement the sacrifice distributions obtained must be integrated over the expectation for the various signal classes. This being the case we are fortunate to discover that the sacrifice is fit well be a flat distribution in NHITs, $(\frac{r}{600cm})^3$, and $\mathbf{u.r}$. This makes

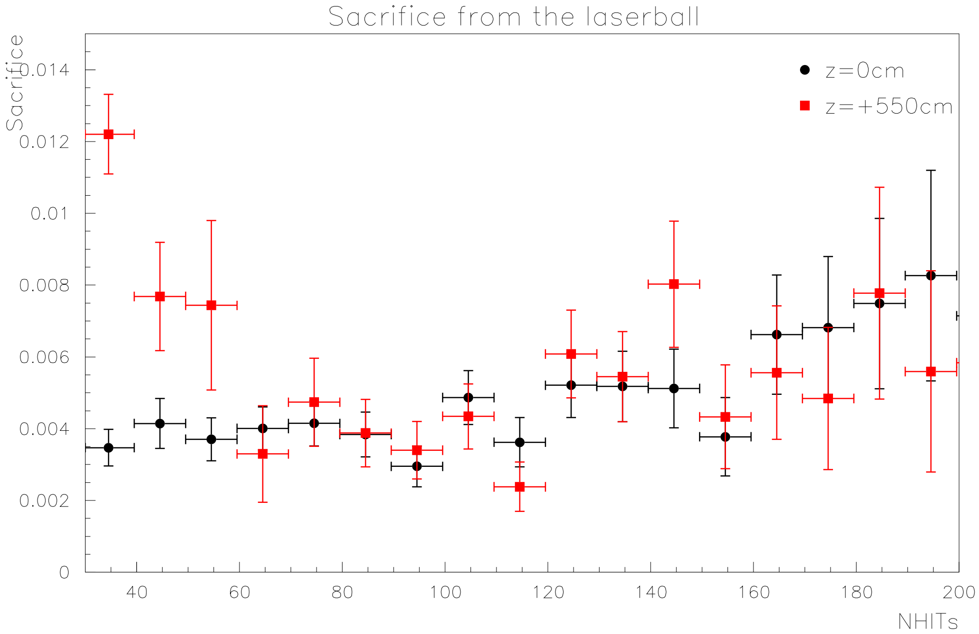

Figure 5.7: The NHIT dependence of the instrumental cuts sacrifice using laserball data.

it easy to integrate the sacrifice with the expected signals to obtain the expected sacrifice on the flux measurement.

The primary measurement used is that from the ^{16}N, with the ^{8}Li events used to verify the continuation of the sacrifice distribution up to the end of the expected energy spectrum. Laser data is not used, but does give us confidence that a gross mistake is not being made at higher NHIT values. At the very end of the spectrum where hep events dominate the sacrifice is measured separately later in this chapter.

The fits used in the sacrifice determination are shown in figure 5.8. These show that in general a flat distribution is a reasonably good fit to the data, although in some cases the deviations from flat are quite large. This must be taken care of in the systematic error. Firstly however the results of the fits should be discussed. These are shown in table 5.7. The results from the the three ^{16}N distributions show very good agreement with each other. This is not surprising as the same data has been used in each case, but this does show that a major mistake has not been made fitting the data to flat distributions. The measurement from ^{8}Li is in good agreement with result from ^{16}N. This is an important observation, as the measurement has been made with a different source, over a year later in a rather different detector. This also justifies the use of the result with the bad card removed for the baseline sacrifice measurement.

The systematic error from the use of ECA validation is obtained via the same

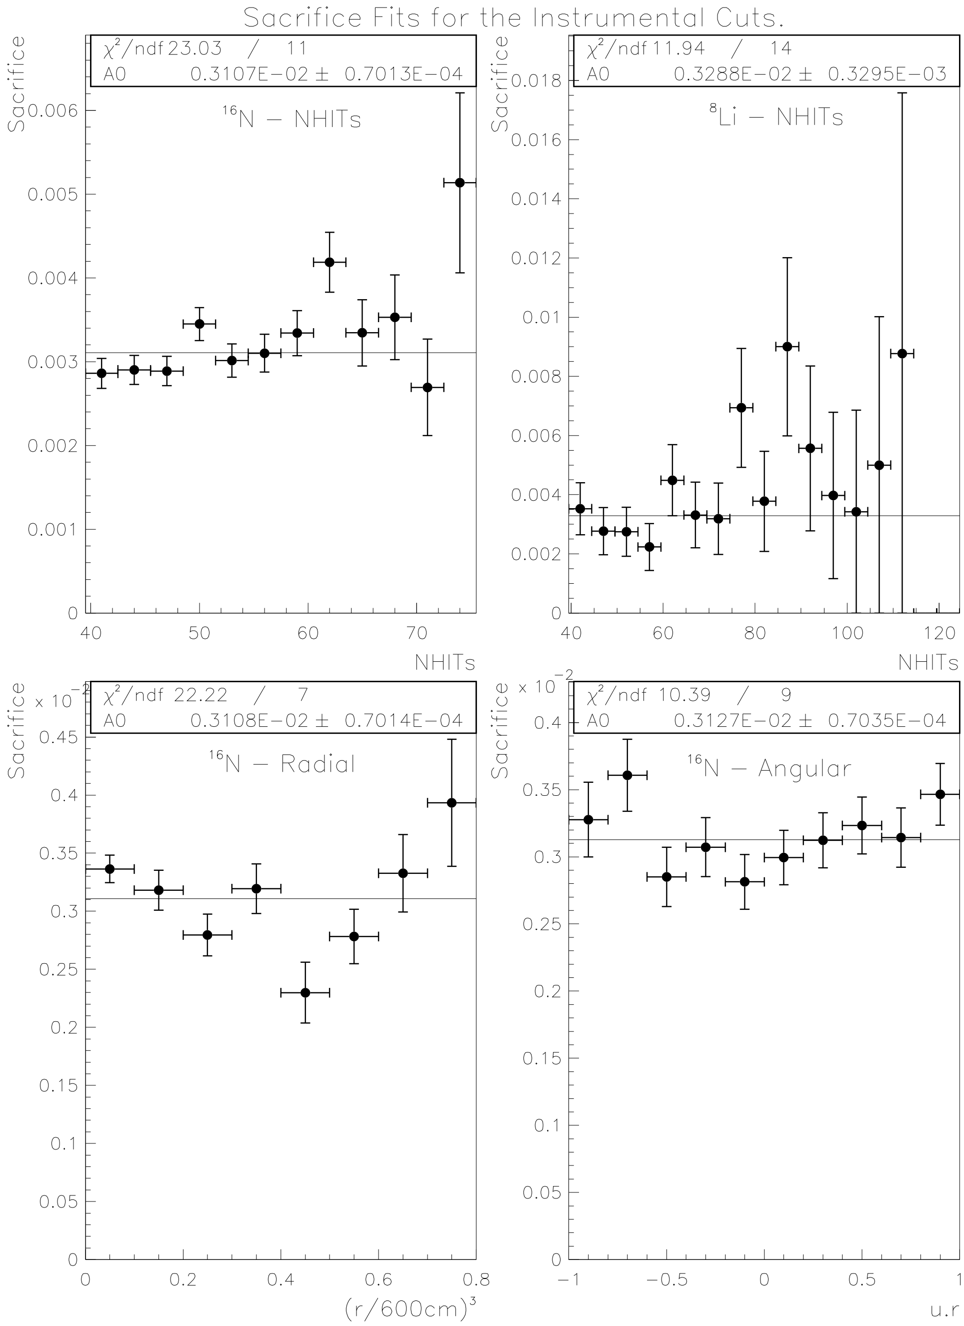

Figure 5.8: The fit of the sacrifice distributions to a flat distribution. The upper plots show NHITs using both ^{16}N (left) and ^{8}Li (right). The lower plots are both based upon ^{16}N with $(\frac{r}{600cm})^3$ on the left and **u.r** on the right.

Distribution	Sacrifice Result /%
NHITs ^{16}N	0.311 ± 0.007
$(\frac{r}{600cm})^3$ ^{16}N	0.311 ± 0.007
$\mathbf{u.r}$ ^{16}N	0.313 ± 0.007
NHITs ^{8}Li	0.329 ± 0.033

Table 5.7: The results of the sacrifice fits for the distributions shown in figure 5.8 including the statistical error from the fit.

technique. The sacrifice from ^{16}N is fit to a flat distribution and the result $(0.290 \pm 0.007)\%$ is obtained. This suggests a systematic error of -0.021% due to ECA validation, for data within the D_2O.

As already stated the deviations from a flat distribution must also be added to the systematic error. The radial distribution shows the biggest deviations, so this distribution is used to determine the systematic error. The technique used to estimate this error is to use the median absolute deviation of the data from the fit value. When this is carried out the systematic error due to deviations from flat is found to be is 0.028%. The remaining systematics in the measurement of the sacrifice of the instrumental cuts are all due to sacrifice stability. These systematics are dealt with in section 5.4.

5.3.3 Fitter sacrifice

After the application of the instrumental cuts the next stage of the analysis is to reconstruct the event. This procedure does have the potential to fail events and therefore introduce sacrifice. This sacrifice needs to be measured. Fitters are briefly discussed in chapter 3.

As discussed the fitter of choice in SNO is known as the path fitter. The path fitter primarily fails events due to problems with the seeding procedure. The path fitter seed comes from either the grid or the time fitter[41]. If a seed cannot be found inside 830cm (the PSUP radius) then the path fitter fails the event. On much rarer occasions the path fitter will itself fail events by a failure in the fitting procedure. Events may also fail one of the figure of merit cuts. The sacrifice of these cuts is correlated with the sacrifice of the high level cuts, and is therefore considered in combination with them. As events that fail these cuts still return a fit vertex they can be ignored while considering the fitter sacrifice.

As with all sacrifice measurements, the NHIT, position and directional dependence of the sacrifice is desirable. The NHIT dependence can be obtained from ^{8}Li across the entire range of interest in NHIT, although with limited statistics and positions as usual. The NHIT dependence of fitter sacrifice from ^{8}Li is shown in figure 5.9. The measurement is statistically limited, but seems to suggest that the fitter sacrifice decreases as more tubes

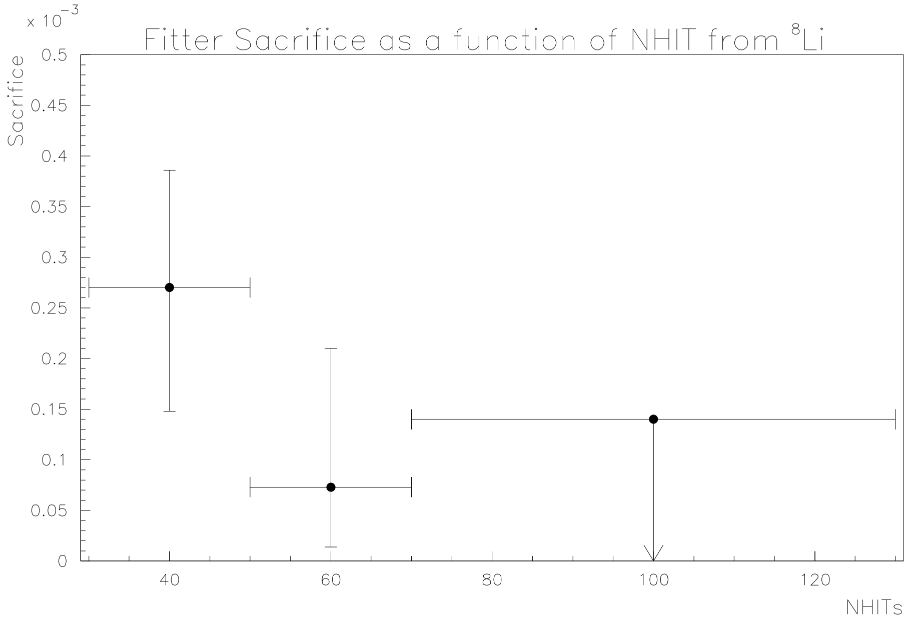

Figure 5.9: The NHIT dependence of fitter sacrifice from ^{8}Li data. For the three highest energy bins there are no events sacrificed, so a combined upper limit is shown. This is a 68%(1σ) upper limit to maintain compatibility with the error bars on the two measurements.

are available for the fit. This is as expected, as with more information it is easier to find a good vertex.

Using this information a measurement at the ^{16}N energy would be an upper limit to a measurement of the fitter sacrifice of the neutrino signal. However the ^{16}N source is the only source that can be used to measure the radial dependence of fitter sacrifice. The fitter sacrifice for every run in the ^{16}N scan taken in November and December 1999 is shown in figure 5.10. This figure can be used to measure the fitter sacrifice inside the fiducial volume at the ^{16}N energy. Conservatively the fitter sacrifice inside the fiducial volume is less than 0.3%. This is very much an upper limit using the NHIT information obtained from ^{8}Li . The sacrifice shows a large radial bias. This is not to surprising as the main cause of sacrifice is the failure of the fitters providing the seed to reconstruct the event inside the 830cm cut off imposed by the path fitter. This occurs due to events on the tail of the resolution of those fitters. This also explains the very high sacrifice in the H_2O region[8].

[8]The extreme size of the sacrifice (approximately 10%) and the relatively unknown behaviour in NHITs suggests that the seeding approach used by the path fitter would need to be revised before an analysis including H_2O data could proceed.

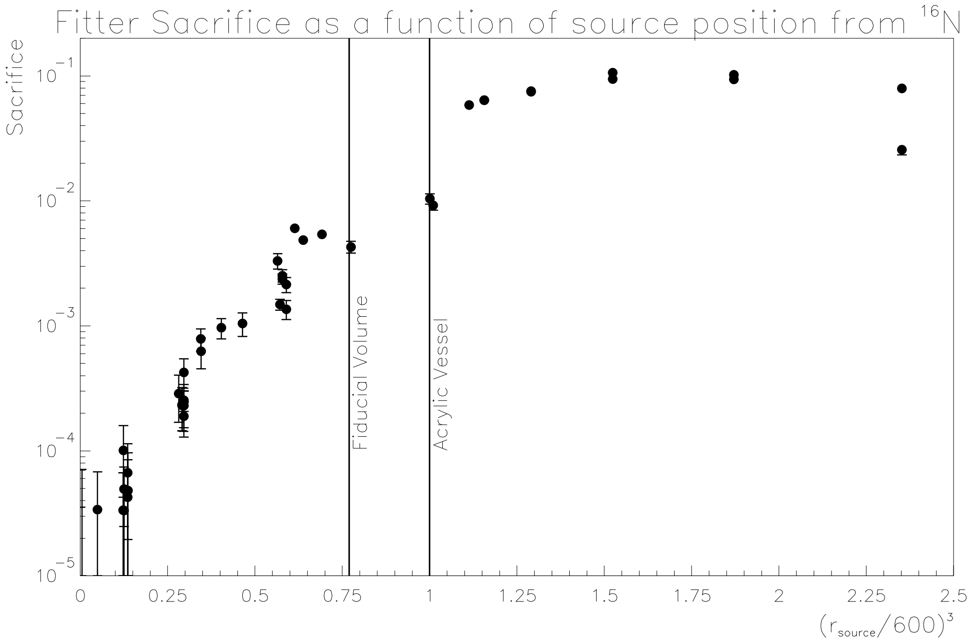

Figure 5.10: The fitter sacrifice as a function of source radius for each run in the November/December 1999 scan. Only events with NHITs$\geq$ 45 are used.

5.3.4 Sacrifice of the high level cuts

The final sacrifice measurement that needs to be made is the sacrifice of the high level cuts. These are the ITR and θ_{ij} cuts, alongside the figure of merit (FoM) cuts from reconstruction. The sacrifice of these cuts is correlated, so the sacrifice of the combination must be measured.

Unlike the case of the instrumental cuts, the high level cuts use only time and geometry, two distributions that are well modelled by the Monte Carlo. The combination of this with the more extreme shape of the sacrifice compared to the instrumental cuts and the difference between γ-rays and electrons for θ_{ij}, means that the Monte Carlo is used to measure the sacrifice of these cuts. This is however not straightforward, as the sacrifice is a measurement of the tail of a distribution. The accuracy to which the Monte Carlo reproduces the main body of a distribution is not guaranteed to be the same in the tail. For this reason it is important to calibrate to Monte Carlo measurement using calibration data.

This however presents a problem. The geometry cuts used are very sensitive to every hit produced in an event, in particular backward light. This means that if the source used to calibrate the Monte Carlo sacrifice has an impact on the event geometry, then the sacrifice measured from the source is not representative of solar neutrino events inside the detector. The ^{8}Li source is the worst affected by this problem. As electrons emerge from the steel they produce all their Čerenkov light within a few centimetres of the source. This results in all the light produced in 2π steradians backwards from the event being blocked. This has a particular impact on the θ_{ij} cut as the backwards light is responsible for the majority of the sacrifice imposed by the high end cut. For this reason the ^{8}Li source produces a serious underestimate of the sacrifice of the θ_{ij} cut. A similar problem is also seen with the FoM cuts. This problem prevents the ^{8}Li source from being used to measure the sacrifice of the high level cuts, leaving just the ^{16}N source. This source does not suffer from the same problem as ^{8}Li as the γ-ray travels some distance from the source before any Čerenkov light is produced, vastly reducing the solid angle for blocked light.

For the high level cuts the triggering of the detector is not as important, so the entire energy calibration scan from December 1999 is used[57]. This provides increased statistics, particularly at high radius, and better sampling around the detector volume than the data cleaning runs used previously. To calibrate the Monte Carlo equal numbers of events from each source position are compared to a data set of Monte Carlo events from the same source positions.

The figure of merit cuts

Although the sacrifice of the combination of the high level cuts must be measured, to gain understanding of the total sacrifice, each cut should be examined individually. The figure of merit cuts are well reproduced by the Monte Carlo as shown in figure 5.11 The figure of merit cuts show a large sacrifice bias at low NHIT. This is a feature that is seen in all the high level cuts, and this would be problematic for the approach used to measure the sacrifice of the instrumental cuts.

The ITR cut

The sacrifice for the ITR cut is shown in figure 5.12. The agreement between data and Monte Carlo is not nearly as good as seen with the figure of merit cuts. This is almost certainly due to the slight discrepancy seen in the late light distribution between data and Monte Carlo, combined with the imperfect time calibrations present only in the data. These effects combine to give the increased sacrifice seen in the data compared to the Monte Carlo. The sacrifice of the ITR cut in both both the data and the Monte Carlo are however small inside the analysis region.

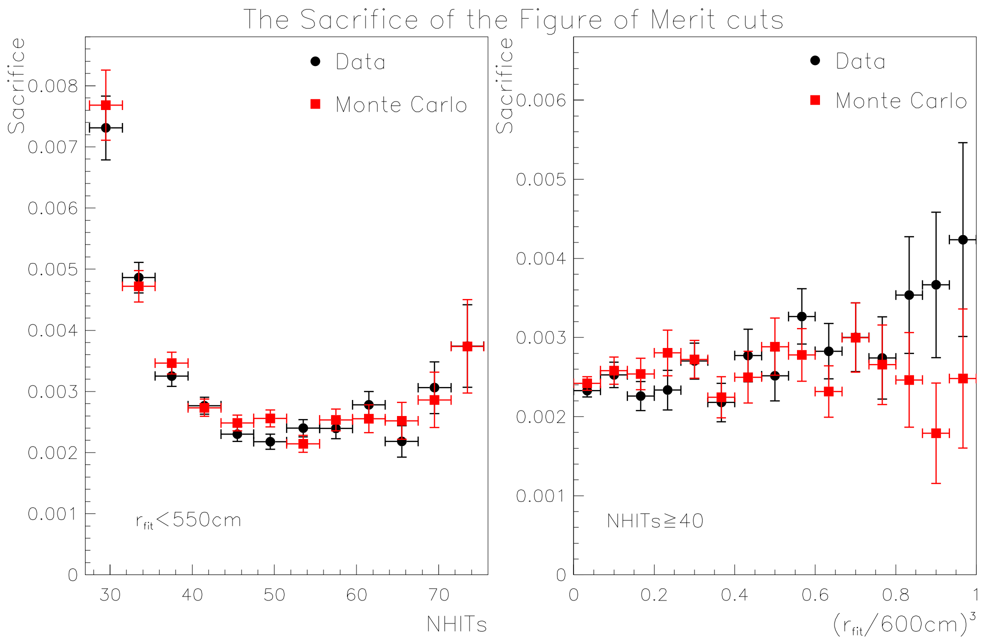

Figure 5.11: The sacrifice of the figure of merit cuts from ^{16}N data and Monte Carlo.

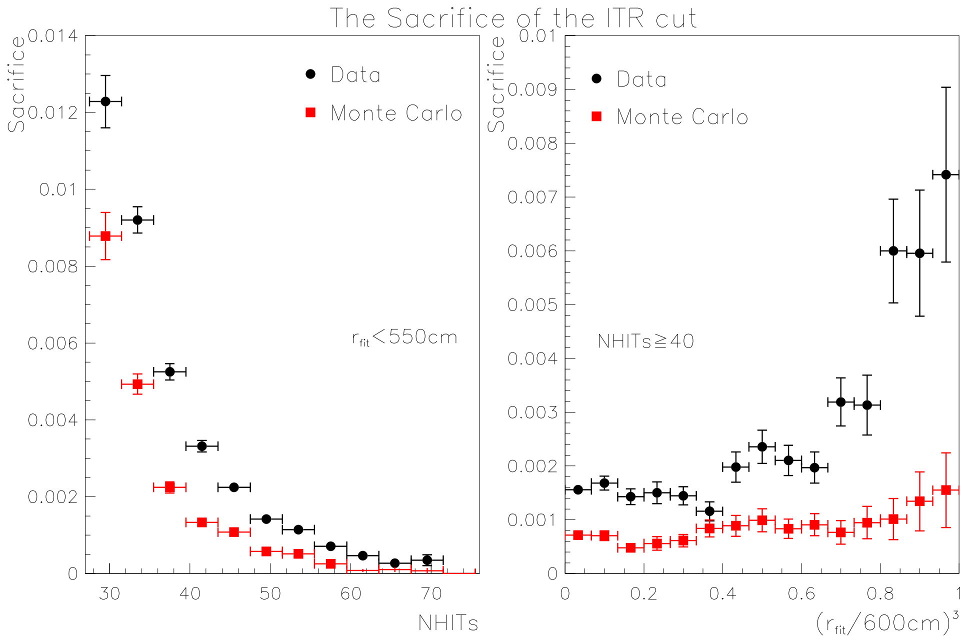

Figure 5.12: The sacrifice of the ITR cut from ^{16}N data and Monte Carlo.

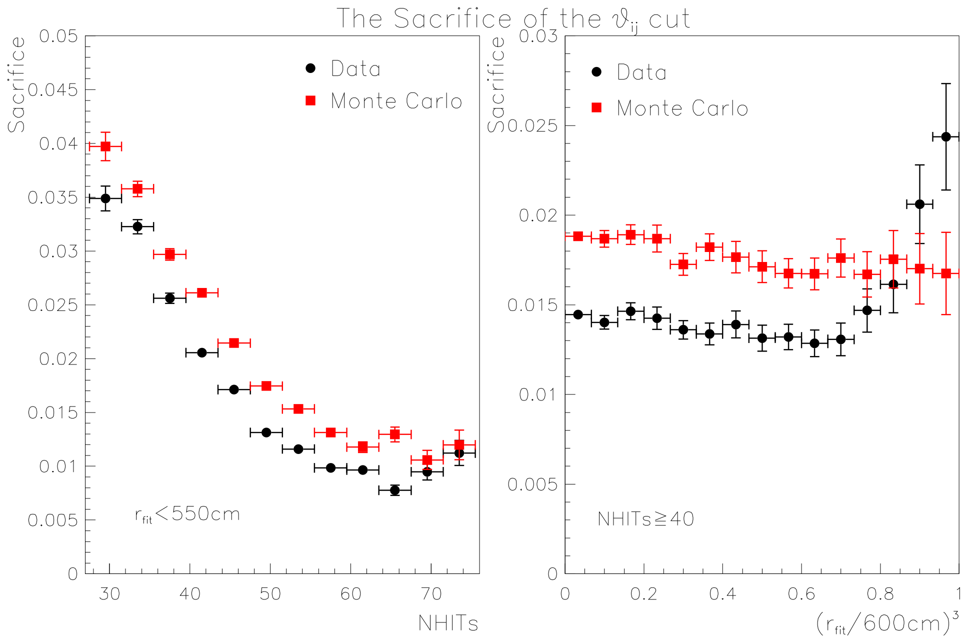

Figure 5.13: The sacrifice of the θ_{ij} cut from ^{16}N data and Monte Carlo.

The θ_{ij} cut

Figure 5.13 shows the sacrifice comparison for the θ_{ij} cut. As was the case for the ITR distribution the Monte Carlo does not reproduce the correct sacrifice. In this case however the sacrifice is overestimated by the Monte Carlo. Also the shape of the radial sacrifice is not well reproduced, the sharp rise as the events move closer to the acrylic vessel is not predicted, but fortunately this rise lies almost completely outside the chosen fiducial volume. The choice of fiducial volume was not made for this reason, but due the lack of understanding of the optical response of the detector and the general lack of calibration source locations outside this volume. These reasons are almost certainly responsible for the failure of the Monte Carlo to correctly predict the θ_{ij} sacrifice outside this region. This failure also justifies the choice of fiducial volume at r=550cm.

The sacrifice ratio

Within the fiducial volume the shape of the sacrifice for the cuts is well predicted by the Monte Carlo, but the value is not. Therefore in order to use the Monte Carlo to measure the sacrifice a correction must be applied. This takes the form of the a scale factor derived from the ^{16}N data already discussed. In order to measure the scale factor and verify that the Monte Carlo does produce the correct shape the sacrifice ratio between data and Monte

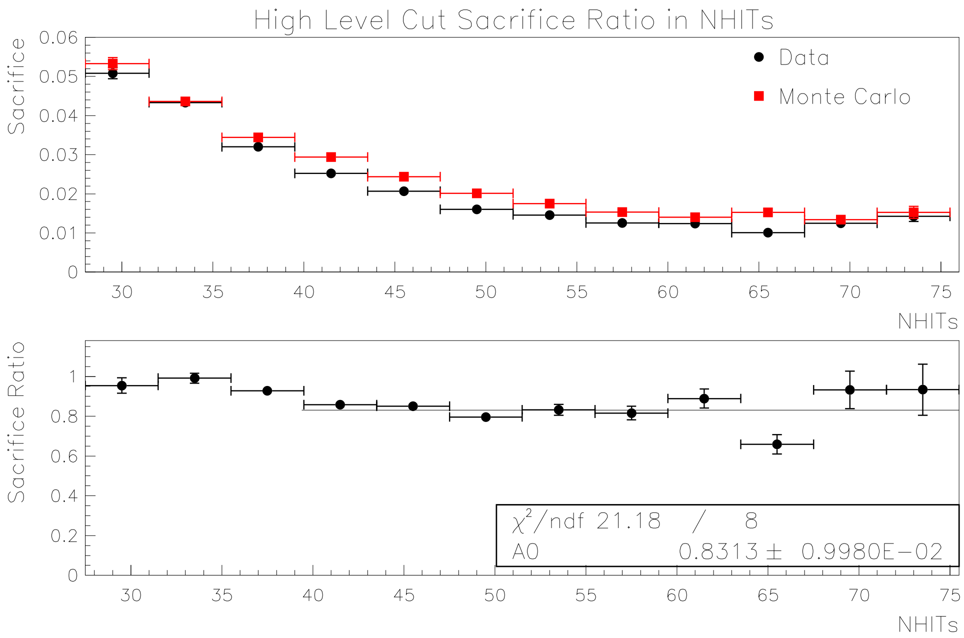

Figure 5.14: The sacrifice distributions and ratio in NHITs for the high level cuts applied to ^{16}N data and Monte Carlo. A fiducial volume cut at 550cm is applied to the data.

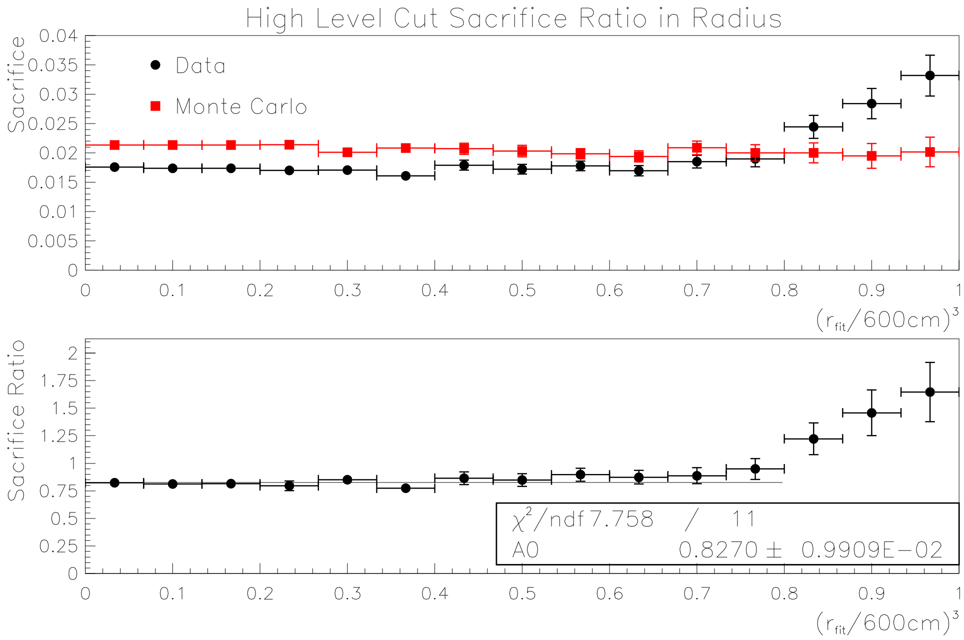

Figure 5.15: The sacrifice distributions and ratio in $(\frac{r}{600cm})^3$ for the high level cuts applied to ^{16}N data and Monte Carlo. A cut at 40 hits is applied to the data.

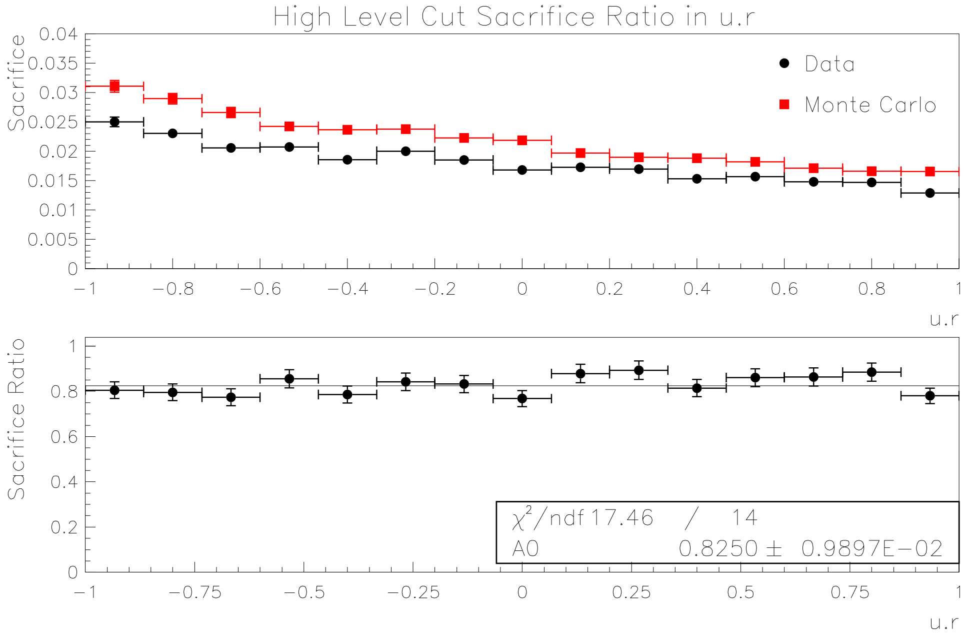

Figure 5.16: The sacrifice distributions and ratio in **u.r** for the high level cuts applied to ^{16}N data and Monte Carlo. A fiducial volume cut at 550cm and an NHIT cut at 40 are applied to the data.

Carlo is studied for the combined set of cuts. The sacrifice ratio distributions in NHITs, $(\frac{r}{600cm})^3$, and **u.r** are shown in figures 5.14, 5.15, and 5.16. Also shown in these plots are the sacrifice distributions for the combined high level cuts. Inside the analysis region[9] the sacrifice ratio fits well to a flat distribution, in each variable of interest. However outside this region the fit breaks down. At low NHIT the difference between data and Monte Carlo seems to disappear, due to a cancellation in the differences from different cuts. At high radius the flat distribution breaks down as the θ_{ij} cut shape change is not predicted by the Monte Carlo.

Despite this the method works inside the analysis region and the scale factor is measured to be 0.83 ± 0.01 where the error is the statistical error from the fit. As was the case for instrumental cuts, the deviations from the Monte Carlo shape are measured by examination of the median absolute deviation from the fit. When measured this turns out to be 0.035 from each distribution. This is added in quadrature with the statistical error to give a final measurement of 0.83 ± 0.04 for the scale factor to be applied to the measured sacrifice from the Monte Carlo.

[9]The analysis region is the same as was used for the instrumental cuts: r<550cm and NHITs≥40.

Energy Threshold	Charge Current /%	Elastic Scattering /%	Neutrons /%
$T_{eff} = 6.75$MeV	0.95 ± 0.07	0.89 ± 0.09	1.43 ± 0.21
$T_{eff} = 5.5$MeV	0.92 ± 0.06	0.89 ± 0.07	1.65 ± 0.12

Table 5.8: The high level cut sacrifice for the various signal classes above two different energy thresholds.

High level cut signal sacrifice

Now that the scale factor is known it is a simple matter to measure the sacrifice of the various expected solar neutrino signals. Monte Carlo data sets are generated for each signal spanning the same time period as the selected run list. The CC Monte Carlo was generated at $10\times$ the standard solar model, while the ES and NC signals were generated at $50\times$. With these scale factors the statistical error in the sacrifice measurement is much smaller than the final systematic error.

Once the Monte Carlo sacrifice from the high level cuts has been determined, the real sacrifice is obtained by the combining the scale factor with the Monte Carlo measurement. The sacrifice of the three signals above two analysis thresholds relevant to CC, ES only, or full CC, NC, ES flux measurements are shown in table 5.8. The errors in this table are a combination of the error already quoted in the scale factor and the statistical errors in the measurement from the Monte Carlo data sets. The increased sacrifice for neutrons is exactly as expected from the difference between γ-rays and electrons from the θ_{ij} cut. The charge current and elastic scattering measurements are consistent both with the reduced energy threshold and with each other. This occurs as the sacrifice is not a strong function of energy in the region where the measurement is made,so the difference in the energy spectrum of the two event classes does not cause a large difference in the integrated sacrifice.

Although it is now possible to combine the sacrifice from the three different sources, one major source of the systematic error in the sacrifice measurement has not yet been considered. The stability of the sacrifice over time is very important. This should be measured before the final sacrifice number is considered.

5.4 Monitoring the sacrifice

The goal of monitoring the sacrifice is to measure the systematic error in the sacrifice measurements due to the change in detector conditions over time. The same measurements need to be made with respect to the position of the sun for study of the systematics of the day-night asymmetry. It is difficult to make these measurements using calibration sources alone, as data is only taken periodically, and is not taken with data cleaning in mind. For example the ^{16}N stability runs are taken with the external trigger enabled. However as they are all taken in the same way the stability of the sacrifice can still be monitored. A bigger concern using calibration sources is the sparsity of data, particularly at the begining of the data set.

Due to the problem with using calibration data attempts are made to monitor the sacrifice of the most important cuts using the data. This can be done using PulseGT data in some cases, in others events from the background wall can be used. These monitoring techniques can be used to check the sacrifice as a function of time, as well as a function of the position of the sun. The stability of individual cuts and of groups of cuts can be studied, but ultimately the overall sacrifice stability is the desirable measurement. This fact needs to be kept in mind as stability studies are made.

The sacrifice stability of some cuts is discussed elsewhere in this thesis. In chapter 4 the neck and OWL cuts are discussed. The OWL cut is studied using PulseGT data. It is found that the cut sacrifice is stable to 3×10^{-5} as a function of time, and that the day-night sacrifice difference is even smaller at 9×10^{-6}. The neck cut is studied in a different way using events from the background wall identified using the $R_{\rm neck}$ parameter. The study was carried out to investigate possible source effects in the measurement of neck cut sacrifice. These effects are not found, and deviations from the mean sacrifice value are even smaller, and the systematics due to stability of the neck cut sacrifice are dominated by other effects. In chapter 6 the livetime corrections from the burst cuts are discussed in detail, including the change in the corrections as a function of time.

5.4.1 Instrumental cuts stability runs

The first study that can be carried out to test the stability of the instrumental cuts is to apply them to the ^{16}N stability runs. These runs were taken to monitor the energy scale of the detector as a function of time. For this reason these events were taken with the ^{16}N source in triggered mode. As previously discussed data taken in this fashion is not ideal for the measurement of sacrifice. However these effects should be the same for each run, allowing these data to used to monitor the stability of the sacrifice.

The sacrifice from each run is shown in figure 5.17. The systematic error due to stability is related to the spread of the points around the best fit. This measurement is found using the median absolute deviation from the best fit value. Using this technique the systematic error is measured to be 10% of the sacrifice.

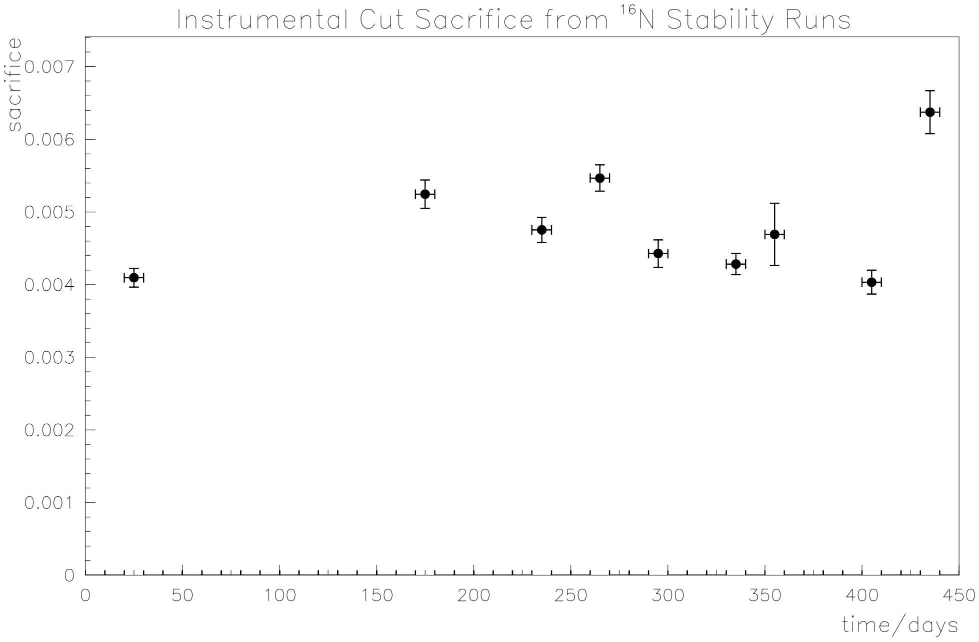

Figure 5.17: The sacrifice measured by the ^{16}N stability runs.

5.4.2 Sacrifice due to sharkfins

Both the QCluster and FGC cuts are liable to sacrifice events due to pile-up with sharkfins. This occurs due to the naturally very high rate of sharkfin events in the detector, and thus sacrifice due to pile-up with a sharkfin is much more likely than sacrifice due to pile-up with another class of instrumental background. This sacrifice acts in addition to the natural sacrifice of the two cuts and is susceptible to any change in the sharkfin rate. The rate of sharkfin pile-up therefore needs to be monitored and can be done using PulseGT data.

Sharkfin-like events in PulseGT can be identified simply as those events that fail the QCluster cut. Any PulseGT events that fail this cut are likely to have done so due to the occurrence of a sharkfin event. Furthermore by not making any further restrictions the analysis stays sensitive to any other class of event that fails the QCluster cut piling up with other classes of events. The distribution of the rate of sharkfin pile-up events as a function of time and solar zenith angle is shown in figure 5.18

As for the OWL cut the variation of the pile-up rate of sharkfins does not contribute significantly to the total error in the sacrifice. The median deviation from the best fit is only 10^{-5}. The day-night sacrifice differential is $(6 \pm 2) \times 10^{-6}$. This will not contribute significantly to any day night measurement.

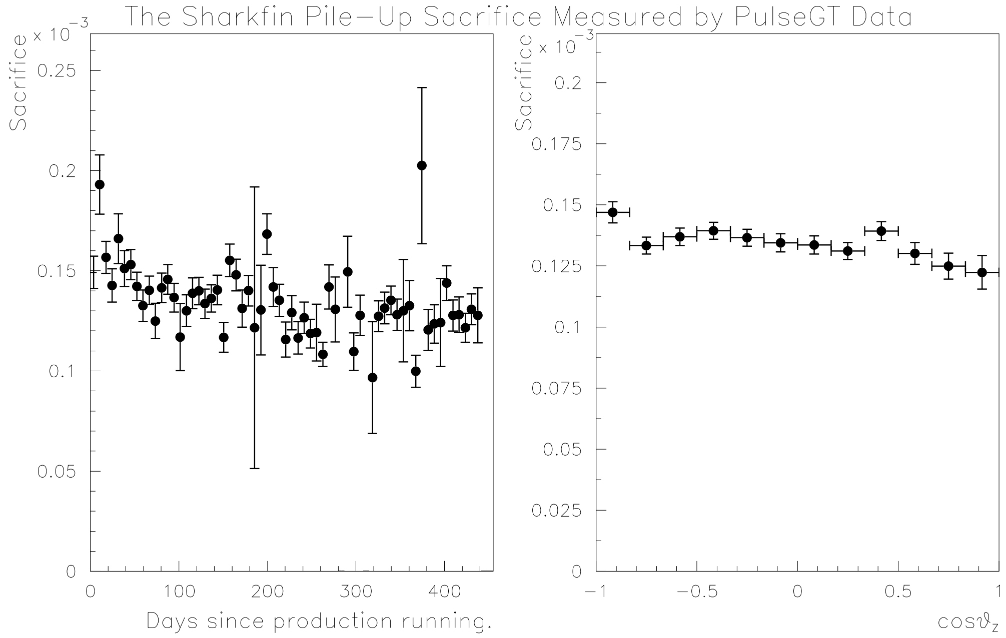

Figure 5.18: The sharkfin pile-up rate measured by PulseGT data as a function of of time and solar zenith angle, θ_Z.

5.4.3 AMB Stability

As the AMB cut has the largest sacrifice of any of the instrumental cuts, the stability of the cut warrants further study. Techniques for carrying out these checks are discussed in [79]. The AMB cut has four parameters that are calibrated using ^{16}N data: the pedestal, the mean, and two parameters which give the functional form for the RMS.

The first of these is the pedestal measurement, which can be tested using NHIT=0 PulseGT data. While this does not give the required pedestal directly[10], any changes in the pedestal should be followed by this measurement. The change in the AMB pedestals as a function of time is shown in figure 5.19. The pedestals are generally stable to 0.5%, except for the peak as a function of time. There is a steady rise in the pedestal after day 200 of production running. The rise corresponds to a 1% change in the pedestal by the end of the data set. Applying these pedestal shifts to the AMB calibration when measuring the sacrifice shows a change in the sacrifice of $\pm 3 \times 10^{-5}$. Once again this is much smaller than many of the other systematic errors on the sacrifice.

The AMB sacrifice depends upon more than the pedestal, these calibration constants can only be measured by the use of calibration data. Unfortunately there is no ideal source

[10]The pedestal is derived from the extrapolation of calibration data to 0 hits.

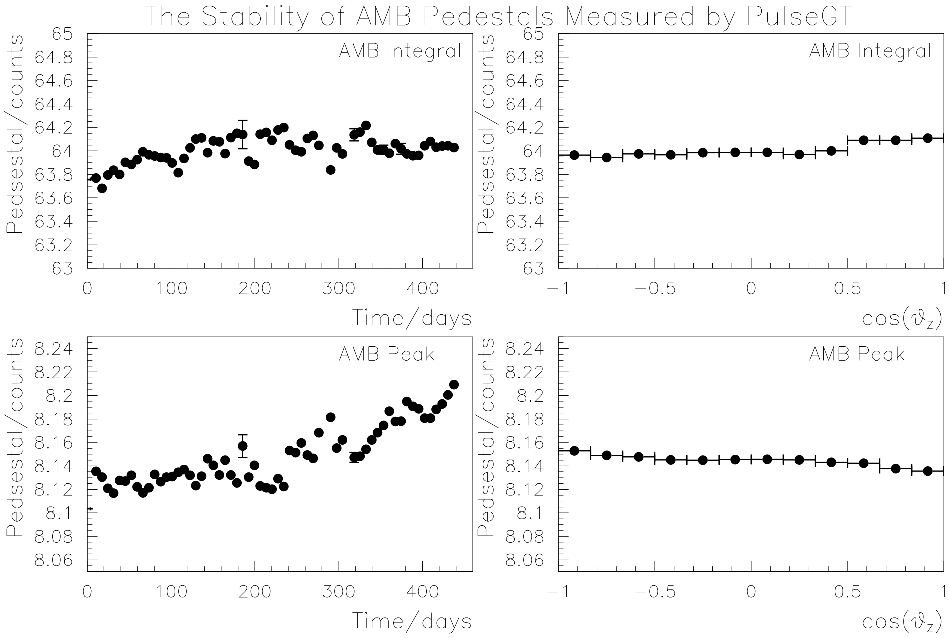

Figure 5.19: The AMB pedestals for the peak and integral measured by PulseGT data, as a function of time and the position of the sun.

data that spans the data set. The ^{16}N stability runs were taken using the external trigger, which means this data cannot be used to measure AMB sacrifice directly as the AMB is dependant on the timing of the hits with resect to the trigger window. These runs do however span the data and can be used to provide a measure of the relative change in the sacrifice. The AMB sacrifice from each run is shown in figure 5.20

The first ^{16}N stability run was taken at the same time as the ^{16}N data used to calibrate the AMB cut and to measure the sacrifice of all the cuts. This point therefore tells us the expected sacrifice of the AMB cut for triggered ^{16}N data. The plot shows that the sacrifice of the AMB cut is always higher than the first point. This suggests one or both of two things:

- The AMB calibration has drifted.
- The effect of the external trigger has drifted.

The first of these effects means that the sacrifice of the cut really has changed, while the second would tell us that this is not the case. The increase in sacrifice should therefore be included as a one-sided systematic error. The mean deviation of the later runs from the the first run shows an increase of 80%. To find the systematic error on the sacrifice from this measurement the sacrifice of the AMB must be understood. This can be found using

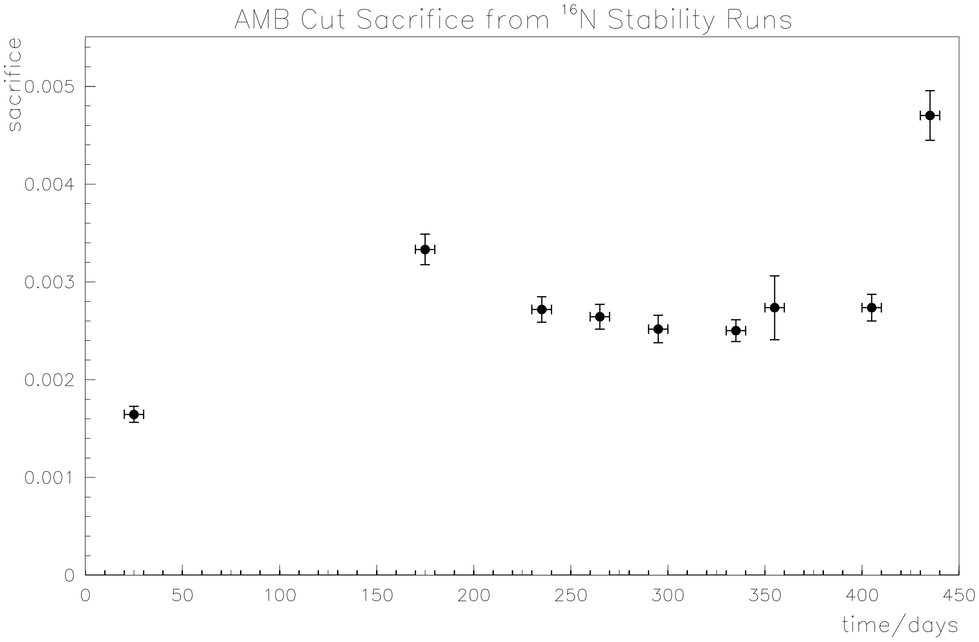

Figure 5.20: The sacrifice of the AMB cut for ^{16}N stability runs.

the same technique as was used for the instrumental cuts. It is found that the sacrifice of the AMB cut is 0.14% giving an additional systematic of +0.11% to the total sacrifice.

5.4.4 QvT stability

The QvT cut is one of the highest sacrifice cuts in the suite of instrumental cuts. It is also one of the least robust cuts, as it is based upon the calibration of a single channel. It is therefore important to ensure that the sacrifice of the cut does not vary by large amounts, and to attempt to measure the variation.

The QvT cut works by searching for a tube with high charge that is early. Both these requirements work by comparing the channel measurements of the highest charge tube to the remaining tubes in the event. The most likely way for an event to fail the QvT cut is for a noise hit with high charge to occur early in the event. While some of these noise hits will have real high charge, a sharkfin for example, some fraction of these hits will be due to mis-measured charges, due to bad ADCs or bad ECA constants. Events containing sharkfin pile-up have already been discussed, and these events are likely to fail other cuts as well. Mis-measured charges are subject to change, and unlikely to cause the event to fail other cuts. The time variation of this effect on the data must therefore be measured. The rate of noise hits on bad channels can be measured using PulseGT data,

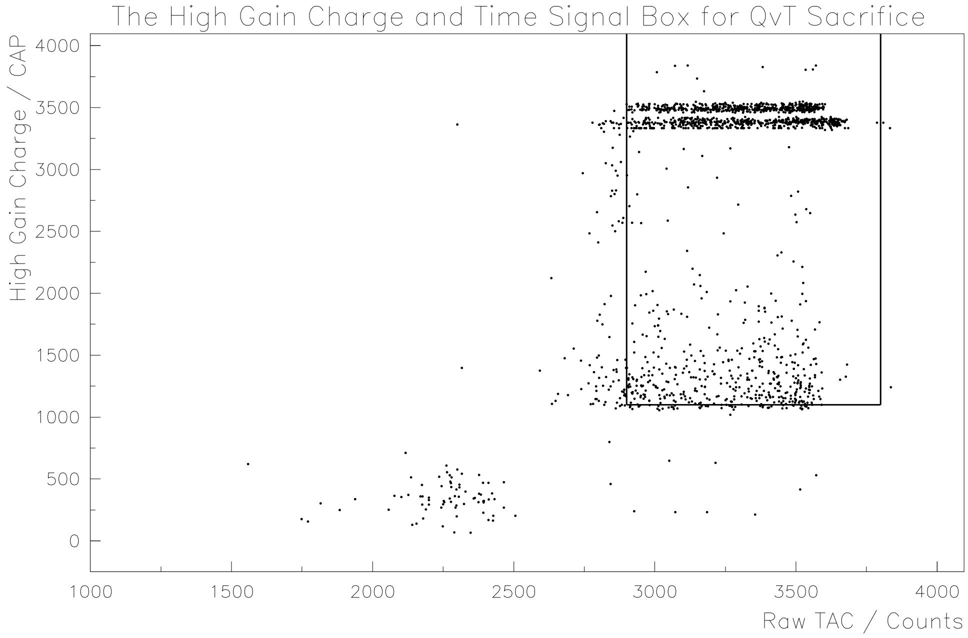

Figure 5.21: The charge and time distributions of the highest high gain charge in D_2O ^{16}N events sacrificed by QvT

and this gives us a measurement of QvT stability.

As stated the QvT cut works by comparing the highest charge tube to the rest of the tubes in the event. This means that it is not possible to define a signal box directly in which to look for bad charges in PulseGT data since there is 'event' to compare the hit to. In order to avoid this, the time and high gain charge distribution of the highest charge tube in sacrificed ^{16}N events is examined as shown in figure 5.21. Only the raw TAC measurement is considered to ensure that effects like the TAC curl are taken care of in the correct way. The plot shows the signal box for the high gain charge cut. In the figure there are hits well away from the signal box, these fail QvT due to the low gain charge analysis. For this reason this charge is also considered in this analysis, using the same signal box. The box boundaries are chosen such that the majority of the hits will be found. There are two exceptions to this, the boundaries for the invalid TAC measurement and negative rail are the same as in the QvT cut. The fraction of PulseGT events where the highest charge tube, either in high or low gain falls in this box is used to monitor the sacrifice. If the event fails the QCluster cut however, the event is not counted, as the effect of these events causing sacrifice has already been considered.

The sacrifice measured in this way is shown in figure 5.22. There are quite clearly two effects. There is a background rate of failures of 0.046%, measured by the median

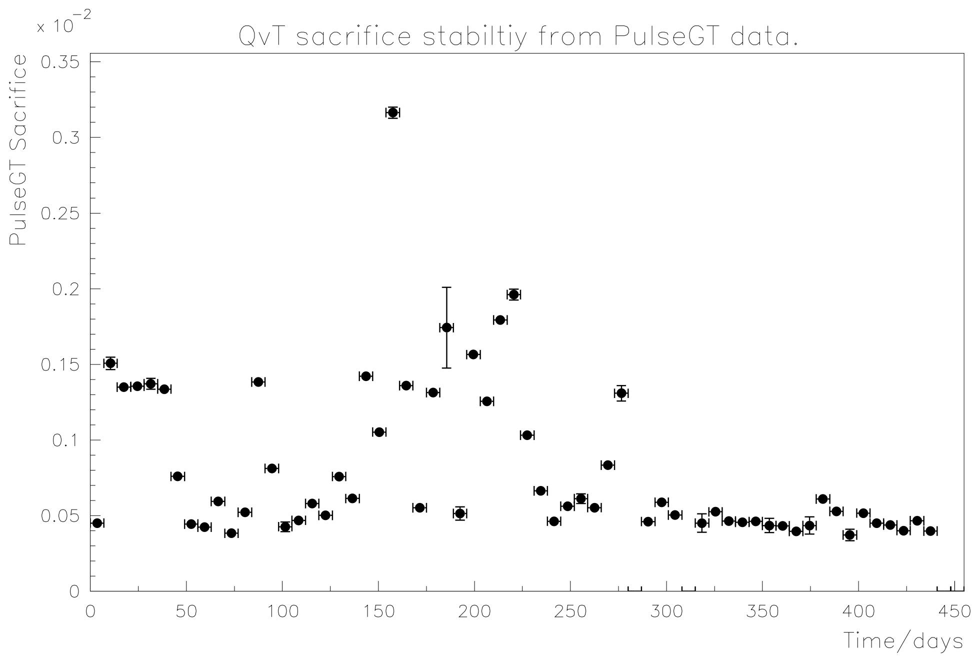

Figure 5.22: The QvT PulseGT sacrifice as a function of time.

value of the data in that region. There are then deviations upwards from this value. These are the noise hits on bad channels. The ^{16}N data cleaning data was taken at the start of December 1999, which corresponds to day 30. At this time the effect of the bad card is quite clearly visible. When the PulseGT measurement is carried out with the ^{16}N data the sacrifice is measured to be $(0.13 \pm 0.01)\%$, in agreement with that measured from the neutrino data at the same time. If hits on the bad card are discounted then the sacrifice is reduced to $(0.03 \pm 0.01)\%$. This agrees within statistics with the baseline measurement in figure 5.22, and with the reduction in the sacrifice measured by exclusion of the bad card in the ^{16}N data. Therefore the excursions above the baseline level give the measurement of the stability of the QvT cut, due to bad channels. As the ^{16}N data with the bad card removed agrees with the baseline level, this suggests that the baseline is not due to hits on bad channels, but is instead due to real hits with high charge, such as a PMT $\beta\gamma$s. This baseline sacrifice is still included in the ^{16}N measurement as the removal of the bad card does not effect this.

The livetime weighted PulseGT sacrifice is 0.083%. With subtraction of the baseline this gives a value of 0.037% excess sacrifice due to bad channels in QvT, weighted over the livetime of SNO. This is a correction factor that should be added to the measurement of the sacrifice. The statistical error in the measurement is tiny due to the the very high statistics of the PulseGT data. The systematic error on this measurement is due to how

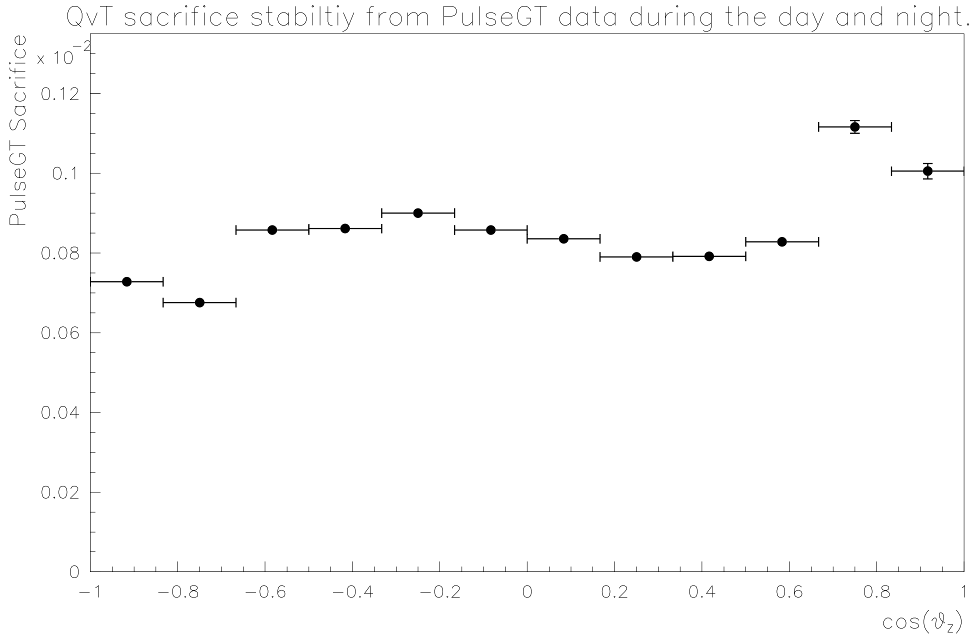

Figure 5.23: The QvT PulseGT sacrifice as a function of the position of the sun.

well the PulseGT sacrifice reproduces the sacrifice in the data. This can be obtained from the ^{16}N measurement of the sacrifice. The cut has intrinsic sacrifice as well as the sacrifice due to bad channels, this is measured using ^{16}N data by the difference between the raw sacrifice and the sacrifice with the bad card removed. The additional sacrifice is $(0.088 \pm 0.007)\%$ compared to the PulseGT measurement of $(0.098 \pm 0.008)\%$. These numbers are compatible, but the scale factor between them is used to correct the correction in the propagation of the error. The final correction is found to be $(0.033 \pm 0.003)\%$. This process works because the hits causing the sacrifice are due to noise, on the bad tubes not direct light. For this reason the extra sacrifice does not scale with NHIT.

The distribution of the deviations from the baseline shows a distinct temporal feature. The last deviation occurs around day 275, the start of August 2000. This corresponds exactly to the fix of the bad ADCs [50]. Given this time dependence there will be difference in the sacrifice between the day and the night, as shown in figure 5.23. The differential sacrifice, after correction the mis-measurement using PulseGT, is $(0.003 \pm 0.001)\%$.

5.4.5 Bad channels and QCluster

When the bad card analysis was carried out with the ^{16}N data two cuts were seen to be effected. The QvT cut showed the biggest effect, but the QCluster also showed a reduction in the sacrifice. Having made the measurement of the effect of bad channels on QvT over time, we can use this measurement to find the correction for QCluster.

When the bad card was removed from the ^{16}N data the QvT correction was $(0.088 \pm 0.007)\%$ compared to $(0.030 \pm 0.004)\%$ for QCluster. This gives a scale factor for the different effect on the two cuts, which can be used to calculate the QCluster correction from the QvT correction already calculated. However when carrying out this procedure the error bars for the two corrections will be correlated, it therefore makes more sense to calculate the combined bad charge correction instead to ensure that the correlated errors are handled correctly in the final sacrifice measurement. The total reduction in the sacrifice upon removal of the bad card is $(0.118 \pm 0.008)\%$. Using this instead of the QvT reduction in the calculation gives a final bad channel correction of $(0.044 \pm 0.006)\%$.

5.4.6 High level cut stability

Just as instrumental cuts are susceptible to increased sacrifice due to bad channels and changing detector conditions, the same is true of the high level cuts. These are all sensitive to the quality of the timing calibrations, through the dependence on the fitter, and in the case of ITR much more directly. Any stability error in the measured sacrifice does not get passed to the Monte Carlo procedure for measuring the cuts. The stability error is instead combined with the final high level cut measurement directly.

The stability of the different high level cuts is shown in figure 5.24. The stability of the θ_{ij} and figure of merit cuts is very good, suggesting that the change in the number of online channels has not had a big impact on these cuts and that neither has the small changes in timing calibration had a big effect on reconstruction. The ITR cut on the other hand shows very large deviations. The ITR cut in its design is less robust than cuts like θ_{ij} due to the aggressive choice for the time window. However the size of the variation is still disturbing, suggesting that a $\pm 0.4\%$ systematic error is required for the stability of the ITR cut.

The systematic error for the ITR cut is by far the largest error seen in the sacrifice analysis, and for this reason it would be nice to attempt to reduce it. However this is not an easy problem. Events in the background wall cannot be used, as it is not possible to define a region in parameter space that is clean enough and has enough statistics to improve upon the error measured by the ^{16}N source. The problem is that events in the D_2O are contaminated by mis-reconstructed events from the H_2O and events in the H_2O

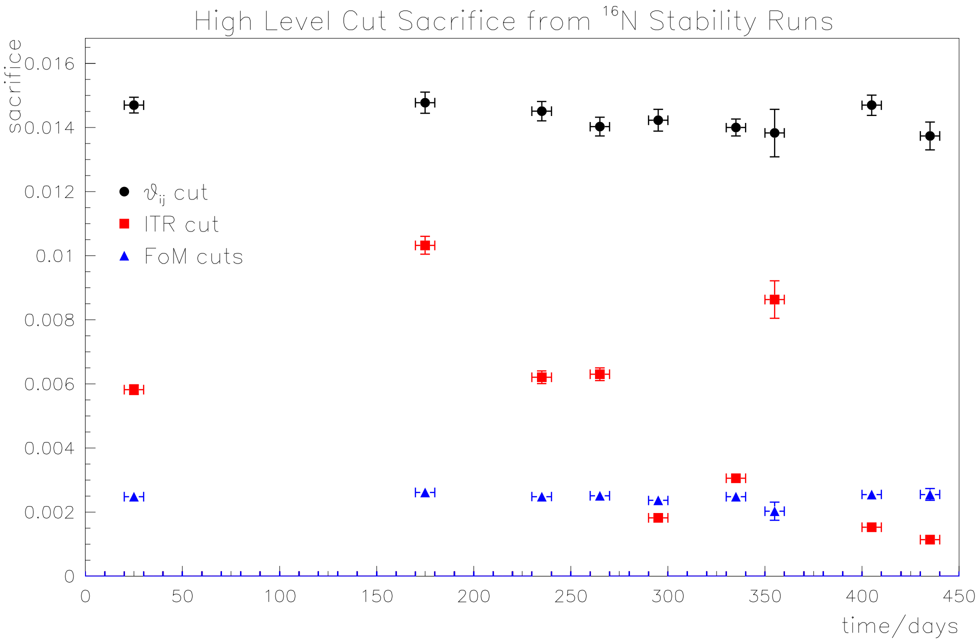

Figure 5.24: The stability of high level cut sacrifice from ^{16}N stability runs.

by events from the PSUP. Another problem with this method is that is suffers from the same problem as the sacrifice measurement at high NHIT from ^{16}N . The events are in the tail of the NHIT spectrum, and are therefore unusual, making them unsuitable for measuring sacrifice, even for just stability.

The inability to find a region in which the ITR stability can be tested is also problematic for the analysis of the day-night asymmetry. Without the ability to check the ITR sacrifice in both data sets independently it cannot be shown that the ITR sacrifice is the same during the day and night. An extreme approach is therefore taken, in which it is assumed that this systematic error cannot be cancelled in the analysis, in other words the measured day and night rates each have a systematic error due to ITR sacrifice of $\pm 0.4\%$ which is independent of the error on the other data set.

5.5 The combined sacrifice measurement

Now that stability studies have been carried out it is possible to combine the sacrifice measurements and errors to find the final sacrifice values for the entire suite of cuts. The value for each signal can then be combined directly with the extracted fluxes, with the errors on the sacrifice adding directly into to systematic error on the flux.

The instrumental cuts were measured to have a sacrifice of $(0.311 \pm 0.007)\%$, with a

Energy Threshold	Charged Current /%	Elastic Scattering /%	Neutrons /%
$T_{eff} = 6.75$MeV	$1.46^{+0.54}_{-0.45}$	$1.40^{+0.54}_{-0.46}$	$1.94^{+0.58}_{-0.50}$
$T_{eff} = 5.5$MeV	$1.43^{+0.54}_{-0.45}$	$1.40^{+0.54}_{-0.45}$	$2.16^{+0.55}_{-0.46}$

Table 5.9: The total sacrifice for the various signal classes above two different energy thresholds.

low-sided systematic error 0.021% due to ECA validation. A number of stability measurements have also been made. Bad charge channels provide a correction of (0.044 ± 0.006)%. There are then the various cut stability measurements that have also been made along with the overall stability from the ^{16}N stability runs. To use both of these would result in double counting so only the individual cuts are used. The stability error is $^{+0.11}_{-0.007}$%, where the high end is totally dominated by the AMB cut. These errors are all independent and can therefore be added in quadrature to give a final instrumental cut sacrifice of: $(0.36)^{+0.11}_{-0.02}$%.

Reconstruction sacrifice has been determined to be < 0.3%, but due to the difficulty in making an accurate measurement, or with combining this with the expected spectrum, no direct measurement has been made. In order to combine this with the other sacrifice measurements, the upper limit is converted into a measurement of (0.15 ± 0.15)%. The sacrifice of the high level cuts is shown in table 5.8. The stability error on these measurements is total dominated by the ITC cut which provides a measurement of ± 0.4%. This error must be combined in quadrature with the errors in table 5.8 to give the total error on the high level cut sacrifice.

The measurements of the sacrifices from the three sources can now be added together to give the total sacrifice[11]. The combination of the errors is not quite so straight forward. The dominant errors on the high level and instrumental cuts are stability errors measured using the ^{16}N stability data. These errors are not independent and are therefore added together. This error is then combined in quadrature with the error from the reconstruction sacrifice to give the total error.

Table 5.9 shows the final combined sacrifice results for each signal class above the two different analysis thresholds. The errors are dominated by the error in the ITR measurement. However the error on the sacrifice is still much smaller than the systematic errors due to energy calibration and reconstruction. Indeed the error is also smaller than the statistical error bar on the CC flux measurement.

[11]Because each the sacrifice from each source is small the corrections due to the statistical correlations (the probability that one event can fail two different sets of cuts by chance) can be ignored.

5.6 Sacrifice at high energy

Although the measurement just made is valid until the end of the ^{8}Li spectrum, there are still interesting events at higher energy. A good example of this would be events due to hep neutrinos. It is therefore important to understand the sacrifice in this region.

5.6.1 High energy sources

At this high energy there are no tagged Čerenkov sources available, but there are still sources of events that can be used. As has been shown previously the laserball does reproduce the sacrifice measured by ^{8}Li and ^{16}N of the instrumental cuts in the centre of the detector. The source cannot be used to measure the sacrifice of the high level cuts, however, or indeed to measure reliably the sacrifice near the acrylic vessel, where the majority of the fiducial volume is. For these two reasons the laserball is of limited use.

We do however have access to the pT source which provides 19.8MeV γ-rays. The principal problem with this source is that the events are untagged. This is serious problem for measurement of the sacrifice, and in practice, this means that an upper limit is the best that can be hoped for. There is also a secondary problem in the use of the pT source. It produces a large number of neutrons in the production of the γ-rays. These can cause problems due pileup with the γ-ray events.

Despite these problems it is possible to carry out a sacrifice analysis using pT data by playing off the instrumental and high level cuts against each other. For both the ^{16}N and ^{8}Li data inside the D_2O these two sets of cuts were uncorrelated in sacrifice. Under the assumption that this statement is still true for pT data each set can be used to provide a set of events from which the sacrifice of the other set can be measured. Using this technique it is possible to measure an interesting upper limit for the sacrifice at these high energies.

In principal it is also possible to measure high energy sacrifice using samples of electrons from muon decay. These can in principal be identified with high purity by the preceding muon. However there are serious problems encountered attempting this analysis. The first is a statistical problem. There are only a small number of these events a year, resulting in a low statistics sample of events. This is not too serious however as the sample of events under study is expected to have even smaller statistics. The major flaw in this approach is due to the response of the detector following a muon. For such high energy events there are always retriggers, and the pileup of a retrigger with the electron is highly likely. This means that the event can have a number of tubes with highly abnormal times and charges, which are exactly what the instrumental cuts use to remove backgrounds. The ESUM signal is also unstable due to the large pulse following the muon, resulting in extremely large AMB sacrifice. These facts combine to make any sacrifice analysis using decay electrons unreliable, and this approach will not be considered further in this thesis.

5.6.2 Data sets and data reduction.

In order to measure the sacrifice from the pT source a data set must first be obtained. The source was run in February 2001, and was placed in three positions on the z-axis, 0, -250 and -500cm. The data taken from these runs is then processed above 140 hits, to begin initial data reduction. With the additional application of an upper NHIT limit at 200 hits there are 17028 events in the initial data sample.

While we want to measure the sacrifice of the entire suite of instrumental and high level cuts, we are starting from a position where the behaviour of cuts is understood at some level. This means that we can apply any cut where the sacrifice is well understood. The following cuts are applied to reduce the data set:

- A reconstruction cut

  As we know the location of the source, a cut on the reconstructed distance from the source is the most powerful way to reduce the number of background events in the data set. By examining the distribution of events around the source the ideal fiducial volume can be found. The relevant distribution is shown in figure 5.25, and the choice is made to cut any event that reconstructs more than 2m from the source. After the combination of this cut with a fiducial volume cut at 550cm there are 4510 events remaining.

- The retrigger cut

  The sacrifice of the retrigger cut is solely due to the coincidences between events, and as such there is no bias to the events removed by this cut. It is therefore safe to apply the cut to the data. With the application of this cut the data set is further reduced to 4019 events.

- The OWL cut

  As discussed in chapter 4 the OWL sacrifice is dominated by noise on the OWL tubes. The OWL cut is very powerful and it is therefore used to reduce the dataset. Despite the power of the OWL cut only a further 5 events are removed. This suggests that the cuts that have already been applied have removed the majority of instrumental backgrounds.

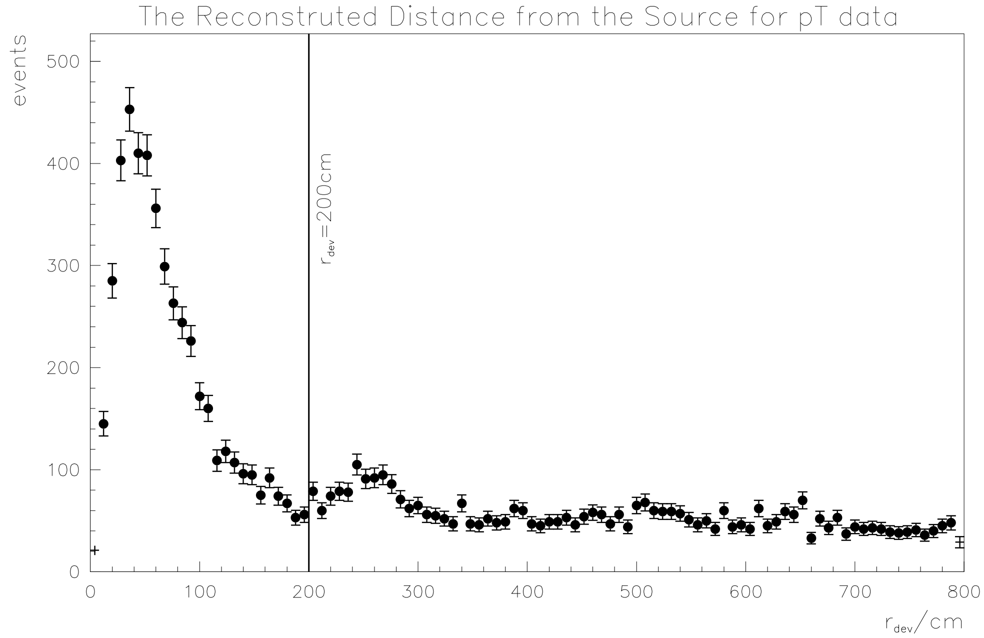

Figure 5.25: The reconstructed distance from the source for events in the pT source run.

- The q/NHIT cut

  The q/NHIT cut has been shown to be highly effective at removing pickup, while accepting events with real light. The sacrifice of the cut is extremely small and if there is any energy bias, it will be to reduce the sacrifice at higher energy, as the spread in the value of R_Q is reduced due to increased hit statistics. The application of the q/NHIT cut does not remove any additional events from the data sample, suggesting that there is no remaining contamination by pickup events. This is not surprising as pickup events would not be expected to reconstruct at all.

- The ITC cut

  This cut was shown with the ^{16}N data to be effective at removing events with pile-up. This is important for the pT study due to the high neutron rate. The sacrifice of the cut is how ever very small, and as is the case for the q/NHIT cut, this should only improve at higher energy. The ITC cut removes 7 events from the sample. These are events where reconstruction has failed to remove pile-up events.

The sacrifice of the remaining cuts can then be measured using this data set of 4007 events.

5.6.3 The instrumental cuts

The sacrifice of the remaining instrumental cuts is now measured using this data sample. Despite the initial data reduction that has occurred, that data is not expected to be a pure sample. Thus in order to make a better measurement the high level and figure of merit cuts are applied to the data. Studies at lower energy suggest that the two sets of cuts are essentially uncorrelated, so the high level cuts are used to clean the data sample, before the sacrifice measurement is made.

At some level the cuts are correlated, and by ignoring this correlation, this study will underestimate the sacrifice. However using the correlation measurements at lower energy this effect is expected to be small. Using table 5.4, 2% of the events failing the instrumental cuts also fail the high level cuts. This measurement can not, unfortunately, be repeated using a different source, so the energy dependence of this value is unknown. There is however no reason to believe that there should be any energy dependence to the correlation. As the goal of this study is to find an upper limit on the sacrifice, a correction for the residual correlation will be added to the final sacrifice number. To do this the 3σ upper limit of the ^{16}N measurement, 3.5% of the instrumental cut sacrifice will be used as this correction.

Once the high level cuts have been applied to the data there are 3847 events remaining. These are used measure the combined sacrifice of the instrumental cuts not used in reduction of the data. The result is shown in figure 5.26. The data fits a flat distribution, with a best fit value of $(0.96 \pm 0.16)\%$. The error bar is the statistical error bar on the fit.

The measured value has to be corrected for the cut correlations, and the OWL cut sacrifice. These corrections amount to 0.04%, giving a measurement of $(1.00 \pm 0.16)\%$ for the instrumental cut sacrifice. However because instrumental background may still have been present in the dataset this measurement is an upper limit.

5.6.4 High level cuts

The same analysis can be carried out to measure the sacrifice of the high level cuts. In this case the suite of instrumental cuts[12] is used to provide the sample of pT events. This process yields 3914 events for the data set. As the correlations between the two sets of cuts has already been corrected for in the instrumental cut study, it does not need to be considered in the high level cut study as well.

Therefore the combined sacrifice of the high level and figure of merit cuts can be measured with no further work. The resulting distribution is shown in figure 5.27. As was the case for the instrumental cuts, the high level cuts are also fit well by a flat distribution.

[12]With the exception of the burst and NHIT burst cuts.

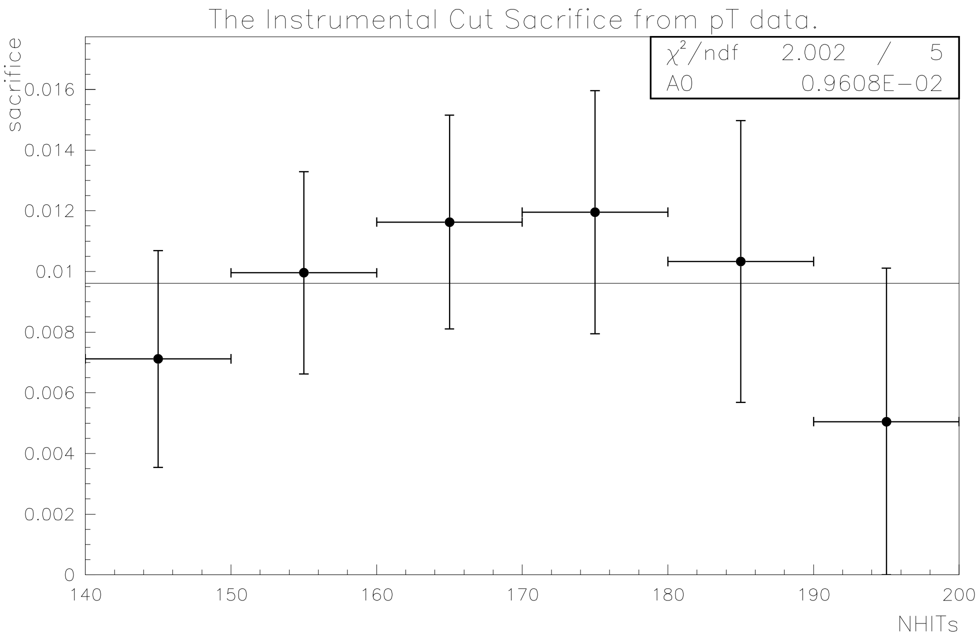

Figure 5.26: The measured instrumental cut sacrifice from pT data.

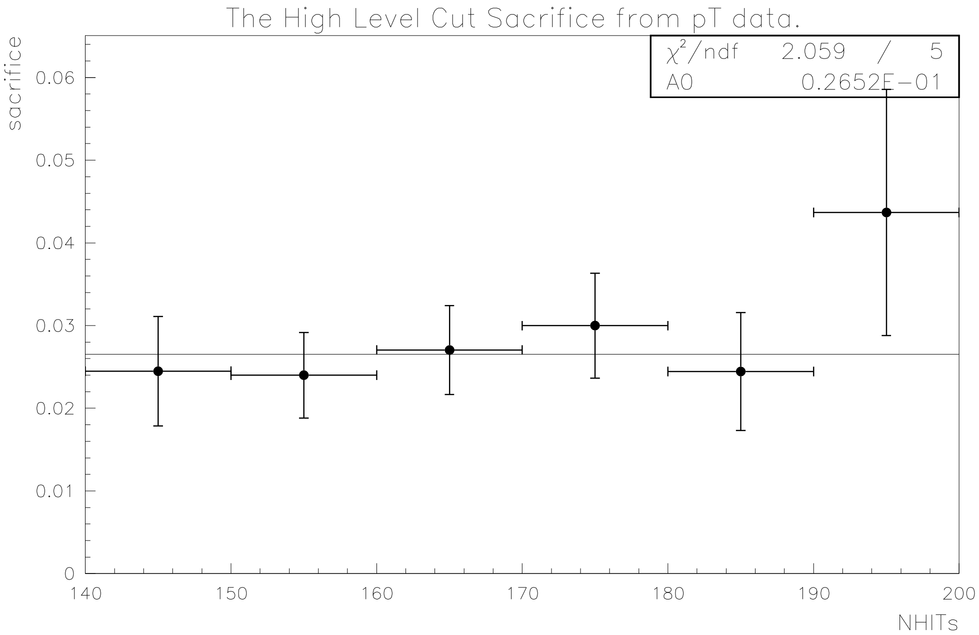

Figure 5.27: The measured high level cut sacrifice from pT data.

In this case the best fit value is $(2.7 \pm 0.3)\%$. This is also an upper limit, for two reasons in this case. Firstly the data sample may contain background events, which increase the measured sacrifice. Also the pT source is a γ-ray source. As has already been discussed the sacrifice of the θ_{ij} cut is greater for γ-rays than for electrons. By using a γ-ray source the sacrifice is overestimated.

5.6.5 Combined sacrifice at high energy

The pT source has been used to set an upper limit on the sacrifice of both levels of cuts at high energy. These measurements can be combined with each other to provide a limit on the total sacrifice at the energy scale of the pT source. Combining the two measurements however is not enough. As was the case for the measurement of the sacrifice for the main neutrino signal effects such as time dependence must also be considered.

Unfortunately, there is no way to monitor the sacrifice at these high energies, so the same results obtained at lower energy must be used. By far the most important of the stability measurements, is the stability of the ITR cut. This was measured to be $\pm 0.4\%$, for ^{16}N stability runs. This error should also be applied to the sacrifice at the higher energy of the pT, as the ITR sacrifice is not a strong function of energy, except at low energy where the statistics in the number of hits is reduced. This error will be added in quadrature to the statistical errors on the two fits.

Taking this into account the measured total sacrifice using pT data is $(3.7 \pm 0.5)\%$. This is a measurement of an upper limit, and should therefore be quoted as such. This is done by using the 3σ upper limit, of 5.2% to provide a *very* conservative bound on the sacrifice at high energy. This value can be used to set a bound on the possible sacrifice within any high level analysis. A low level bound can also be set by the use of the same technique and numbers used for the sacrifice measurements for the charged current, neutral current and elastic scattering signals. When this is carried out above 12MeV a sacrifice of $(2.5 \pm 0.5)\%$ is found with the error again dominated by the ITR cut. This measure provides a lower bound on the sacrifice of hep events in the high energy region.

5.7 Contamination

As well as measuring the sacrifice it is also necessary to measure the contamination. This is the residual number of instrumental backgrounds inside the neutrino signal after the application of all the cuts. In chapter 4 the contamination of neck events was measured, however the contamination of all background classes is the important quantity to understand. Rather than attempt to measure the contamination of all classes individually the best way to do this measurement is to include all backgrounds together. The method that has been developed to do this is known as a bifurcated analysis [78] [80] [57].

The bifurcated analysis works by comparing the effect of two cuts (or sets of cuts) on a data set and using each cut to calibrate the other. Provided that the two cuts are independent the leakage of each cut is measured allowing the contamination to be determined. By using sets of cuts that remove all the instrumental backgrounds (at least at some level) then this technique is sensitive to all backgrounds, and the resulting contamination measured being the total contamination.

The basics of a bifurcated analysis are quite simple. Assuming two cuts 1 and 2, applied to a data set of T events then the following equations are true provided that the two cuts are uncorrelated:

$$\alpha = \epsilon_1 S + x_1 B \tag{5.1}$$

$$\beta = \epsilon_2 S + x_2 B \tag{5.2}$$

$$\gamma = \epsilon_1 \epsilon_2 S + x_1 x_2 B \tag{5.3}$$

$$T = S + B \tag{5.4}$$

where S is the total number of signal events, B is the total number of background events, ϵ_i/x_i is the acceptance of cut i for signal/background, and $\alpha/\beta/\gamma$ are the number of events pass cut 1/cut 2/both cuts. Given that we can measure the sacrifice of the cuts using the methods outlined earlier in this chapter, we now have 4 equations in 4 unknowns. This can then be solved analytically to give us the contamination.

In order to measure the contamination the cuts are split into three different sets. The first of these reduces the data set and includes the fiducial volume and energy cut. The other two sets are used in the bifurcated analysis. One of these sets is the combination of the ITR and θ_{ij} cuts. The other set of cuts that is used is the maximal set of instrumental cuts that are independent of time and geometry distributions. Cuts that use these distributions including QCluster and QvT and neck cannot be used, as they could introduce unwanted correlations that spoil the bifurcated method. The three sets of cuts are shown in table 5.10. The sacrifice of the two sets is determined using the method described earlier in this chapter and in chapter 6 for the burst cuts. The sacrifice required is the sacrifice after the application of the pre-applied cuts and this is taken into account using the cut correlations. The sacrifice of set 1 is $(0.7 \pm 0.4)\%$ assuming the signal is dominated by electrons, as is the case above 6.75MeV. Set 2 has a sacrifice of $(0.35 \pm 0.10)\%$, dominated by the AMB and NHIT burst cuts.

Set 1	Set 2	Pre-applied
ITR	Retrigger	$T_{eff} > 6.75$
θ_{ij}	Burst	r<550cm
	Q/NHIT	junk
	AMB	QvT
	OWL	Crate Isotropy
	NHIT Burst	FTS
	OWLEHi	Neck
		ESUM
		QCluster
		ITC
		FGC
		Figure of Merit Cuts
		Muon Follower Short

Table 5.10: The three sets of cuts used in the bifurcated analysis.

		Set 2	
		Passed	Failed
Set 1	Passed	1196	7
	Failed	3	16

Table 5.11: The result of applying the two sets of cuts to the data set for the bifurcated analysis.

Of the total data set there are 1222 events that pass the pre-applied cuts. The result of applying the two sets of cuts to this data set are shown in table 5.11. Using these results and equations 5.1,5.2, 5.3 and 5.4 the contamination ($x_1 x_2 B$) can be found. This is found to be 0.1 ± 0.3 events. This result is consistent with zero and for this reason it is turned into an upper limit. Thus a 3σ limit on the contamination is found to be 0.1%. This number is very small and shows that the data cleaning process has done the job it was required to do. The value of the contamination is also small enough that it can be safely ignored in the remainder of the analysis.

The contamination has been measured inside one particular analysis region, but there is nothing special about the choice of the fiducial volume and energy threshold from the point of view of the contamination measurement. The bifurcated analysis technique can be used in any region where the sacrifice is understood. Thus, the energy threshold can be lowered to the second energy threshold used in the sacrifice calculation without difficulty. When this is carried out the contamination is found to be less than 0.2% at the 3σ level. This is still a negligible contribution to the signal, demonstrating the effectiveness of the data cleaning process in this extended data set.

The bifurcated analysis technique does suffer from some problems that prevent it being used in all analysis regions. As already stated the sacrifice of the two sets of cuts must be understood before hand. This prevents this technique being used in regions where the sacrifice techniques presented break down. The requirement that the two sets of cuts are orthogonal is also problematic, if this requirement is not met the contamination measured is too small preventing an upper limit from being determined. Techniques based upon relaxing the applied cuts have been used to search for correlations [57], but none have been observed. This is enhanced by the use of instrumental cuts based solely on charge, rate and veto tubes.

The final problem with the bifurcated analysis for the contamination measurement is that any non-instrumental background that fails the cuts forms a background to the instrumental contamination. Provided this background is small this is not a problem, this is however not always the case. Mis-reconstructed events fail the high level cuts and therefore cause a background. This is not problematic in the D_2O, but is a serious problem when extending this method into the H_2O region. In this case mis-reconstructed H_2O $\beta\gamma$s cause the instrumental contamination to be seriously overestimated.

5.8 Conclusions

In measuring the sacrifice and contamination we gain an understanding of the cuts that are used to remove instrumental backgrounds. By obtaining such small values for both measurements it is also shown that the data cleaning process is doing exactly what it has been designed to do. The one sacrifice measurement that has not yet been determined is the sacrifice of the cuts based upon time correlations. These are discussed in the next chapter.

Chapter 6

Livetime

SNO is an experiment that is designed to measure the rate of solar neutrino interactions. It is therefore important to have a very good knowledge of the total exposure time for the sample of events used in any analysis. Livetime studies are however more than measuring the total exposure time. To study any time dependent phenomena such as a day-night asymmetry the livetime must be understood in a differential fashion. The measurement is further complicated by the presence of cuts that introduce livetime corrections when applied to the data. These corrections also need to be understood, both in a differential way and with respect to the sacrifice of the cut.

6.1 Raw livetime

In order to measure livetime there is one piece of equipment that cannot be done without: a clock. The clock provides a time reference to each event seen by SNO and thus allows the livetime to be calculated.

6.1.1 The clocks

There are two clock systems used in the SNO electronics system. The first and most important of these for measuring livetime is the 10MHz Clock. The 10MHz clock time is stored on a 53 bit register, which will store times of up to 28 years before roll-over. This is much longer than the projected lifetime of the experiment. The most important thing about the 10MHz clock however is that the clock signal is transmitted through a fibre optic from a GPS receiver on the surface to the detector [38]. The 10MHz clock signal is therefore tied to the GPS system providing long term accuracy to the 10MHz clock. Because of the long validity of the clock and its accurate nature the 10MHz clock is used to control most aspects of data analysis that rely on knowledge of the time of the event.

However there is a second clock in the system that is more precise than the 10MHz clock and is used to measure inter-event timing.

The 50MHz clock is stored on a register of only 43 bits, which means that the clock can run for only 2 days before rolling-over. The clock signal comes from an external oscillator into the trigger system and is not tied to the GPS system. However because of the higher frequency of the system this clock is more useful than the 10MHz clock for analysis involving inter-event timing at short time scales. The 50MHz clock while more limited than the 10MHz clock can also be used to verify measurements made using that clock, and to replace it in situations where 10MHz clock data is not available, such as when there are problems with the GPS receiver.

6.1.2 Raw livetime determination and verification

The raw livetime of a run is a very simple measurement to make. The time of the last event is compared to the time of the first event, and the time difference between them gives the run time. Because events in a run can be stored out of time order, care has to be taken to ensure that the earliest and latest events are used. Carrying out the measurement this way does introduce a very small error, as the start and end of the run are defined by instants in time, not events. There is a small period of time after the run starts, before the first event, the same is the case at the end of the run. The maximum time between events is 0.2s due to PulseGT events, so the mis-measurement due to using this technique is very small[1].

The measurement is carried out using the 10 MHz clock to give the official raw livetime for each run. The livetime of the run is then stored in the RLDT bank (see appendix B) for that run for later use. The distribution of livetime for the runs in the neutrino data (see appendix C) set is shown in figure 6.1. The total livetime measured by the 10MHz clock is 245.6 days. The three spikes seen in the run time histogram at 24h, 36h and 96h are due to the various maximum run times that have applied through the operation of SNO.

The 50MHz clock, while not as accurate as the 10MHz clock, due to not being fixed to the GPS, can still be used to verify the raw measurement made by the 10MHz clock. The procedure for measuring the livetime by the 50MHz clock is similar to that for the 10MHz clock, except that, due to the short roll-over time, care must be taken to find these roll-overs and account for them in the livetime calculation. Operationally this is accomplished by looking for a shift from the end of the register to the beginning, and turning a counter. In order to avoid double counting due to event disordering, a lockout

[1]With 410 runs in the data set, the livetime could be up to 164s, or $8 \times 10^{-4}\%$ longer than that calculated by this method.

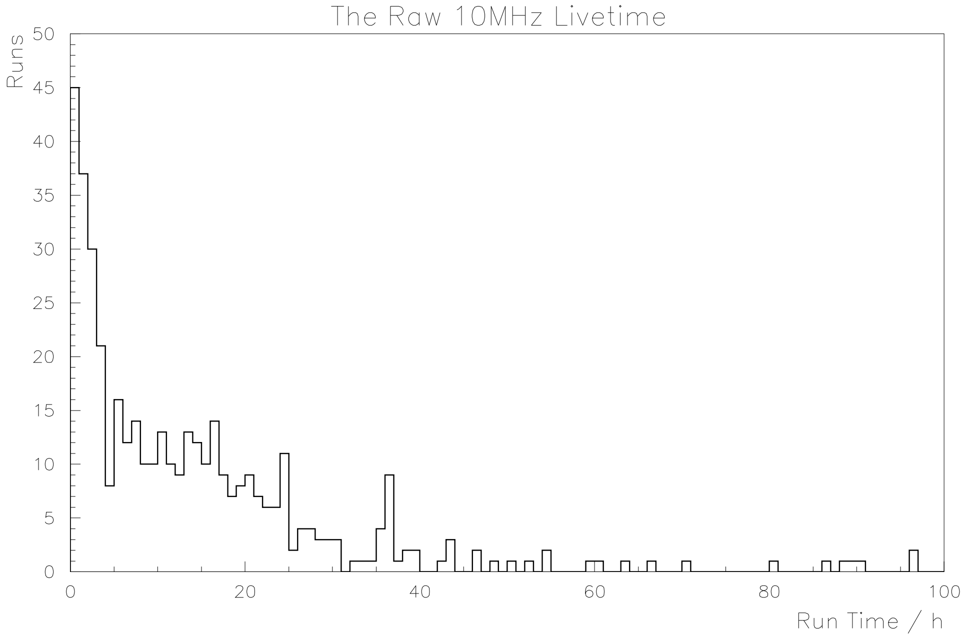

Figure 6.1: The livetime of the runs in the neutrino data set using the 10MHz clock.

on this search is imposed for 50000 events after each roll-over is found. The resulting livetime distribution can be compared to that from the 10MHz clock for a comparison.

The total live time measured by the 50MHz clock is 16.7 minutes smaller than that measured by the 10MHz clock. There is however an additional complication, the 50MHz clock does not run at exactly 50MHz, but is slightly slow. This can be seen by examining the fractional difference in measured run times, shown in figure 6.2. There are some visible outliers in the plot, but the majority of the distribution if highly peaked. Only by zooming in on this peak as shown in the right hand plot can the systematic difference between the two distributions be seen. This difference is due to the 50MHz clock running a slightly lower frequency. This can be measured by examination of the right hand plot and the result is that the 50MHz clock runs at 49.9995MHz.

This effect must be corrected for before a comparison of the two clock run time can be made. Once this correction is made the difference in the livetime measurements is reduced to 12.8 minutes. This figure comes from the outliers seen in figure 6.2. It is unknown which clock is at fault for each of these outliers, however the clocks are separate systems, so the chance that both develop a fault at the same time is small. The difference in the measurements made by the clocks gives us an estimate of the error in the raw livetime measurement by the 10MHz clock of 0.004%.

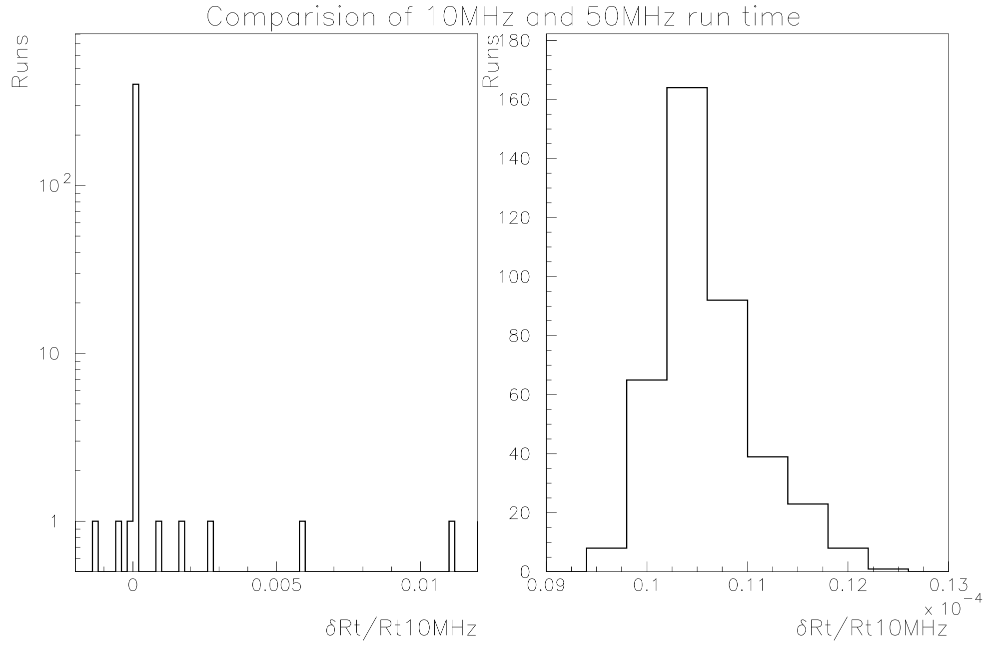

Figure 6.2: The fractional difference in the run time measurements from the 10MHz and 50MHz clocks.

Verification using PulseGT

SNO data is always accompanied by data taken from the pulsed global trigger (PulseGT). This is an unbiased trigger that takes a reading of the state of the detector at a rate of 5Hz. This trigger is used primarily to measure the noise rate of the PMTs and to monitor the detector. However it can also be used to verify the livetime due to the fixed rate.

Unfortunately life is not that straightforward. There are 105945429 PulseGT events in the data set. This corresponds to a livetime 8.9 hours smaller than the measurement of the total raw livetime. The livetime measurement is complicated by three different problems. Firstly the PulseGT is tied to the 50MHz clock, which means that the frequency of the clock must be corrected for, this however is a small correction.

The second problem is that when the NHIT monitor[2] is run, it will steal the PulseGT triggers and they are lost from the analysis. These events can be recovered by the use of the pedestal trigger in NHIT monitor events and the hit pattern of the pedestals. Recovery of these events provides another 154820 events, or 8.6 hours of livetime. This leaves a residual discrepancy of 19 minutes. This is almost the same size as the clock

[2]The NHIT monitor a tool that is used to monitor the NHIT trigger. This is carried out by causing channels on crate 18 to fire artificially. These events should cause the NHIT trigger to fire, and the NHIT monitor checks that this takes place. The artificial events used in this process are called pedestals.

discrepancy.

The final problem, and the one that stops the PulseGT being used to measure the livetime precisely is that PulseGT triggers can also be stolen by real events coincident with the PulseGT event. If an event trigger occurs at the right time the PulseGT trigger is missed. In normal running this would be expected to happen approximately once in every 40000 events. This is dependent on the rate, as the coincidence rate is enhanced during bursts. This does not however stop PulseGT data being used to verify the livetime after the corrections have been applied. Assuming the 19 minutes of livetime unaccounted for were due to the loss of PulseGT events in this way this corresponds to a capture rate of 1 event in 18000 PulseGT events, close to the expected capture rate. The remaining discrepancy could be due to further missed PulseGT or the livetime calculation from the clock. However at this point the difference between the two measurements is very small and can be ignored.

6.1.3 Differential livetime

While the total livetime is the most important value to know in the SNO experiment, knowledge of the distribution of the livetime is vital for understanding the neutrino signal in a more complete way. For example the livetime distribution throughout the year is important for studying the stability of the signal, and for looking for seasonal effects. To look for a day-night effect the distribution of livetime as a function of the position of the sun is required.

The livetime in SNO is defined by the start and stop times of each run. While it is a simple matter to turn these measurements into a total livetime number measuring the differential livetime is a more complicated process, in particular while studying data sets produced by passage through the data splitter as is done in the SNO blindness scheme.

One procedure that is used to calculate the differential livetime for any of the required distributions, is to step through all valid times defined by the run start and end times. Whenever the current time crosses a boundary in the required distribution, the amount of time stepped through to reach this point can added to the livetime histogram. The accuracy of this procedure is dependent upon the size of the time step used. Typically the time step will be chosen to be much smaller than the size of the bins in the histogram of the distribution under study. A time step of 0.1s will be used in this thesis. The modification of the livetime by the data splitter is easily added to this procedure. Instead of valid time periods defined by runs, the time periods defined by the valid data sets can be used instead.

This procedure is particularly useful when considering the livetime as a function of a variable that does not have a one-to-one mapping to universal time. A good example of this is the position of the sun. The sun is in a slightly different position in the sky at the same time from one day to the next. This means that a specific position of the sun cannot be mapped in a direct way to time of the day. However using this procedure the livetime as a function of the position of the sun is as simple to calculate as any other distribution.

Figure 6.3 shows the raw livetime distribution as a function of the cosine of the solar zenith angle ($cos(\theta_Z)$), and compares it to the same distribution for the data taking period. The fraction of total livetime is considerably greater during the night than during the day. This is seen more clearly in the livetime fraction considered as function of time of day, rather than using the position of the sun. The reduced livetime fraction during the day, is caused by activities such as calibration, which are carried out primarily during the day.

Having calculated the raw livetime distribution as shown in figure 6.3, the separate day and night livetimes can be found. For the neutrino data set there are 101.0 days of livetime during the day and 144.6 days during the night.

Verification of the differential livetime calculation

As well as being used to verify the total livetime, PulseGT data can also be used to check that the livetime distributions have been calculated in the correct way. In this way the PulseGT data is used to verify the differential livetime code. The results of this process are shown as the ratio of observed PulseGT events to the expected number based upon the livetime. The results of this procedure are shown in figure 6.4.

The figure shows that with the exception of four bins the differential livetime is verified to better than 0.005%. This small difference can be attributed to stolen PulseGT events. For the other four bins the deviation is larger. The biggest is a difference of 0.03%. The difference between this value and the typical value is not due to extra stolen triggers. It may point to a problem with this verification method, where extra PulseGT events have been lost for an unknown reason, or it could point to an error with the raw livetime measurement. As stated before the PulseGT discrepancy is of the same size as the clock discrepancy. A mistake by the 10MHz clock with the start or end time of a run would lead to problems of this kind in the differential livetime calculation, and the discrepancy noted by the 50MHz clock. However these problems are very small and can be added to the systematic error without difficulty.

This would be the end of the livetime calculation were it not for the use of cuts that cause livetime corrections, both at the instrumental level and in the higher level analysis. To obtain a true measure of the livetime these corrections must be taken into account.

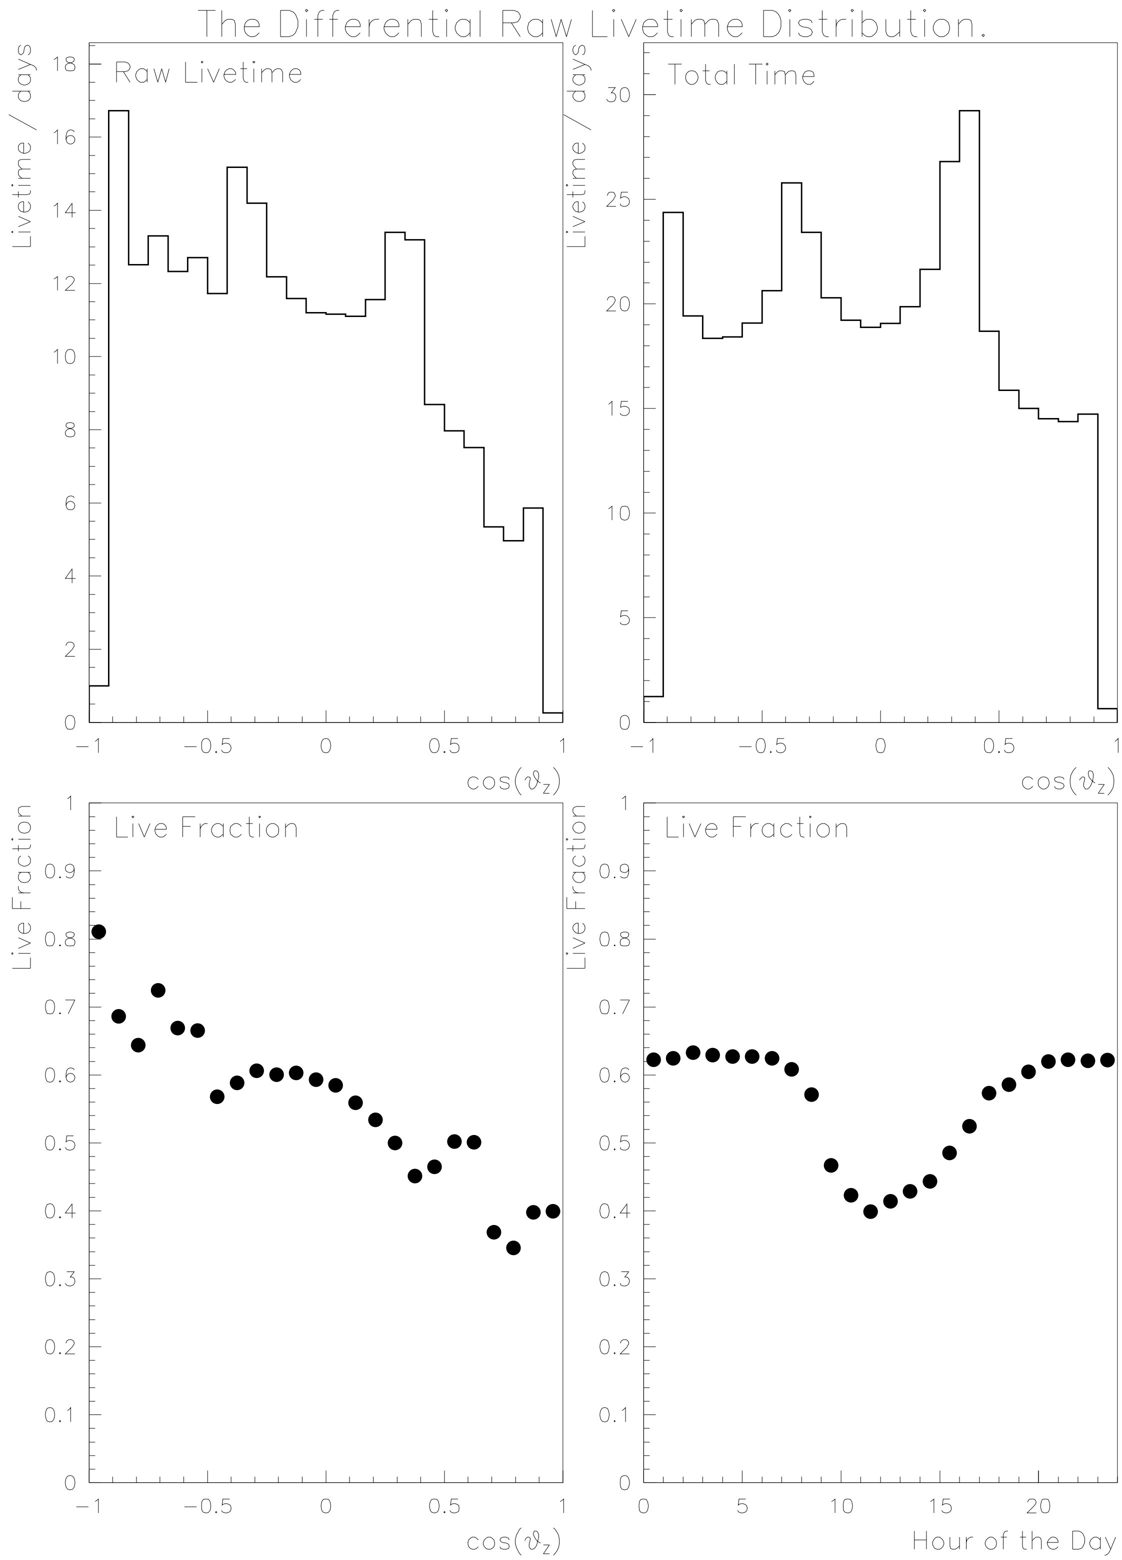

Figure 6.3: The raw livetime distribution of the selected neutrino data set, compared to the time period of data taking. The top left plot shows the distribution of raw livetime. The plot in the top right shows the same distribution of the entire period of data taking. The lower left plot, shows fractional livetime compared to the solar position, and in the lower right the fractional livetime is shown for the time of day.

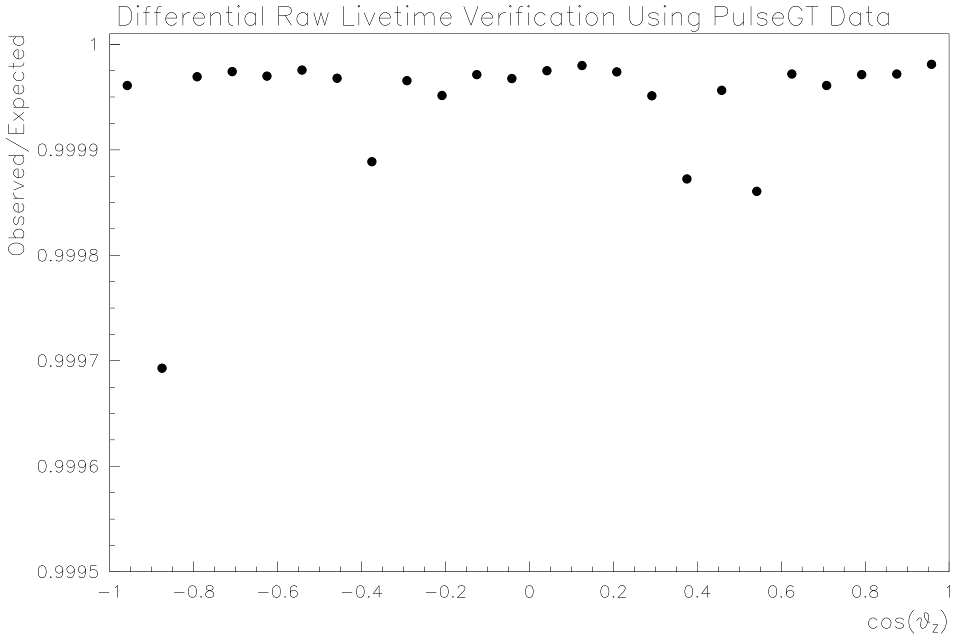

Figure 6.4: Verification of the differential livetime distribution using PulseGT data.

6.2 The Burst Cuts

The burst cuts are designed to remove periods of excessively high rate in the detector. These periods of high rate arise because many of the instrumental backgrounds seen in SNO occur in bursts. The use of burst cuts to remove these events not only provides protection against leakage through the instrumental cuts, but also periods of time during which the detector is unstable. The burst cuts also have the potential to remove any class of physics event that have time correlations, such as anti-neutrino events. The use of burst cuts provided protection against these types of event appearing inside a sample of solar neutrinos, but does mean that a different set of instrumental cuts must be used while attempting an analysis for these signals.

The suite of instrumental cuts contains three burst cuts, which work on very different time-scales. The shortest of these is the retrigger cut. This particular burst cut is specifically designed to remove retriggers from the dataset. For this reason the burst cut operates as a follower cut and not a burst cut in the usual sense. The next cut is the burst cut, with the NHIT burst cut operating over the longest time scales. These two cuts remove genuine bursts from the data.

6.2.1 Burst cuts and randomly distributed data.

The burst cuts work in a different way to the other instrumental cuts. They do not search for particular events, so any event that occurs can trigger a burst cut. This means that the cuts do not have a sacrifice in the way that event based cuts do. Instead the burst cuts cause a livetime correction, or dead time in the data. In order to understand the measurement of the dead time imposed by the burst cuts the effect of them on randomly distributed data needs to be understood.

The first thing that needs to be decided while attempting to understand the dead time is the position of the beginning and end of the time window that is removed by a burst cut. There are a number of choices that can be made.

- Burst Only.

  The time window removed goes from the first event of the burst to the last. Thus, for a burst with N events.

$$t_{cut} = t_N - t_1 \tag{6.1}$$

- Time windows above threshold.

  The time window removed covers all possible time windows where there are more than the threshold number of events. Thus, for a burst of N events where the burst cut removes n or more events in t_{win}.

$$t_{cut} = (t_{N-n} + t_{win}) - (t_n - t_{win}) \tag{6.2}$$

- Time window in which events would be cut.

  This is the largest time window. This will tag events inside a time window during which a trigger event would have been cut, had it occurred. With the cut at this setting, time windows around a burst that are below threshold are cut. Thus once again for a burst of N events, where there the burst cut removes n or more events in t_{win}.

$$t_{cut} = (t_{N-(n-1)} + t_{win}) - (t_{1+(n-1)} - t_{win}) \tag{6.3}$$

The choice made, was to use the time window above threshold approach (equation 6.2). The removal of bursts only does not give an accurate estimate of the time lost due to the removal of the burst. There is a period of time outside the burst where if an event occurred it would be removed. The third option cuts periods of time which are below threshold, this is unsatisfactory, as it blurs the burst threshold. The second option however gives an accurate representation of the amount of data removed by a burst cut, so this is the time window used.

In order to understand the effect of burst cuts on randomly distributed data it is first necessary to understand the δt$_n$ distributions expected for this data. The occurrence of events in SNO in the absence of bursts and events from the pulsed global trigger is a Poissonian process. This means that the statistics for these processes are well understood [81]. The δt$_n$ distribution follows that for random stochastic processes, the probability density function (PDF) for which is shown in equation 6.4, where δt_n is the time difference from an event until the n^{th} follower and R is the average rate of events.

$$p(\delta t_n) = \frac{R^n (\delta t)^{(n-1)} e^{-R\delta t}}{(n-1)!} \tag{6.4}$$

Using this equation it is possible to calculate the sacrifice and the fractional livetime correction imposed by a burst cut. These two numbers are crucial to the understanding of the action of burst cuts on the data. The sacrifice is fraction of events removed by a cut, while the fractional livetime correction is the fraction of time removed by the cut. These two numbers do not have to be the same.

Case study: The two event burst cut

It is an informative process to show how the sacrifice and livetime correction can be calculated for a given burst cut. To do this a burst cut with a threshold of two events shall be used. The burst cut will remove all time periods where two or more events occur inside a time window t_{win}. The assumption that $Rt_{win} \ll 1$ will be made, this will be the case for all burst cuts used in SNO. Having made this assumption the dominant contribution to random failure of the burst cut is due to occasions where two events pile-up within t_{win}. The case of the pile-up of three events within t_{win} can safely be ignored.

Given that only two event bursts are important equation 6.4 can now be used. For a two event burst the time to the first follower δt is the variable of interest. Thus equation 6.4 simplifies to equation 6.5.

$$p(\delta t) = Re^{-R\delta t} \tag{6.5}$$

In the general case $n = (N-1)$, where N is the burst threshold. With equation 6.5 it is then possible to calculate the probability that an event has a follower within t_{win}, by integrating the PDF.

$$p_{\delta t < t_{win}} = 1 - e^{-Rt_{win}} = Rt_{win} + O((Rt_{win})^2) \tag{6.6}$$

This equation is simply related to the sacrifice of the cut. The event can be cut by the burst cut by either having a follower within t_{win}, or following within t_{win}. For this reason the sacrifice of the cut is twice the value of the probability shown in equation 6.6. In the general case, for $Rt_{win} \ll 1$ the sacrifice will be $(n+1)p_{\delta t_n < t_{win}}$. The sacrifice of the cut only tells one part of the story. It is also necessary to know the size of the livetime correction imposed by the cut, in burstless data.

The burst cut will remove a specified amount of time each time it is triggered. This time is shown in equation 6.2. For the case where there are no real bursts in the data this becomes on average $t_{cut} = 2t_{win} - < \delta t >$. This equation holds for any size of burst cut. $< \delta t >$ is calculable in a straight forward way from equation 6.5.

$$< \delta t >= \frac{\int_0^{t_{win}} R\delta t e^{-R\delta t} d\delta t}{\int_0^{t_{win}} Re^{-R\delta t} d\delta t} = \frac{t_{win}}{2} + \frac{O((Rt_{win})^2)}{R} \tag{6.7}$$

Thus for an average accidental burst $t_{cut} = \frac{3t_{win}}{2}$.

In a run of time T there will be RT events. This can be combined with equation 6.6 to calculate the number of bursts in the run. This is to first order R^2Tt_{win}. It is now possible to find the expected fractional dead time:

$$\frac{T_{cut}}{T} = \frac{3(Rt_{win})^2}{2} \tag{6.8}$$

It can be seen that in the case of data without bursts, the sacrifice of the cut is not equal to the fractional dead time measured by the cut. This means by measuring the fractional dead time introduced by the cut, the sacrifice has not been measured. The sacrifice is much greater, as the time removed is by definition, the sum of time periods which preferentially contain events, or in other words the burst cut removes time periods in a biased way. When the data does contain real bursts the fractional dead time will grow accordingly while the sacrifice due to rate fluctuations remains constant. In the case where the dead time due to real bursts is much greater than the dead time due to coincidences then the sacrifice of the cut is measured by the fractional dead time. At this point the sacrifice due to coincidences becomes a correction to the dead time. However some of that correction is already included in the dead time, so the extra sacrifice is quoted as a systematic error on the deadtime.

The procedure laid out in this section can be extend to any length of burst and the expected sacrifice and livetime correction can be calculated. These numbers can then be taken into account when calculating the livetime correction due to the cut.

6.2.2 Implementation of the burst cut

The burst cut runs as part of the random access analyser (RAA) within SNOMAN [41]. The RAA is a subrun based process that works in five stages. It allows a double pass over the data for time dependent analyses. These analyses are carried out by subprocessors of which the burst cut is one. The five stages of the RAA each carry out a different operation within the SNOMAN job.

- Initialisation Phase

  This mode is used to initialise the RAA and any active subprocessors. This means that any constants are read in and counters are initialised.

- Save Phase

  This mode is used to carry out any part of the analysis that requires use of the full event data. This phase is also used to fill the summary record and the event buffer used by the RAA.

- Analyse Phase

  This mode is where the bulk of the power of the RAA lies. It is here that time correlations are studied using the summary record. The summary record is a collection of numbers connected to each event that can be traversed in both directions without the need to reload events. The results of the analysis are stored by the use of bit masks stored inside the summary record.

- Restore Phase.

  This mode is used to restore events from the buffer to allow further processing of them, perhaps dependent on the results of the time correlated analysis. The results from the analyse phase are written to the event data at this point to make them available for outside analysis.

- Termination Phase

  This mode is used to report any processing statistics from the RAA.

The burst cut does the majority of its work inside the analyse phase using the summary record to find and tag events that satisfy one or more of the burst criteria. Because the burst cut can search for events satisfying up to ten different criteria, the cut has its own word in the summary record to store the burst cut results. The results are stored in bits 1-10 of this data word. The remaining bits are used for other purposes.

The algorithm for finding bursts is quite simple. First the events are scanned to see if they satisfy the conditions for the event to trigger the cut. There are currently three conditions tested before further analysis of the event.

- The triggering of the event.

  The burst cut is designed to work on events that are not forced triggers. This means that PulseGT and pedestal events are ignored by the burst cut to avoid these events having an impact on the burst cut. This condition is also used to avoid orphan events from any burst analysis, as these events do not have a valid time.

- The results of previous burst cuts.

  The burst cut runs each cut independently. This means that by specifying the shorter timescale cuts first the result can be known when longer timescale cuts are processed. With this knowledge the events that fail the shorter cuts can be avoided. This is particularly useful when it comes to avoiding retrigger events which from the point of view of the other burst cuts are artificial. By using this setting, events that cause retriggers are not given extra weight by other burst cuts. The burst can be made to avoid events failing any combination of previous burst cuts.

- The NHIT threshold .

  There is also the option of requiring that the event pass a specified NHIT cut. This allows the cuts to look for bursts of events that are above our analysis threshold. This option is easily disabled by setting the threshold to zero.

Once the events have passed initial filtering the remaining events are passed onto the burst finding algorithm. The algorithm has two inputs used to find the bursts in run. These are a threshold number of events and a time window. These two inputs define a threshold rate, above which events are removed. The summary record is searched until a period of time is found that satisfies the requirement for a particular burst. The burst cut will then expand the time window around the burst until the entire burst is contained within. The start and end times of the cut window are set as shown in equation 6.2. All events within the time window whether or not they pass the initial cuts are then marked as failing the burst cut. The one exception to this is if the first event switch is set. This switch turns the burst cut into a follower cut. If this is set the first event of the burst is not tagged by the burst bit, but is instead tagged by the first event bit, bit 0 of the burst cut status word. As knowledge of the removed data is required when calculating the dead time for a suite of cuts, the events defining the time window (t_{N-n} and t_n) are tagged in the burst cut status word. This allows the analysis of the cut times to proceed without the need to reapply the cuts.

The RAA works on at a subrun[3] level. This means that effects at subrun boundaries can be missed as conditions outside the subrun are unknown. For this reason the burst cut will also tag events within 1 minute of the beginning and end of each subrun. This tag provides a check against extra events that could be missed by the cuts that run within the RAA. The time window is chosen at 1 minute as this is the length of the muon follower long tag, the longest in a cut that uses the RAA. Subrun boundary events are tagged by bit 31 of the burst cut status word.

The burst cut collects statistics to report the number of events failing each cut, and the dead time imposed by each. These statistics are reported at the end of the job, and the dead time is also passed the the runlog analysis system, to be stored in the RLDT bank.

6.2.3 The Retrigger cut

After any high charge or NHIT event in the SNO detector there are always a number of events that follow. Usually these events occur 500ns later, and are characterised by a flat time spectrum. The burst cut is the best way to be certain of cutting these events. This does present a problem, however, in that the burst cut identifies and marks a burst of data. The event that caused the retriggers may have nothing wrong with it, and indeed cutting these events would introduce unnecessary NHIT bias to the sacrifice. For this reason the retrigger cut tags the first event of the burst independently, and this is not removed. The event is instead tagged by the first event tag. With this implemented the retrigger cut operates as an event follower cut, not a true burst cut.

As with all burst cuts, the retrigger cut has the two basic parameters, the time window and the event threshold. As events may have only one retrigger, the event threshold is fixed at two events. The next step is to determine the time window needed to cut the retriggers. For this application the time window should be as long as needed so as to cut all the retriggers, because although most of these events occur at 500ns this is not guaranteed. In order to do this a data set must first be constructed. A number of neutrino runs were selected, these are shown in table 6.1. The data from these runs is then pre-processed, by removal of PulseGT and orphan events. Following this step the δt distribution is found, and is shown in fig 6.5. The data is as expected an exponential, with the exception that the first bin is very highly populated. These are the event retriggers. By examining the distribution below $10\mu s$, the time window of the cut can be determined.

This figure shows that looking below $10\mu s$ most of the events occur within the first microsecond, but a significant fraction occur later than this but within $5\mu s$. Indeed there

[3] A subrun is a single data file. A run can consist of up to 1000 subruns, however the largest run in the data set consists of 13 subruns.

Run	Run length/h	Run	Run length/h
10040	35	10133	22
11269	18	11275	36
11291	19	11303	17
11310	20	11312	29
11313	13	11363	16
11366	26	11368	21
11371	60	11384	29
11389	16	11393	24
11399	96	Total	497

Table 6.1: The runs used to determine burst cut parameters.

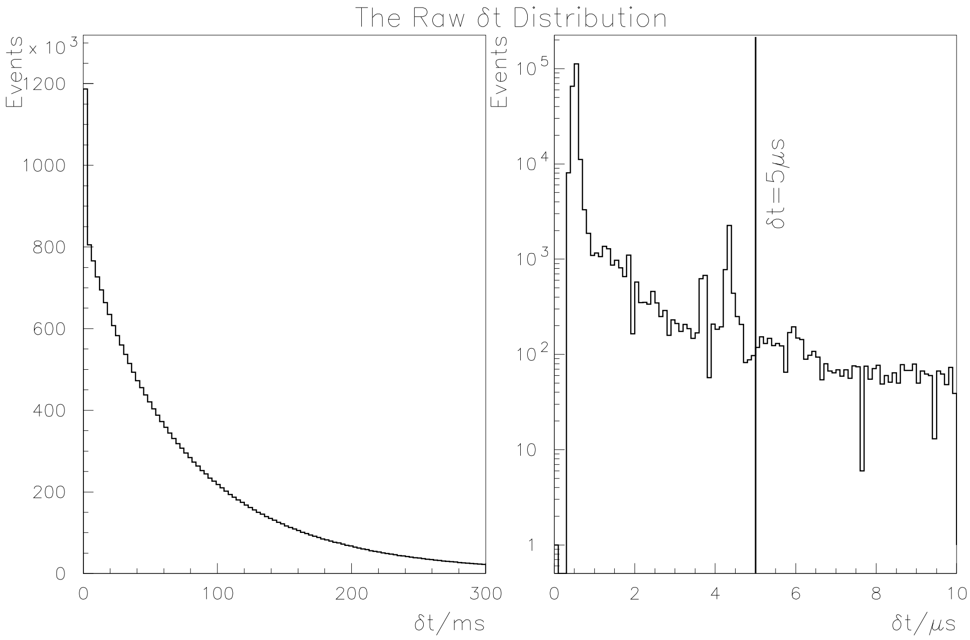

Figure 6.5: The raw δt distribution of events.

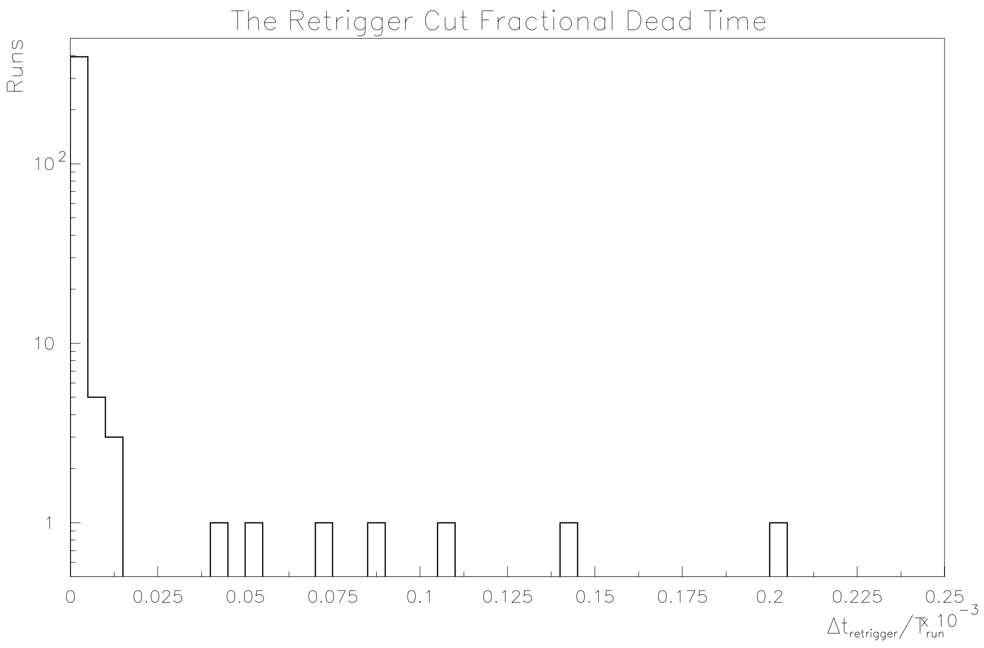

Figure 6.6: The fractional dead time from the retrigger cut.

is a second peak of retrigger events occurring with a time gap of $4.3\mu s$. These events look like PulseGT, but have NHIT20 triggers associated with them and are due to pickup on the NHIT20 trigger. However they still typically follow large events just like ordinary retriggers, so they should still be cut in the retrigger cut. The third peak at $3.8\mu s$ is due the combination of both retrigger classes in the same burst.

After this the events fall off to a constant background rate. Thus the retrigger cut is placed as two events within $5\mu s$. Placing this cut here removes all the peaks shown in fig 6.5 with minimal effect on other events. Events after this are likely to be cut by subsequent burst cuts but are not considered to be retriggers.

The retrigger is a burst cut despite its mode as a follower cut. This means that it does introduce a certain amount of dead time into the data. The fractional dead time introduced per run is shown in figure 6.6. The plot shows that the dead time is dominated by a small number of runs containing large bursts. Because of runs like these it is important to find the total dead time introduced over the entire data set, rather than consider some mean correction. The total amount of time removed by this cut is 13.3s, or 6.3×10^{-5}%. This is a very small amount of time. However as has already been discussed the sacrifice and the fractional deadtime are not the same for burst cuts.

The detector takes data at an average rate of 16Hz, 5Hz of which is PulseGT data. The retrigger cut is an event follower cut, so the sacrifice of the cut is simply Rt_{win}. For this cut this gives a sacrifice of $5.5 \times 10^{-3}\%$. The theoretical deadtime due to coincidences is much smaller than this at, $3 \times 10^{-8}\%$. This shows that the measured deadtime is overwhelmingly due to bursts, but the sacrifice is dominated by coincidences, but the level of this is very small. The sacrifice will be included in the systematic error on the livetime.

6.2.4 Genuine bursts of data

There are many ways the SNO detector can see bursts of data other than retriggering. These include after a seismic event, or spallation after a muon has passed through the detector. In particular large bursts of flashers and neck events are a common occurrence. Burst cuts are again an ideal way to remove these events. There are two different approaches used to try and tag these bursts. The first approach is to look at all the events and to find a suitable time scale and event threshold. The second approach is to apply a cut to reduce the data sample and then look for a burst. The cut is applied so that the burst cut works in the solar neutrino regime.

In both cases the cuts are designed to look for bursts of data and not following events. This means that the first events of any burst is given no special status and is removed accordingly. Also to avoid preferential removal of events that cause retriggers, retriggers are ignored by the cuts. This means that the cuts can cope with the artificially high rate caused by the retriggers.

Cut with the entire data set - The Burst cut

To determine the parameters needed here the time difference between one event and the nth following event, δt_n is examined. This is expected to have a peak at small times that indicates where the cut should be placed. The threshold number of events inside the time window should be more than two, as the two events burst size is used in the retrigger cut. However a small burst size (and therefore time window) is desired, to be complementary to the cut in the solar neutrino regime. So an event threshold of three was chosen and the time window determined.

To achieve this the data shown in table 6.1 was further pre-processed to remove retrigger events. The resulting δt_2 distribution is shown in figure 6.7. In this case the peak at small times showing the bursts is not visible on the plot showing the entire distribution. However by examination of the distribution below 4ms the peak containing the bursts of events is clearly visible. The choice of this cut is made to cut the obvious peak. At larger times the tail of the peak is obscured by the background rate of events, making the exact placement of this cut difficult. The choice of 1ms was made to attempt to maximise the removal of this peak, without introducing unnecessary sacrifice. With

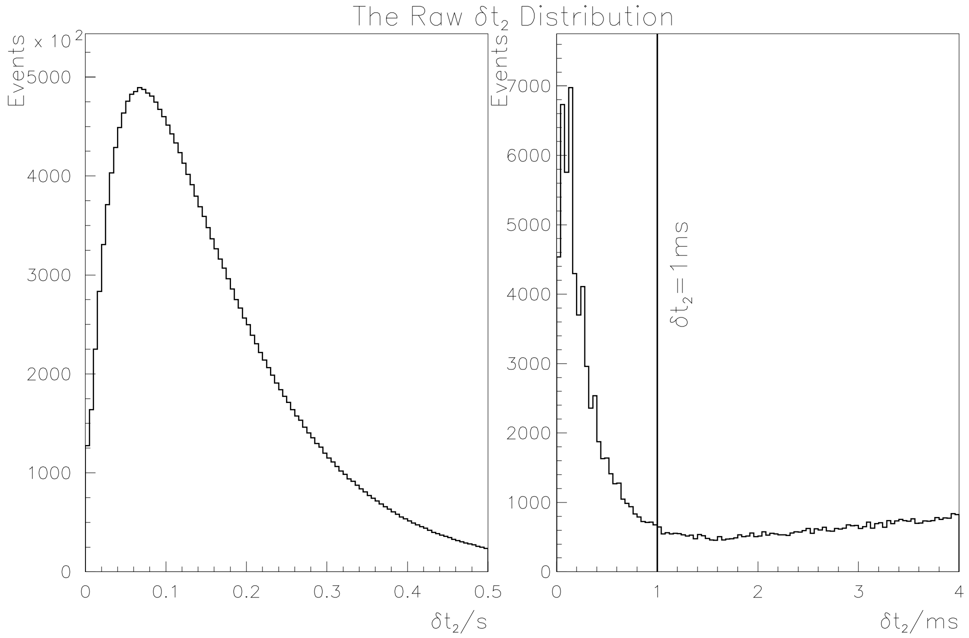

Figure 6.7: The raw δt_2 distribution, with retriggers removed.

hindsight it may be preferable to cut at 1.5ms, however at time of writing the cut values are set in stone. The burst cut therefore searches for three events within 1ms.

The distribution of fractional dead time introduced by the cut is shown in figure 6.8. As was the case for the retrigger cut the total dead time is dominated by the few outliers, where large bursts have occurred. The total dead time introduced by this cut is 15.8 minutes, or $4.5 \times 10^{-3}\%$.

Following the procedure laid out in section 6.2.1 the fractional dead time and sacrifice can be calculated. It is found that the theoretical fractional dead time is $9 \times 10^{-5}\%$, much less than the measured dead time. This statement shows that the data at these timescales is dominated by real bursts and not coincidences, thus the fractional deadtime does give the correct measure of sacrifice due to the bursts. The theoretical sacrifice due to coincidences is 0.02%. This will be added as a one sided systematic error to the livetime measurement as some of this will be included in the measured fractional dead time. As was the case for the retrigger cut the data is dominated by bursts, but the sacrifice due to coincidences between events is still larger than the amount of time removed due to those bursts.

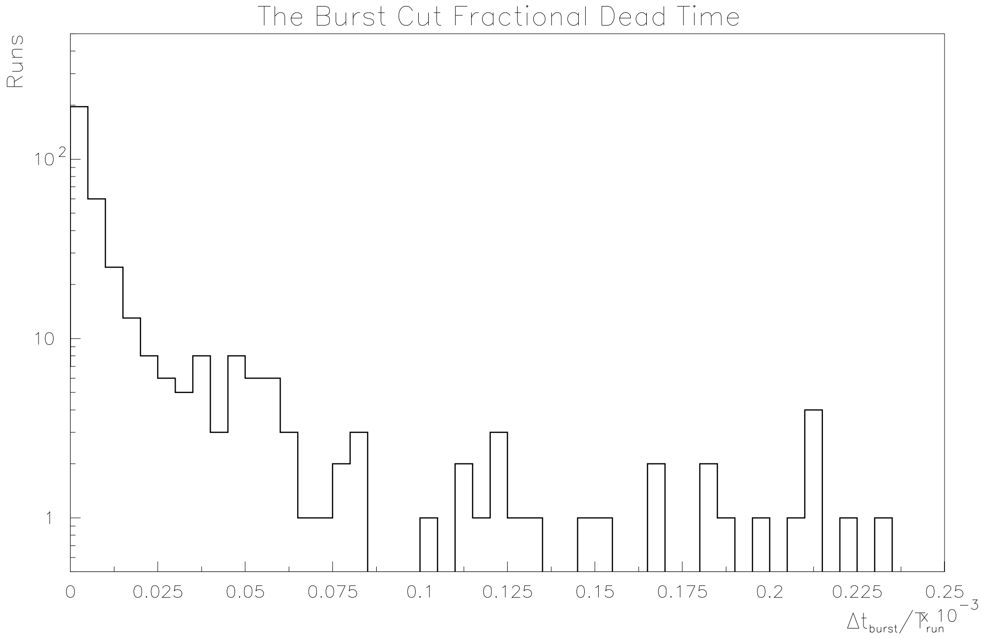

Figure 6.8: The fractional dead time per run from the burst cut.

Cut with a reduced data set - The NHIT burst cut

The second approach to genuine bursts of data is to look at a subset of that data. By far the simplest way to do that is to apply an NHIT cut. This cut can then work over a much longer time-scale. This cut is explicitly designed to look for flasher or neck bursts, so a suitable NHIT cut can be placed at forty hits. The cut then works on events that fall into the analysis regime for solar neutrinos. This also means the cut will also remove physics events that occur in bursts such as neutrons following a muon, or a supernova if one occurs.

Once again the same procedure as before was used to calculate the needed time window for this cut. Due to the nature of the cut a larger number of events will also be needed inside that window to identify the burst. This number must be large enough to begin to separate non-burst events from the small time peak, but small enough so that sensitivity to smaller bursts is not lost. This is particularly important as neck events often appear in small bursts. Taking these bursts into account the threshold was chosen to be six events.

Figure 6.9 shows the time structure for events with more than forty hits. The bursts in the data are very clearly visible in the large number of events in the first bin. Examination of the distribution below 20s shows that there are two contributions to the

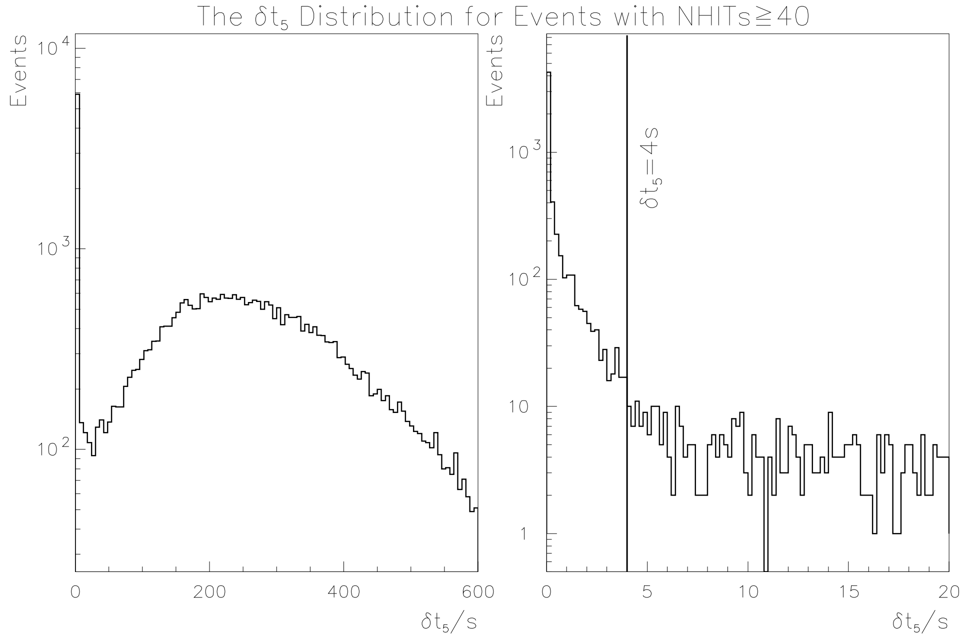

Figure 6.9: The raw δt_5 distribution, with retriggers removed and NHITs$\geq$ 40.

number of events in this bin. The first contribution is visible in the still highly populated first bin. The second contribution is occurring over a longer time scale and is visible as the rise in event numbers towards small values of δt_5. Given the time scale of bursts of flashers and neck events the contributions towards the excess at small δt_5 are entirely expected. The timescale in a flasher burst is of the order of ms, while that in a neck burst is two orders of magnitude greater.

Because of the time scale of the cut the amount of introduced dead time is a more important consideration than for the Burst cut. For this reason the threshold is set at 4s, removing the bulk of the burst data, but attempting to keep the sacrifice of the cut at an acceptable level. As the time window is increased above this, the NHIT burst cut fails the sacrifice criteria required for instrumental cuts as the fractional deadtime grows.

The distribution of fractional dead times for the NHIT burst cut, shown in figure 6.10, shows that for this cut the sacrifice introduced is much more of a problem. This cut removes 8.4 hours, or 0.14% of data, on a par with the AMB and QvT cuts for sacrifice. As was the case for both the burst and retrigger cuts there are again a number of runs where the dead time is much larger due to large bursts during the run.

Once again we can also calculate the theoretical sacrifice and fractional dead time of the cut. The rate used in this calculation is different to that for the retrigger and burst cuts. For the NHIT burst cut the rate of trigger events is 0.02Hz. At this rate the NHIT

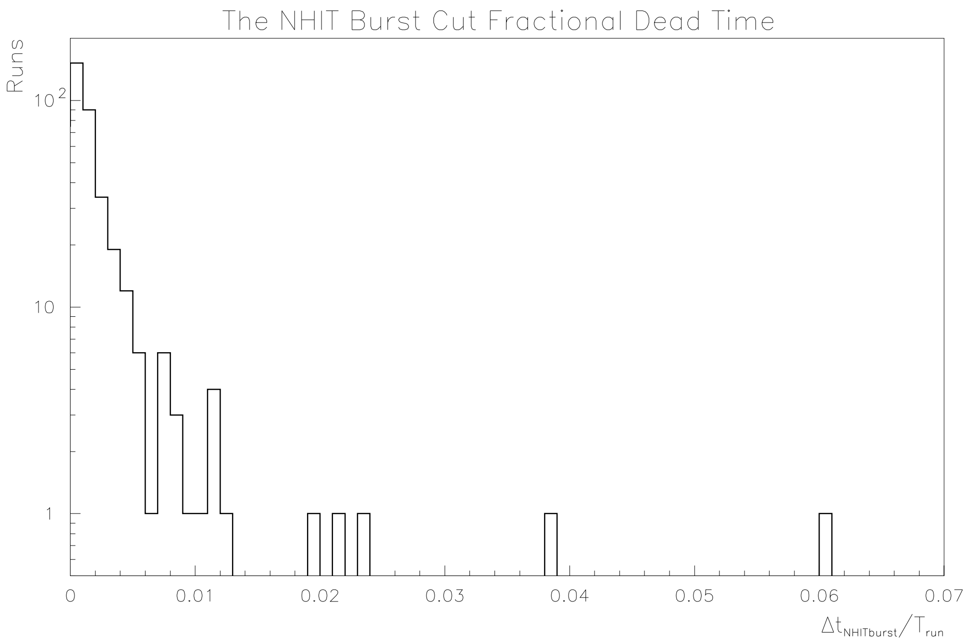

Figure 6.10: The fractional dead time per run from the NHIT burst cut.

burst cut will only cut a tiny amount of data due to coincidental bursts. However the sacrifice does depend upon the fifth power of the rate, so a small change in the rate of events above forty hits can lead to a much bigger change in the theoretical sacrifice of the cut. The theoretical sacrifice is $2 \times 10^{-5}\%$, much less than the measured fractional dead time, and small enough that large changes in the rate of trigger events are required before it becomes important. Therefore for the burst cut with the largest measured fractional deadtime we find that, unlike the retrigger and burst cases, the fractional deadtime really does measure the sacrifice of the cut.

6.3 The muon follower cut

At the depth of over 2km underground the through-going muon flux in the detector due to cosmic rays is vastly reduced. However the flux is not reduced to zero and there are approximately three through-going muons per hour in SNO. Each of these muons has the potential to cause spallation inside the detector. The spallation products can then create a background to the solar neutrino signal. There are two classes of spallation products seen in the SNO detector.

The first and most important of these are neutrons created by the muon. These can occur in very high multiplicities, with bursts of several hundred detected neutrons not unknown. The time scale for these bursts is of the order of the neutron capture time in the detector, approximately 50ms in pure D_2O. The high multiplicity events are removed by the various burst cuts but these cuts are not aggressive enough to remove events when the muon only creates a small number of neutrons. Potentially bursts of up to 4 neutrons could pass the burst cuts. These neutrons provide a very important potential background to neutral current measurements, however due to the ease of tagging of the neutrons, they provide a very nice way to contaminate the data in a controlled, but unknown fashion to carry out a blind analysis of the neutral current signal.

Spallation products are also created by breakup of the ^{16}O nucleus by the muon. These spallation products will decay and create backgrounds to the neutrino signals. The number of events caused by this process is expected to be small (10 to 100 events per year [82]), this number of events still would cause a significant potential background. Some of the products created in this process have relatively long lifetimes, and these lifetimes are used to set the size of the follower cut. In particular the ^{16}O breakup can create ^{16}N[4], which has a half-life of 7.13s. With these decays in mind the muon follower cut is set at any event within 20s of an identified muon. This is almost three half-lives of this particular decay product, removing approximately 90% of these decays. Studies to look for these spallation events[83] find no evidence for their appearance, so the 20s window is sufficient.

However, in order to allow systematic checks there are two time windows of tagged events after each identified muon, known as muon follower short and long. Muon follower short is set at the 20s window discussed and is used as the muon follower cut. Muon follower long tags 1 minute after each muon, and is used to look for any differences in this time window that may indicate that the window has been set incorrectly.

6.3.1 Muon identification

Muon identification is carried out as described in [82]. The muon tag is specifically designed to identify through-going muons, although it does have some power to remove partially contained events as well. Fully contained events are not tagged. Followers from these events must be removed by the burst cuts. However since the rate of fully contained events is extremely small, in the presence of burst cuts followers from these events do not cause problems. There are a number of simple cuts carried out to identify the muon.

- The NHIT threshold

  Through-going muons are high energy events, and even in the case where they just graze the PSUP they cause a large amount of light to be emitted in the detector.

[4] ^{16}N is actually created by muon capture on ^{16}O

The choice of threshold is made in order to remove as much instrumental background as possible, while keeping all potential muons. The threshold is therefore set at 150 hits.

- The NOWL threshold

  As a tag designed to identify through-going muons, the muon tag uses the OWL tubes extensively. The through-going muon will create large amounts of light in the outer volume of H_2O, where it can be detected by the OWL tubes. The use of these tubes provides very good discrimination against many backgrounds that occur inside the PSUP. External backgrounds such as neck events are not however effectively removed by this cut. The OWL threshold is set to 5 hits on the OWL tubes.

- OWL retrigger lockout

  Retrigger events will often cause readout in the OWL tubes due to pickup, or late light in the case where light is emitted in the outer H_2O volume. These events cause a background to the muon tag, and are removed by applying a retrigger lockout of 5μs after an event above NOWL threshold.

- The Neck cut

  The muon tag uses the neck cut. This is discussed extensively in chapter 4.

- The time RMS cut

  Many background events to muon identification, contain light emitted over a long period of time. This causes the RMS of the PMT timing to have very large values. This can be used to reduce the contamination of events selected by the muon tag. To achieve this end a cut is placed at 90ns.

The muon tag is used in this way to identify muons and allow the followers to be tagged. The tag is designed to be as close as possible to 100% efficient. This means that the false tag rate is relatively high. However since the sacrifice of the muon follower cut is easy to measure and understand the false tag rate does not cause a problem from a sacrifice point of view unless the rate becomes unacceptably high.

6.3.2 Implementation of the cut

The muon follower cut, like the burst cut works within the RAA. This avoids any possible problems with event disordering within the data. The muon identification is carried out during the save phase of the RAA process, and muon events are tagged using the RAA

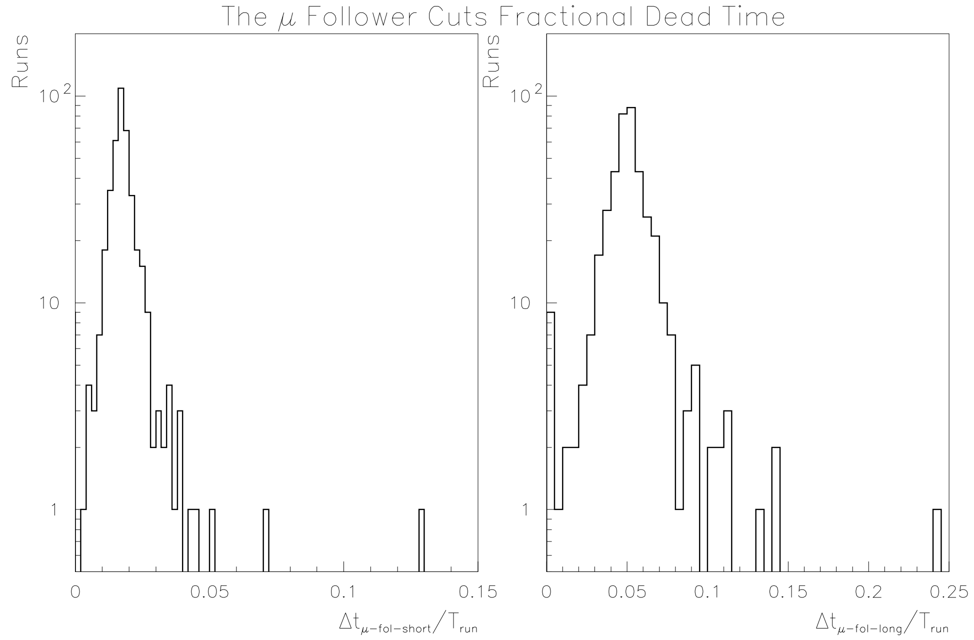

Figure 6.11: The fractional dead time distributions from the muon follower short and long cuts.

status word. The followers are identified during the analysis phase using the summary record. Two further bits are used in the RAA status word to record the results of the follower cuts.

The dead time introduced by the muon follower cuts is also calculated during the analysis phase. While this is a much simpler process than that for the burst cuts, as neutrinos will not trigger the cut, there are still problems that must be overcome. The dead time from the cut is simply twenty seconds (or one minute) for each event identified by the muon tag[5]. There are however two corrections that must be made to this number. The first is a correction for the end of a subrun. The cut cannot extend over a subrun boundary, so this must be accounted for. The second correction is for overlap between time windows. While muon events are unlikely to occur in such close proximity, false tags can and do occur in bursts.

The fractional dead time distributions from the two time windows is shown in figure 6.11. These plots show that the dead time for the muon follower cuts is distributed in a much more Gaussian way than for the burst cuts. However as for the burst cuts the total dead time is still the important number. The muon follower short cut removes 102 hours of data from the data set. This corresponds to a sacrifice of 1.73%, which is much

[5]This is because the time windows are unbiased from the point of view of solar neutrinos.

bigger than any of the instrumental cuts, but is the best understood sacrifice of any cut. The amount of dead time introduced by the cut is very close to that expected from the raw muon rate. This tells us that the problem of mis-tagging in the muon tag is not a big problem, primarily because mis-tagged events occur in bursts. The muon follower long cut removes 298 hours, or 5.05% of data, providing a sample of 196 hours for systematic checks. As is the case for the burst cuts the dead time from the muon follower cuts is stored in the RLDT bank. The results of the muon follower cut are stored in the DAMN 0 banks alongside with the results of the instrumental cuts.

6.3.3 Muon followers and blind analysis

Because muon followers in SNO are predominantly neutrons these events can be used to conduct a blind analysis of neutral current data. The neutrons produced by the muons are created with the same spatial distribution as neutrons from the neutral current interaction. The temporal and energy distributions are different, but these events still provide a very similar signal to neutral current events. Thus by allowing neutrons from muons into an analysis data set in a unknown amount, a blind analysis of the neutral current signal can be carried out. This is an approach taken by SNO for analysis of data taken during the salt phase.

In order to achieve this the muon follower cut has been modified for salt data. Events falling inside a specified time window following the muon are no longer tagged by the muon follower short, or long cuts. They are instead tagged by a blindness tag. The analysis of data can therefore continue in the same way as D_2O analysis, with the application of the muon follower short cut. However this data set now also contains neutrons from muon. Once the analysis procedure is complete these can be removed using the tag to find the true signal.

6.4 Dead time and orphans

The measurement of livetime can also be effected by dead time caused not by cuts applied to the data, but by the detector itself. The dead time introduced in this way is due to the loss in efficiency due to very high data rates.

When there is a very large burst of events (usually due to instrumental backgrounds) the number of PMT bundles sent to the builder can overwhelm the program. Typically in normal operation the builder will hold events in memory for a sufficient amount of time, for the entire event to be received and built. However during periods of very high rate, the builder memory can become full, causing the events to be flushed from the buffer early. In the extreme case events headers can loose all their associated PMT bundles, and are output as orphans. When this happens the detector becomes inefficient, and these periods of time can be classed as dead time.

An analysis of the amount of dead time introduced in this way is presented in [84]. This analysis is based upon the proportion of the PMT bundles appearing as orphans as function of time. The analysis shows that 3.078 hours of the dataset are lost in this way. The estimate presented is an overestimate of the total time lost. This is in part because the method requires coarse time bins, and the binning technique results in over estimation of the dead time. Once the livetime corrections due to cuts have been applied this estimate becomes a larger overestimate, as some of this lost time is removed by the various burst cuts.

Due to the overestimation of the of the dead time, it is accounted for by subtraction of half the dead time, from the livetime. The other half is added to the systematic error. Thus the correction applied to the livetime is $(-0.026 \pm 0.026\%)$. There is no information about the distribution of this through the data. It will therefore be assumed that these occurrences are distributed evenly throughout the data set.

6.5 Livetime corrections

The use of the burst cuts and the muon follower short cut to obtain our standard neutrino data set means that the raw livetime numbers discussed in section 6.1 are no longer valid. Instead the raw measurements must be corrected for the livetime corrections imposed by the cuts. The correction must be carried out in both a total and differential way.

6.5.1 Correlations between corrections

There is an additional complication involved with the application of four different corrections to the data. The cuts that cause these corrections can be correlated with one another. This is particularly important as the NHIT burst and muon follower cuts, which impose the two largest corrections are correlated by design. The correlation between these cuts is twofold. Firstly as discussed bursts of neutrons are known to follow some muons. These bursts of neutrons are one of the targets of the NHIT burst cut giving an overlap between the corrections imposed by the cuts. Further overlap between the cuts arises from mis-tagged muons. The mis-tagged events are typically instrumental backgrounds, which can occur in bursts. There is overlap between each of the different cuts, meaning that the livetime corrections must be corrected for these overlaps.

This process is carried out within the RAA, in the same process that creates the RLDT banks. During the operation of the burst cuts, the events which can be used to find the start and end of each burst have been tagged inside the burst cut status word.

Cut	Correction	Fractional Correction
Retrigger	13.3 s	6.3×10^{-7}
Burst	15.8 mins	4.5×10^{-5}
NHIT Burst	8.4 hours	0.0014
Muon Follower Short	102.1 hours	0.0173
Combined Correction	109.7 hours	0.0186

Table 6.2: The livetime correction imposed by the various cuts, together with the combined correction.

In a similar way the muon events have also been tagged. These tags can then be used to create a table containing the start and end times of the cut periods. This table is then scanned and overlapping periods are combined together until the table contains a list of non-overlapping time periods within which each instant has been removed by one or more of the four cuts.

Once this procedure has been carried out the the total combined livetime correction can be calculated. The result for the entire neutrino data set is shown in table 6.2. The overlap between the cuts is actually much smaller than was expected, just 0.01% of the total livetime, or 9% of the NHIT burst correction. This correction is then applied to the raw livetime to give the true livetime of the data set. Thus the neutrino data analysed contains 241.0 days of data.

6.5.2 Differential Livetime Corrections

The livetime correction is not expected to be constant as a function of time. This is particularly true as the data taking progresses, as the detector is understood in a more complete way and the backgrounds (particularly rings of fire) are reduced. However the corrections are not expected to be constant throughout the day either. Bursts would be expected to be higher during the day due to low level lab activity and D_2O circulation. There is also the phenomenon of the burst of flashers that occur due to blasting in the mine, which typically occurs in the early hours of the morning.

The procedure used to calculate the total correction can also be used to calculate the livetime correction in a differential way. The table created containing the periods of time removed by the cut can be scanned as a function of any desired variable and the correction can be binned just as was carried out for the livetime. This procedure is carried out binning the data as a function of the position of the sun and of the hour of the day. The results of these scans are stored in the RLDT banks.

The fractional livetime correction for the position of the sun, time of day, and for each month of data taking are shown in figure 6.12. As the livetime correction is dominated by the muon follower short cut, which is in turn dominated by the through going muon rate, there is no large scale deviation from the mean. However there are some systematic fluctuations seen in the data. These deviations are caused by variations in the NHIT burst correction.

These distributions show that the day tends to be slightly noisier than the night. It is also seen that the detector has become much quieter throughout the data taking. With the exception of May 2000 the noisiest data occurs in 1999. This shows the decrease in rate of background events as the detector became better understood and we were able to stop the ring of fire events within the detector. May 2000 is a month of exceptionally low livetime due to a number of problems at the time, most particularly with the chiller system. It is therefore of no surprise to discover that the number of bursts is higher during this period, due to increased pickup and the reduced livetime. The changing rate of bursts throughout the day and over time has an impact on the livetime corrections required in the study of the day-night effect. However given the results shown in figure 6.12 these effects can be accounted for. It is found that the correction during the day is 1.94% compared to 1.81% during the night. These correction can be applied to the raw night and day livetimes to give the true livetimes of these two data sets. There are 99.05 days of day data and 141.99 days of night data.

As has been done for the distribution of the raw livetime, the corrected livetime distribution can be checked using PulseGT data. PulseGT events, while not triggering any of the cuts, will be tagged by the cuts if they lie inside the time window removed. By counting events that pass these cuts the corrected livetime distribution can be verified. The repeat of the earlier analysis is shown in figure 6.13.

The agreement between the observed and expected number of events get worse as the livetime corrections are applied. There are now a number of bins where the number of PulseGT events is less than expected. The reason for this change is unknown. The scale of the effect is that the livetime correction measured by the PulseGT data differs by 1% in a few bins. However as the correction is itself of order 1%, the effect of this difference on the total corrected livetime is much smaller, and is added to the total systematic uncertainty on the livetime measurement. The biggest deviation now seen in a bin is now 0.04%. This sets the the scale of the systematic error on the understanding of the differential corrected livetime.

6.6 Conclusions

The results of the determination of the livetime are shown in table 6.3. The raw livetime is measured using the start and stop time of each run in the dataset measured by the

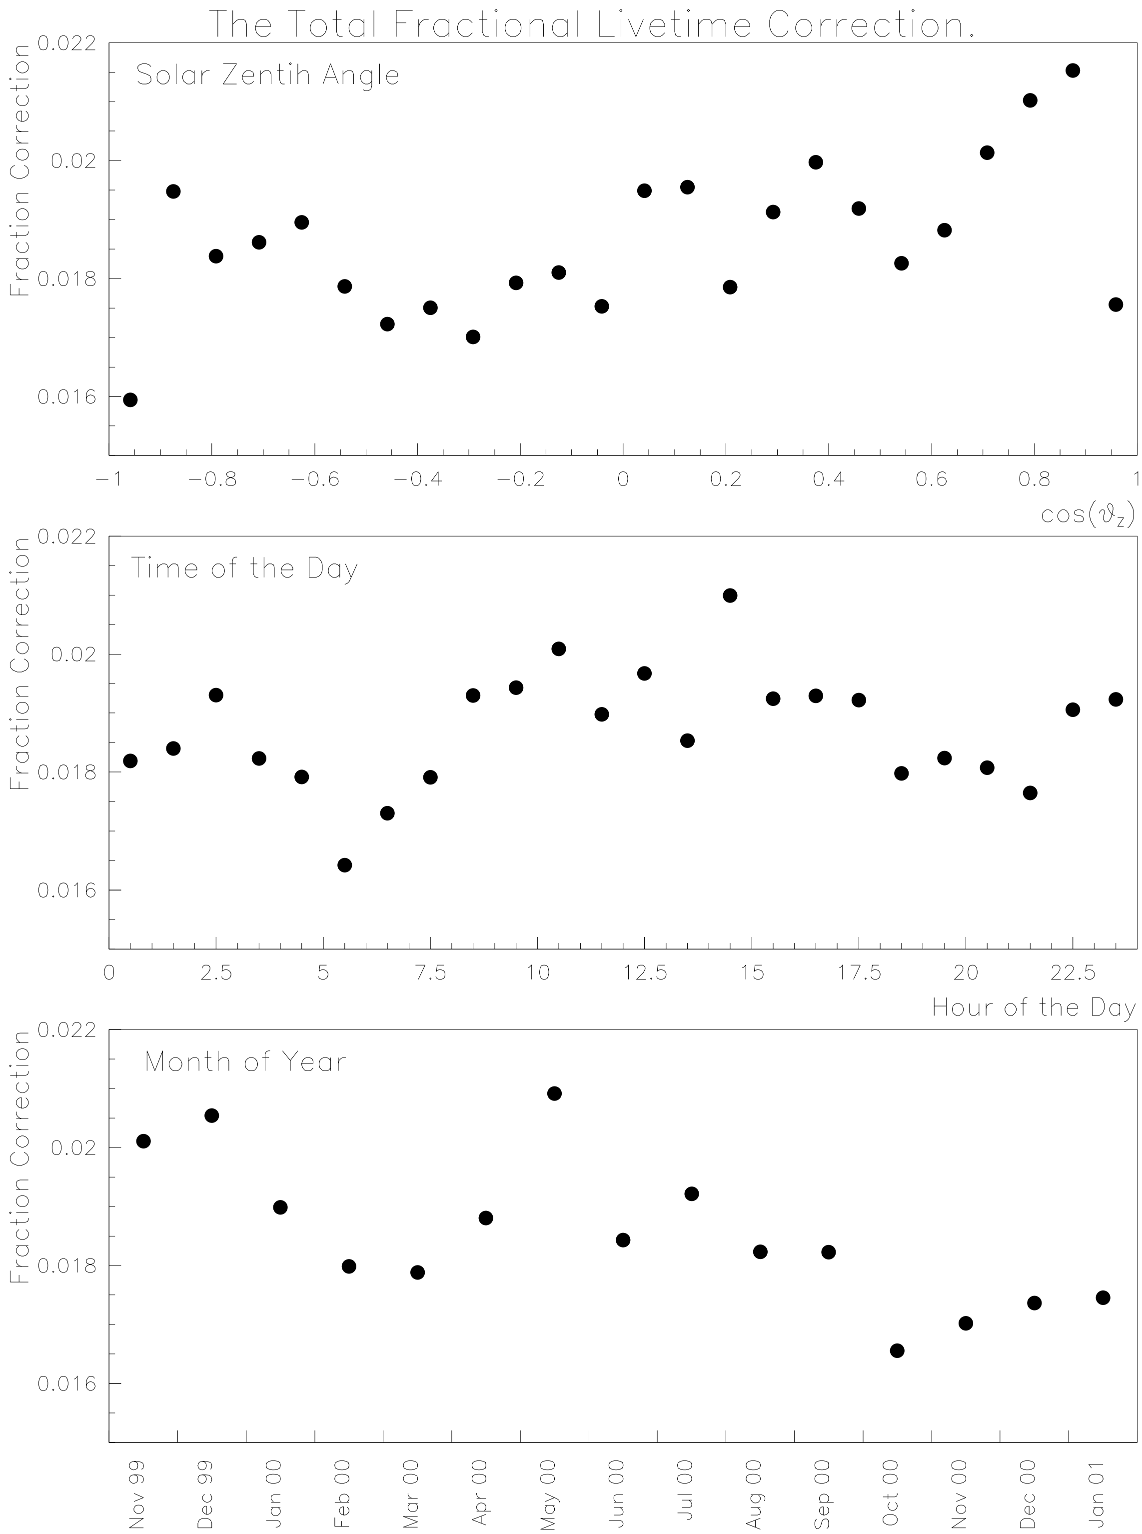

Figure 6.12: The distribution of the fractional livetime corrections as a function of the position of the sun, time of day and month of data taking.

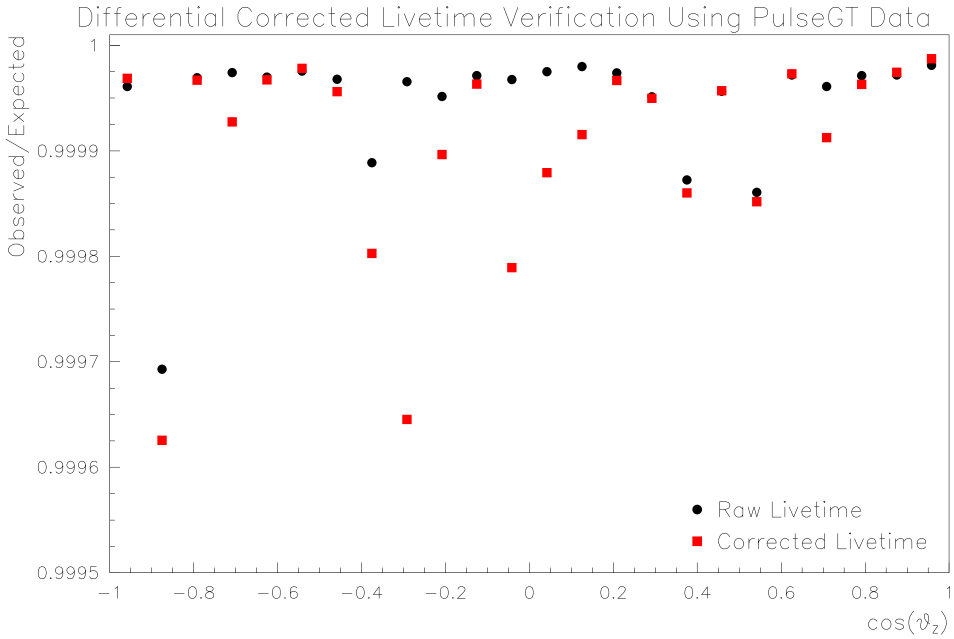

Figure 6.13: The ratio of observed and expected PulseGT as a function of solar zenith angle for the raw and corrected livetime.

Cut	Day	Night
Raw Livetime	101.0 days	144.6 days
Livetime Correction	47.0 hours	62.7 hours
Corrected Livetime	99.1 days	142.0 days
Non-Blind Data	67.3 days	97.9 days
Blind Data	31.7 days	44.1 days
Day-Night Open Data	64.6 days	93.1 days

Table 6.3: Summary of livetime results

10MHz clock. The raw livetime measurement is modified by the presence of cuts that impose livetime corrections. The burst and the muon follower cuts combined impose a correction of 1.8%. There is the further complication of the livetime correction, imposed by the efficiency loss due to bursts of orphans.

The livetime corrections are responsible for the majority of the error in the understanding of the livetime. There are errors due to both the measurement of the corrections, and theoretical errors due to the use of the burst cuts. These errors can be combined in quadrature to give an error in the livetime measurements of $^{+0.032}_{-0.038}\%$. These errors are very small showing that the livetime is one of the better understood properties of the neutrino data set. These errors are very similar in size to those measured for the differential livetime using the PulseGT data, suggesting that these errors are also applicable for those distributions. However to be conservative an error of $\pm 0.04\%$ can be quoted for each bin in the livetime distribution.

The determination of the differential livetime forms an important part of the analysis to determine the day-night asymmetry. The measurement of this quantity for the charge current flux will presented in the next chapter.

Chapter 7

Measuring the Day-Night Asymmetry

As described in chapter 2 observation of a day-night asymmetry in the solar neutrino signal would provide very strong evidence for MSW neutrino oscillations. Measurement of the asymmetry can also strongly constrain the regions of neutrino oscillation parameter space for solar neutrinos.

7.1 The day-night asymmetry

In order to obtain a measurement of the solar neutrino flux, as presented in [58], a set of tools were developed to separate the different signal contributions. This process is known as signal extraction. In the D_2O phase, SNO measures three signals at once via Čerenkov light. These are the charged current (CC) and elastic scattering (ES) signals, and a contribution from neutrons. The neutrons come from two principal sources, the neutral current (NC) interaction of solar neutrinos and photo-disintegration due to γ-rays from the decay products of uranium and thorium in the D_2O. While with an effective kinetic energy threshold at 6.75MeV the neutron contribution to the signal is small these events must be included in the signal extraction as otherwise they would still provide a significant background. Signal extraction uses a maximum likelihood technique to measure the contribution of the three signals based upon the distribution of events in energy, reconstructed radius and direction with respect to the sun. The systematic errors in the measured fluxes come from the various inputs into this procedure, such as energy scale and resolution, or the properties of the event reconstruction procedure. Other effects, such as those due to instrumental backgrounds do not impact upon the signal extraction procedure directly, but can still introduce a systematic.

The extension of this technique to the measurement of the day-night asymmetry is straightforward. The data is split into two (or more) data sets depending upon the position of the sun, and the signal extraction procedure is carried out. When splitting the data into smaller data-sets, the livetime of each has to be measured. This is carried out as described in chapter 6. The day-night asymmetry is expressed as shown in equation 7.1

$$A_{ND} = \frac{\phi_N - \phi_D}{\frac{\phi_N + \phi_D}{2}} \tag{7.1}$$

where $\phi_{N/D}$ is the measured flux during the night/day. With A_{ND} defined in this way the asymmetry can vary between ± 2. The error on the asymmetry is given by the equation:

$$\sigma_{A_{ND}} = \frac{4}{(\phi_N + \phi_D)^2}\sqrt{((\phi_D \sigma_{\phi_N})^2 + (\phi_N \sigma_{\phi_D})^2)} \tag{7.2}$$

There is one principal problem with the A_{ND} parameter, for large values of the asymmetry, it is not distributed in a Gaussian way. However the parameter

$$\Delta_{ND} = \phi_N - \phi_D \tag{7.3}$$

is. The significance from zero of the Δ_{ND} can therefore be used to estimate the significance from zero of a measured asymmetry for any value of A_{ND}.

In the construction of these measurements of the asymmetry the systematic errors naively cancel producing a purely statistical measurement. If this were the case then the measurement of the day-night asymmetry would be a simple matter. However it is important to ensure that these errors really do cancel and when they do not, the error must be carried through the day-night analysis to produce an accurate measurement of the systematic error on the day-night asymmetry. A systematic error that does not cancel is known as a differential systematic. The differential systematics due to livetime and sacrifice have already been discussed. In order to measure the day-night asymmetry all differential systematics must be measured and dealt with correctly. The table of systematics on the measurement of the charged current flux (table 7.1) shows the various effects that impact on the measurement of the flux. Only if an effect has a significant error associated with it must the differential systematic be measured. There is for example little point in attempting to measure the differential systematic in instrumental contamination, as given the size of the contribution this would be dominated by the statistical errors for many years worth of data.

There are two principal types of measurement for which the differential systematic must be measured: those due to energy and those due to reconstruction. The remaining systematics are either theoretical and therefore cannot have a differential systematic, or are too small to have an impact. It can however be important to show that a differential systematic is small for an effect like sacrifice to give added confidence to the day-night result.

Error Source	CC error /%	ES error /%
Energy Scale	$^{+6.1}_{-5.2}$	$^{+5.4}_{-3.5}$
Energy Resolution	±0.5	±0.3
Energy Scale Non-linearity	±0.5	±0.4
Vertex Accuracy	±3.1	±3.3
Vertex Resolution	±0.7	±0.4
Angular Resolution	±0.5	±2.2
High Energy γ-rays	$^{+0.0}_{-0.8}$	$^{+0.0}_{-1.9}$
Low Energy Backgrounds	$^{+0.0}_{-0.2}$	$^{+0.0}_{-0.2}$
Instrumental Contamination	$^{+0.0}_{-0.2}$	$^{+0.0}_{-0.6}$
Trigger Efficency	±0.0	±0.0
Livetime	±0.1	±0.1
Sacrifice	$^{+0.7}_{-0.6}$	$^{+0.7}_{-0.6}$
Earth Orbital Eccentricity	±0.2	±0.2
^{17}O, ^{18}O	±0.0	±0.0
Experimental Uncertainty	$^{+7.0}_{-6.2}$	$^{+6.8}_{-5.7}$
Cross Section	±3.0	±0.5
Solar Model	$^{+20}_{-16}$	$^{+20}_{-16}$

Table 7.1: The contributions to the systematic error on the measurement of the CC flux presented in [58]

Other effects such as livetime and the orbital eccentricity of the Earth become more important in the measurement of the day-night asymmetry. Differential livetime systematics are measured in chapter 6. The orbital eccentricity of the Earth is included in the Monte Carlo when the signal PDFs are generated to remove this source of error. This chapter will describe methods used to measure the important differential systematics, and the measurement of the day-night asymmetry in the raw neutrino data and the charged current and elastic scattering signals.

7.1.1 Data blindness for day-night analysis

Data analysis for the day-night asymmetry is carried out firstly by the use of an open data set that is different to that used for the flux analysis. Only data taken before July 1st 2000 is open for analysis. This provides 157.7 days of data for analysis, 64.6 days of which are day data. This data was analysed and the analysis must be frozen before the remainder of the data is examined, a process known as opening the box. Blindness restrictions do not apply for data outside the signal region. Thus the entire data set can be used for systematic checks.

7.2 Reconstruction systematics

There are two principal properties of reconstruction: vertex accuracy, a systematic shift in the reconstructed vertex, and vertex resolution. Of the two the vertex accuracy has the much bigger impact on the measurement of the flux, as this parameter is directly related to the number of target atoms for solar neutrino interactions. In order to measure these properties between day and night a point source of events present throughout data taking is desirable. Fortunately such a source exists, a patch of radio-activity on the acrylic vessel known as the Berkeley blob.

The Berkeley blob is so called because it was first discovered in studies carried out at LBL [85],[86]. These studies shows that the blob is a radio-active source equivalent to $9 \pm 3(stat.)^{+20}_{-5}(sys.)\mu g$ of thorium. The blob can then be used to calibrate the radio-activity of the acrylic vessel. It also provides exactly the type of standard candle required to measure the differential systematics due to reconstruction.

7.2.1 Identifying events from the Berkeley blob

Events that pass all the data cleaning cuts between 30 and 50 hits from the neutrino data are used and selected if they reconstruct within 50cm of the acrylic vessel. The blob is now clearly visible in the reconstructed angular co-ordinates of the events shown in figure 7.1. Also visible in this plot is a demonstration that the low energy backgrounds are not completely randomly distributed throughout the detector. In particular there appears to be an increase in the background seen around the equator. This plot can also be used to define the final analysis box for the blob: $0.65 < \cos(\theta) < 0.85$ and $0.3 < \phi < 0.7$.

As the Berkeley blob is a patch of extra radioactivity on the acrylic vessel it is not possible to identify the events uniquely. Instead a subtraction technique must be used. This is carried out by subtracting the events from the same size region from a different quadrant of the detector from the blob. As the backgrounds are not spread completely isotropically around the detector it is important that the backgrounds are the same as the ones in the region of the blob, with the exception of the blob itself. For this reason a box is chosen for this subtraction at the same height. Boxes in this region all cancel the residual background within statistics. The final cancellation box uses the same cuts as the blob, but with $x, y < 0$, or $0.3 < (\phi - \pi) < 0.7$ instead.

7.2.2 Reconstruction properties of the Berkeley blob.

Now that events from the Berkeley blob can be identified, the events can be used to test reconstruction systematics. The reconstruction of the events in x,y,z, and radius is shown in figure 7.2. The events can be fit well by a Gaussian in each co-ordinate. The width

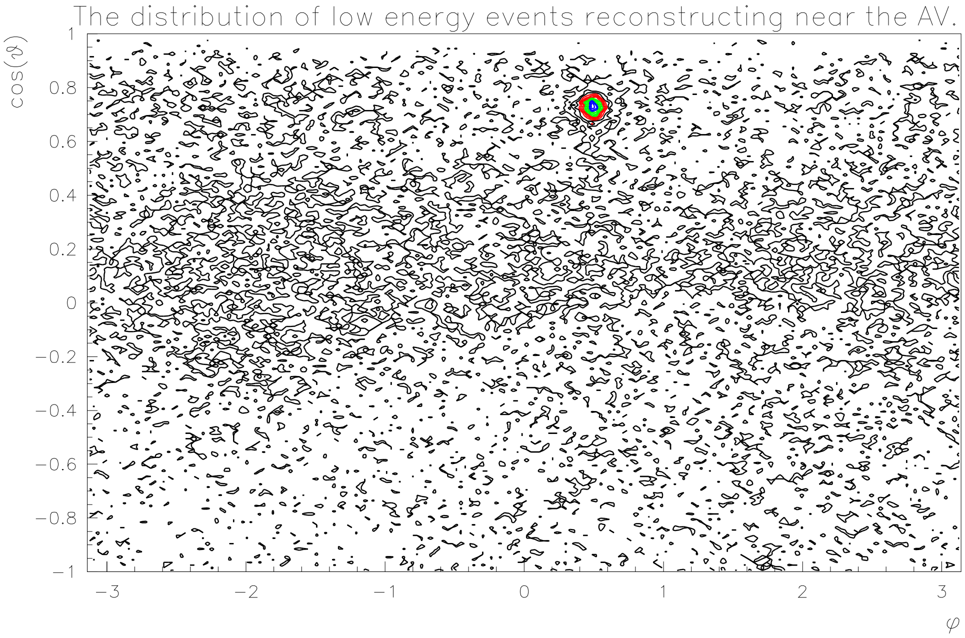

Figure 7.1: The reconstructed θ, ϕ distribution of low energy events near the acrylic vessel.

of the Gaussian is related to both the distribution of the radio-active contaminant, and the resolution of the fitter. Due to the unknown distribution of the contaminant and the unusual optics at the acrylic vessel, the blob cannot be used to measure fitter properties. However the blob is assumed to be distributed in the same way at all times, so the blob can be used to look for systematic changes in reconstruction as a function of time. In particular the blob is an excellent monitor to look for a systematic change in position resolution, or vertex accuracy between day and night.

Any shifts in vertex accuracy or position resolution should show up in the Gaussian fit to the radius as a change in the centroid and width respectively. In order to see if any effect is significant, the data is split into two in a number of different ways to compare with the day-night shift. In order to make a fair comparison the other splits are made randomly and in a way independent of any possible day-night effect. The results of the tests on both the centroid and width are shown in table 7.2.

The size of the shifts in both the centroid and the width compare well with the shifts seen when splitting the data set in the other independent ways. The shifts seen are also compatible with zero within the errors of the fitting procedure. These observations tell us that using the Berkeley blob the differential systematic errors in reconstruction are very small, and limit any residual vertex shift to less than a 0.2% radial shift.

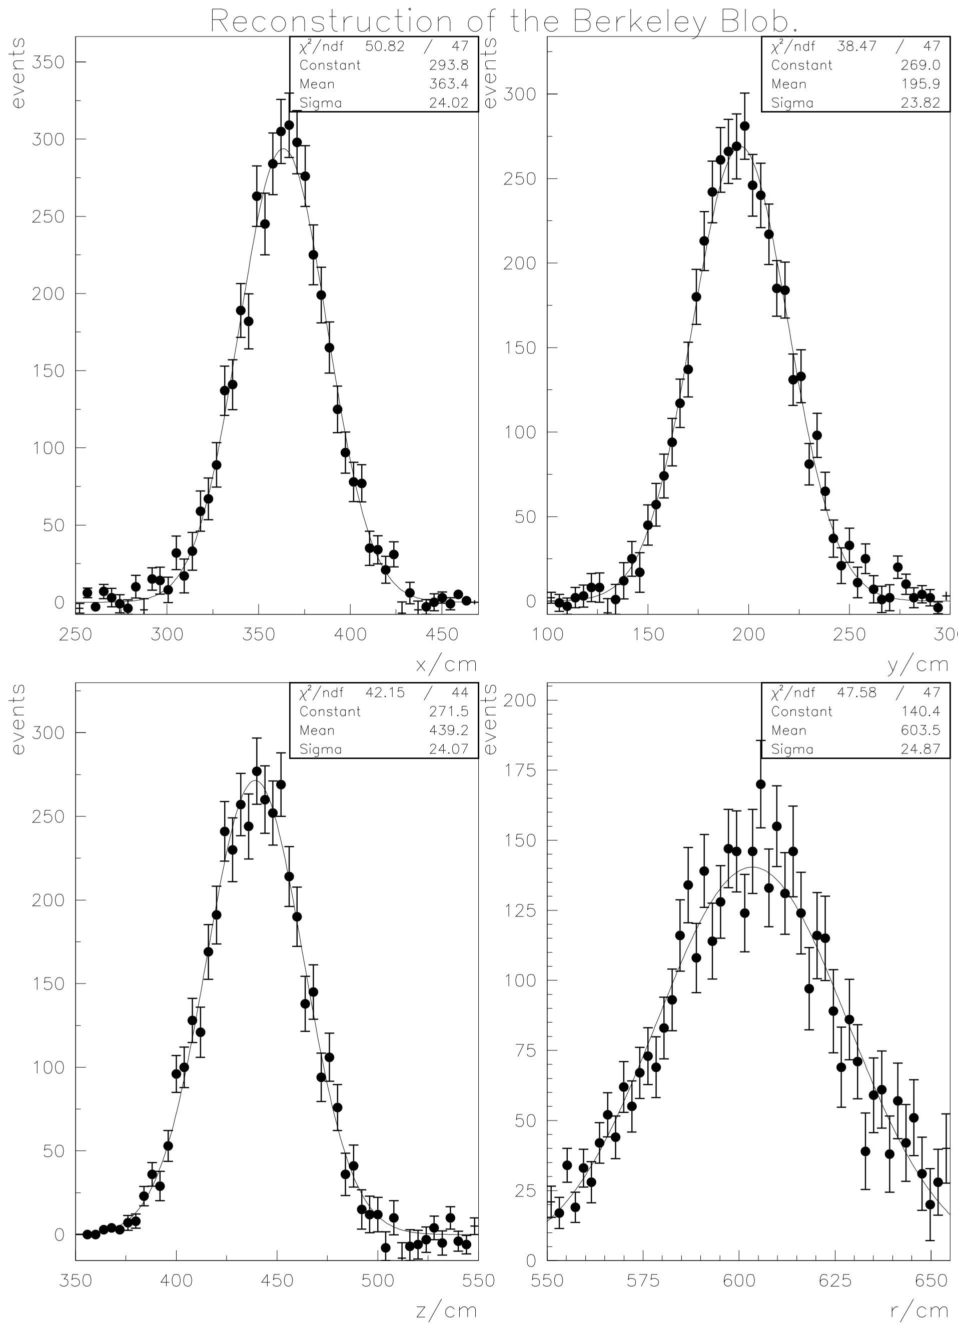

Figure 7.2: The reconstruction properties of the events in the Berkeley blob.

Method of Splitting	Shift in the centroid/cm	Shift in the width/cm
Day-Night	-0.3 ± 1.3	-1.0 ± 1.3
East-West	0.7 ± 1.3	-1.4 ± 1.3
Even-Odd Julian Day	-2.0 ± 1.3	-0.8 ± 1.3
Even-Odd GMT Hour	0.5 ± 1.3	1.0 ± 1.3
Even-Odd GMT Second	-0.3 ± 1.4	0.2 ± 1.3
Even-Odd GTID	1.4 ± 1.3	2.3 ± 1.3
Even-Odd Run	-0.8 ± 1.3	-1.0 ± 1.3

Table 7.2: The shifts in the centroid and width of the Berkeley blob, splitting the data set in various ways.

7.3 Energy Scale

The principle uncertainty in the analysis of data from SNO is the energy scale. For the study of the day-night asymmetry it is important to demonstrate that the energy scale is the same during the day and night. There are two effects that contribute to the differential systematic in energy scale, the energy drift and the absolute scale. The absolute scale in this case corresponds to the short term difference between day and night, perhaps due to changing lab conditions.

The drifting energy response of the detector means that the day and night energy scales can be different when integrated over livetime, even without any short term change. How well this can be corrected for depends upon how well we can understand the energy drift. This is a very important potential systematic for the day-night analysis.

There is no ideal way to look for short term variations in a way that spans the entire data set and does not also include the energy scale drift. However on at least one occasion a calibration source has been run overnight allowing a one off search for these effects to be made. Standard candles provided by muon followers and the Berkeley blob allow further measurements and verifications of this to be made.

7.3.1 The energy scale drift

At various times during data taking the ^{16}N source has been deployed into the centre of the detector in order to monitor the energy scale. It was found that the response of the detector changed as a function of time, a phenomena known as the energy scale drift. The black line in figure 7.3 shows the energy scale drift measured by the ^{16}N runs. The best fit to the drift is a linear piecewise model, with no drift after day 284 of data taking. One other thing this figure clearly shows is the unfortunate gap in the data between the first and second runs. This gap means that for data taken in this period there is more

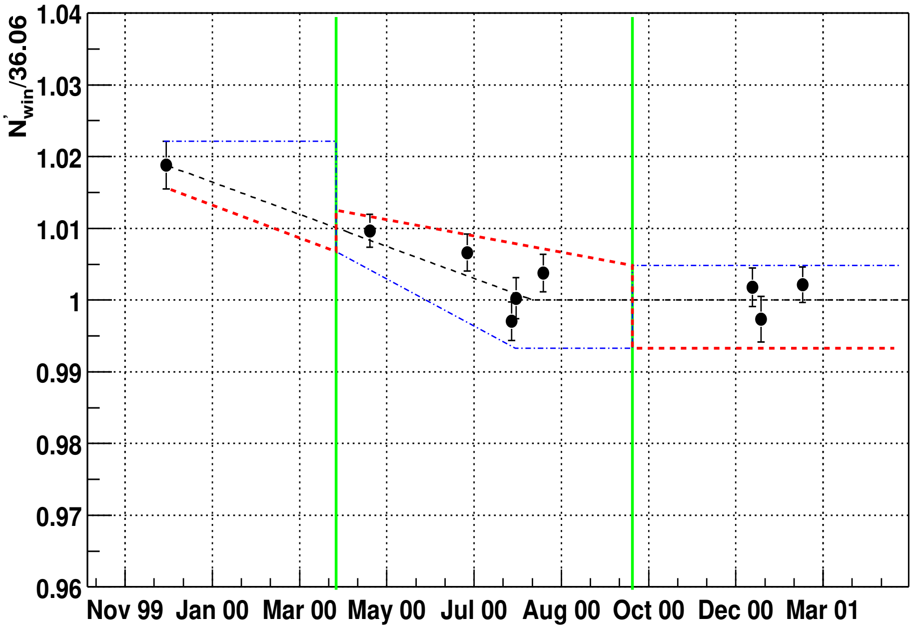

Figure 7.3: The energy scale drift for RSP using ^{16}N stability runs. For a description of the different fits to the drift see the text. Courtesy of [87].

uncertainty in the shape of the energy scale drift than is desirable.

If the energy scale drift is corrected for with no error, then the drift *cannot* introduce an artifical day-night asymmetry into the data. We do not however understand the drift completely and it certainly cannot be corrected for without error. Thus due to the way the data samples the drift a difference between the energy scales during the day and night will be produced. A limit on the size of this effect can be found using a technique designed by Scott Oser[87]. By constructing drift models that overestimate the energy scale during the summer[1] and underestimate the data during the winter, and vise versa, a limit can be found for this effect that maximises the problem for the day-night analysis. If the effect is small enough this technique is sufficient as a measurement of the systematic error from the drift.

The two extreme models are shown in red and blue in figure 7.3. These models are unphysical at the winter and summer boundaries (the green lines in the plot), but this does not matter as the analysis is attempting to set a limit, not make an accurate measurement. These different drift models can be used to directly measure the effect of the drift on the day-night asymmetry. The shift in the measured fluxes and asymmetry

[1]When the day is longer than the night

Condition	Start Time	End Time	Events
Day	17:53 - 25/05/2000	18:30 - 25/05/2000	1498
Night	20:00 - 25/05/2000	05:00 - 26/05/2000	22124
Day	07:30 - 26/05/2000	09:55 - 26/05/2000	6106

Table 7.3: The division of run 11850 into day and night segments.

using the alternative drifts rather than best fit gives a measure of the systematic error due to the drift.

7.3.2 The neutron source

At one point during the data taking the weak Cf source was left in the detector overnight. This occurred on the night of the 25th of May 2000, in run 11850. This occurrence provides a unique opportunity to verify that there is no day-night effect in the energy scale. As the source was only in for one night this does not check the energy scale as a function of time, however it does allow a number of the most troublesome day-night systematics to be checked. During that night, the lab was left in the same state as any other with the systems in the usual night configuration. This means that data from a calibration source can be used to verify that the night configuration does not differ from the day when most calibration data is taken.

The Cf source is a small amount of californium encapsulated within acrylic. The source provides neutrons via spontaneous fission of the Cf nucleus. The source also produces a large background of events from the fission products. These do not reconstruct a large distance from the source, and can be removed by the use of a radial cut around the source. Unfortunately the position of the source is not known[2]. For this reason the location of the source must be fit before the backgrounds can be removed.

The definition of day and night in this analysis differs from the usual one, based on the position of the sun. Instead day and night are defined by the lab conditions, which the source is being used to test. The run is split into three parts, based upon the conditions in the lab at the time. Two of these at the start and end of the run are considered day data. The remaining part is considered night data. Data is rejected where lab conditions are uncertain. The division of the run is shown in table 7.3.

The neutrons from the Cf source are not tagged, so event selection is more complicated than for the ^{16}N or ^{8}Li sources. Firstly the instrumental cuts and high level cuts are applied. The exception to this is that the NHIT burst cut is not applied, increasing

[2] The reason the source was left in the detector overnight was because the manipulator system developed a fault. The same fault resulted in the unknown location of the source.

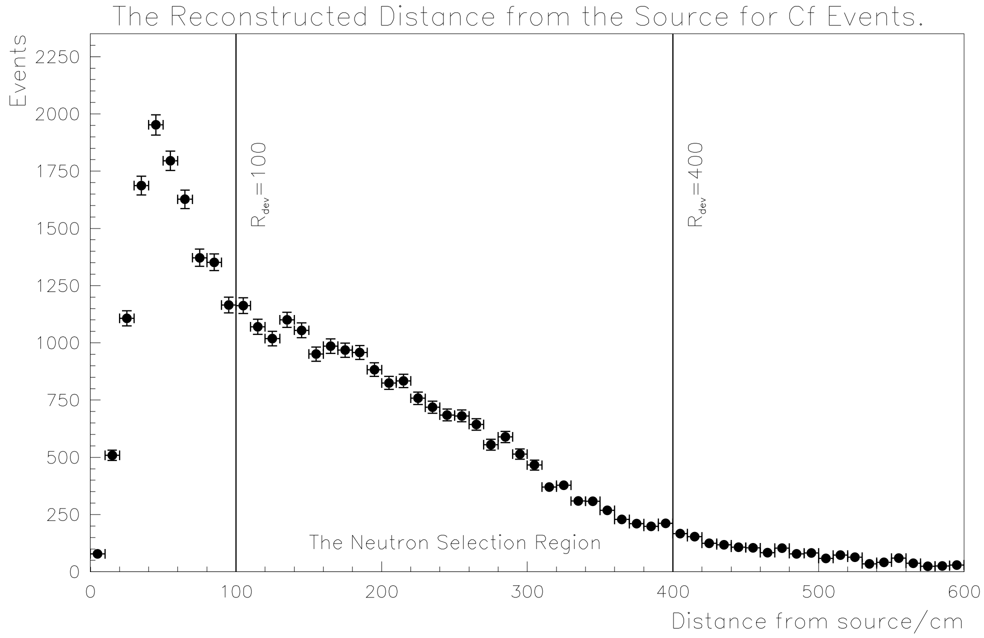

Figure 7.4: The reconstructed distance from the source for events in the Cf source run.

the available statistics, as the rate and multiplicity of the source means that a significant fraction of neutrons are removed by the cut. The next step is to find the source position. This is carried out by fitting the x,y and z co-ordinates of the source independently. The best fit position is (2.4,-22.8,269.7)cm. A fiducial shell is then set up around the source to find the neutrons. The inner radius is set to 1m, while 4m is used for the outer radius, as shown in figure 7.4. The excess of events inside the the 1m cut is due to fission products from the source. A fiducial volume cut at 5.5m is also applied, for comparison with other analyses.

Once this selection procedure has been carried out the NHIT and energy distributions are as shown in figure 7.5. By fitting the distribution to a Gaussian the centroid and width can be found. These are directly related to the energy scale and resolution. Splitting the data into day and night segments, the consistency of these results can be checked. The results are shown in table 7.4. The measured energy scale shift is $(0.34 \pm 0.66)\%$ using NHITs and $(0.62 \pm 0.55)\%$ using RSP energy, with the energy response less during the day than the night. However both of these measured values are compatible with zero.

The use of the neutron source in this way has two potential problems. The first is that the data was not taken at an ideal time, as May 2000 was not the best month

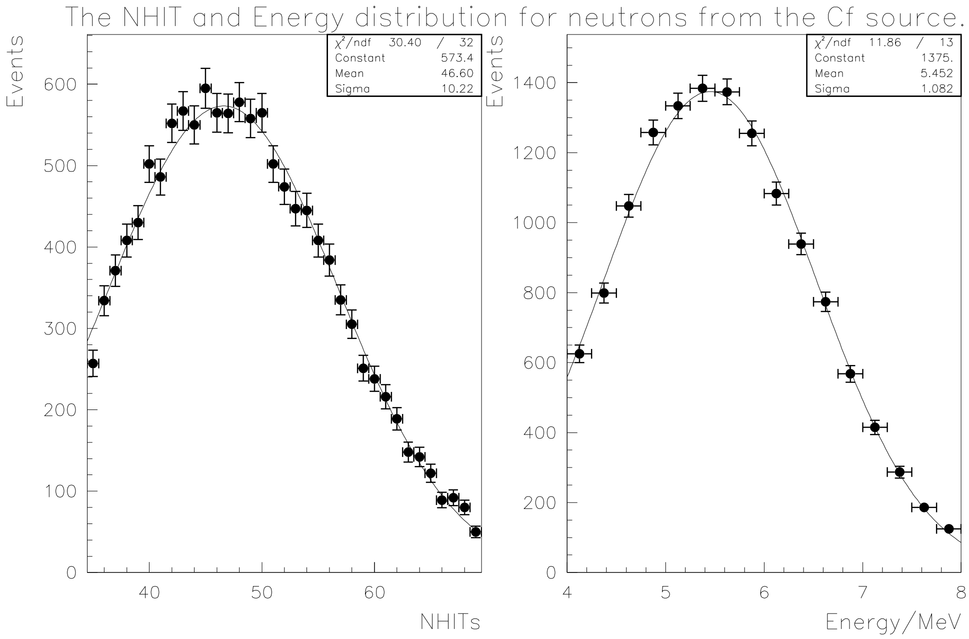

Figure 7.5: The NHIT and energy distributions for neutrons from the Cf source.

	NHITs		*RSP Total Energy/MeV*	
	Centroid	Width	Centroid	Width
Day	46.49 ± 0.27	10.08 ± 0.25	5.426 ± 0.026	1.075 ± 0.024
Night	46.65 ± 0.16	10.25 ± 0.15	5.460 ± 0.015	1.083 ± 0.014
Day-Night	-0.16 ± 0.31	-0.17 ± 0.29	-0.034 ± 0.030	-0.008 ± 0.028

Table 7.4: The energy distribution of neutrons from the Cf source, during the day and night.

for data taking as there were a number of problems, particularly with the chiller[3] at this time. On the night in question, the lab temperature was somewhat higher than usual. The neutron source also only tests one particular day, it does not test the entire data set. The run still provides an important check of the day-night effect on the energy scale in a way that is decoupled from the energy drift problem.

7.3.3 The Berkeley blob

Just as the Berkeley blob can be used for tests of reconstruction it can also be used to test energy scale. These tests are not as straightforward, however, because unlike most event sources used to estimate energy the only events analysed for the Berkeley blob are those in the tail. This causes problems as it is not possible to separate changes in energy scale from changes in rate. This means that the blob analysis could be complicated by any possible lifetime effects in the contaminant such as a half-life comparable to the lifetime of the experiment, or more likely a slow cleanup of a surface contaminant. There are additional computational complications that appear in the energy analysis of the Berkeley blob that are related to the analysis threshold. Events are only analysed above 30 hits, but due to the drifting detector gain and changes in the number of online tubes a 30 hit threshold does not translate into a fixed energy threshold. This does not however prevent the blob from being used to test for short term variations in energy scale, particularly to look for a different energy response between day and night.

The approach that is instead taken with the analysis of the blob, is to treat each run as a low statistics experiment on the blob. The use of individual neutrino runs means that the 30 NHIT analysis threshold provides a consistent energy threshold throughout each experiment, despite the fact that the energy threshold in each experiment is slightly different. The use of a short time scale also mitigates any long term variations in the behaviour of the blob or the detector. If there is a change in energy scale between the day and night this appears as an asymmetry in the rate of the events between the day and the night. Each experiment measures this only very poorly, but in combination this becomes a powerful measure of this systematic.

In order to measure the rate of the Berkeley blob runs with more than two hours of livetime in both day and night were selected from the neutrino run list. This gave 184 runs for the analysis, corresponding to 80% of the total livetime. Events were then selected as for the reconstruction test, for both the blob window and the subtraction window, and a simple counting experiment carried out. For each run the difference in the rate of events from the blob between day and night is calculated and stored. The results of this analysis for each run are shown in figure 7.6.

[3]The chiller is a piece of hardware within the environmental control systems in the lab.

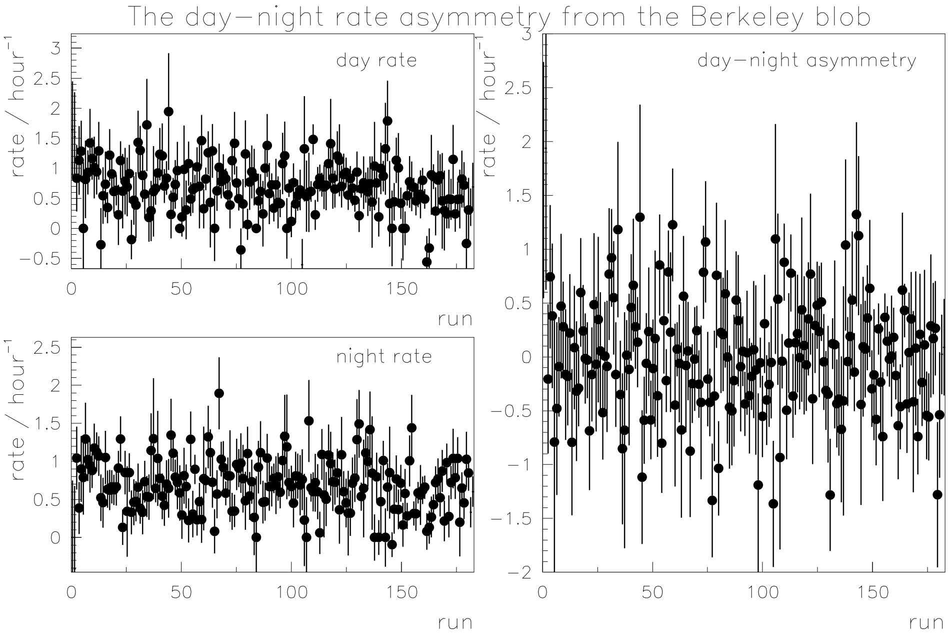

Figure 7.6: The rate of events from the Berkeley blob in day and night, and the asymmetry on a run by run basis.

In order to measure the overall asymmetry these results must be combined in the correct way. This is carried out in the usual way, using the following equations:

$$\delta^{bb}_{DNk} = (\frac{n_{Nk}}{t_{Nk}} - \frac{n_{Dk}}{t_{Dk}}) \tag{7.4}$$

$$\Delta^{bb}_{DN} = \frac{\sum_k (\frac{\delta^{bb}_{DNk}}{\sigma^2_k})}{\sum_k \frac{1}{\sigma^2_k}} \tag{7.5}$$

$$\frac{1}{\sigma^2} = \sum_k (\frac{1}{\sigma^2_k}) \tag{7.6}$$

where $n_{D/Nk}$ is the number of events in the day/night in run k, and $t_{D/Nk}$ is the livetime for the day/night in that run.

Once the results have been combined the residual day-night asymmetry in the rate of Berkeley blob events is found to be 0.009 ± 0.034 events/hour, compared to an average rate of 0.7 events/hour. The asymmetry is compatible with zero, but the sensitivity of this measurement can be used to limit any possible day-night changes in the energy scale.

In order to turn this measurement into a limit on any possible change in the day-night energy scale, the spectrum of the events must be considered. The NHIT spectrum of the events from the Berkeley blob is well fit by a falling exponential ($e^{-\alpha E}$). Using this fact it is then possible to calculate the change in (NHIT) energy scale required to cause the observed asymmetry. The fractional change in energy scale is given by

$$\frac{\delta E}{E} = \frac{\Delta^{bb}_{DN}}{R^{bb}\alpha E_T} + O(\frac{\Delta^{bb}_{DN}}{R^{bb}})^2 \tag{7.7}$$

where R^{bb} is the rate of events from the blob, and E_T is the energy threshold. Using this procedure the day-night asymmetry of the Berkeley blob limits the energy scale difference to $< 0.5\%$ at the one σ level.

7.3.4 Muon followers

Muon followers are dominated by neutron events. These events are present throughout the data taking of SNO. They can therefore be used to show whether the energy scale of the detector is consistent between the day and the night.

Muon follower events are selected by the use of the muon short cut. The instrumental and high level cuts are then applied. The exception to this is that the NHIT burst cut is not used in this analysis. This procedure selects a data set dominated by neutrons created by the passage of the muon through the detector. At lower energy the background wall is still important, and must be accounted for in the measurement of the energy scale.

The energy can be monitored by a Gaussian fit to the energy estimator for day and night data sets. As the livetime split in the data set is a 40:60 split between day and night, the statistics in the night sample are greater than the day. This results in a day measurement of lower accuracy than the night. The same is true for any systematic check that relies on data taken with the neutrino data.

The systematic is checked for both of the available energy estimators, NHIT and RSP energy. Both of these are corrected for the energy drift in order to avoid artificial day-night energy scale differences. The NHIT and energy distributions for the muon followers are shown in figure 7.7. The plot shows that the background wall becomes important at 4 MeV, or 35 hits. The Gaussian fit is not carried out in this region. The results, when the data set is split into two, are shown in table 7.5. There is no observable day-night effect in the energy scale measured by the muon followers. However this is not a strong statement due to the poor statistics of the muon followers. There are approximately 1000 events in use for the determination of energy scale, explaining the 2-3% error in the measurement. This error is larger than that determined for the energy scale. The poor statistics means that the muon followers can only be used to show that the day-night energy scales are

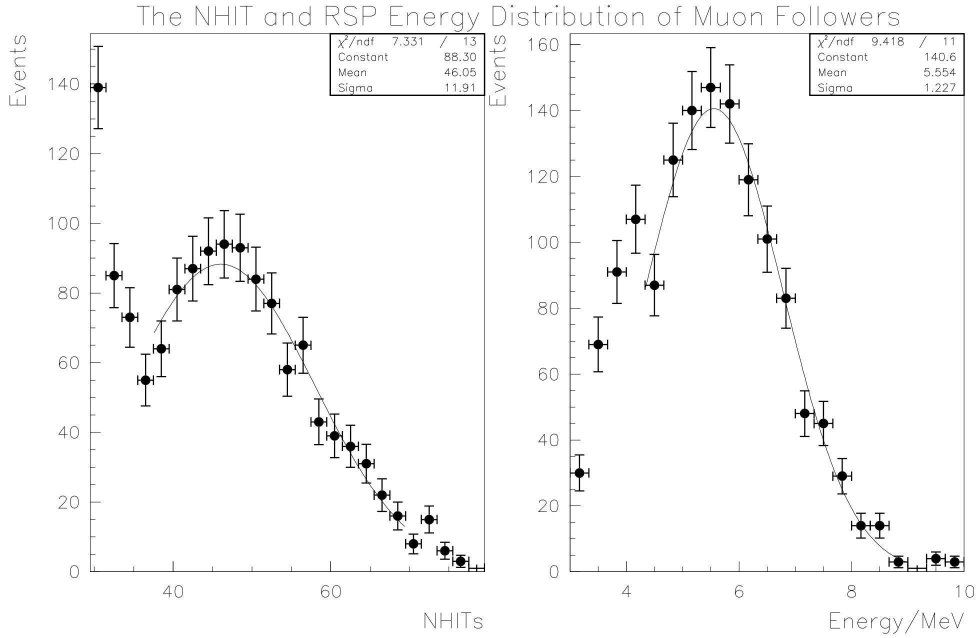

Figure 7.7: The NHIT and total energy distributions for muon followers

	Day	Night	Day-Night
NHITs	45.5 ± 1.8	46.1 ± 1.1	-0.6 ± 2.1
RSP Total Energy	5.49 ± 0.14	5.57 ± 0.07	-0.08 ± 0.16

Table 7.5: The centroid of the muon follower energy distributions during the day and the night.

consistent, as they do not contribute significantly to a combined measurement of the day-night differential systematic in energy scale.

7.3.5 The final energy scale systematic

The final energy scale systematic is split into two parts: the drift and a global shift in the scale between day and night. The drift is limited in such a way as to maximise any possible day-night asymmetry, while still being consistent with the data. This results in two different drift models that can be used to measure the systematic error due to this effect.

Three independent methods have been investigated to search for a global energy scale shift. Of the three, the use of the Berkeley blob is the most convincing as it does not have the problems of the other techniques. The muon followers do not have sufficient statistics to adequately probe interesting scale changes and are coupled to our understanding of the energy drift, while the neutron source provides data from only a single point in time. Each of the three measurements does not find a significant change in the energy scale, but the Berkeley blob sets the best limit to this, as well as the most convincing. Thus the size of this effect can be limited to a change of 0.5% in the energy scale from the day to the night.

7.4 The raw day-night asymmetry

At the time of writing the SNO collaboration is still operating the strict blindness policy for the day-night analysis. This forbids analysis of data from the signal region taken after June 2000[4]. The analysis in this thesis will therefore only present a measurement of the day-night asymmetry using open data. The use of only the open data does restrict the sensitivity of the day-night analysis, however, overall the open data taken corresponds to approximately half the total pure D_2O data set[5], the final statistical error bars can therefore be expected to be reduced by a factor of approximately $\sqrt{2}$ in the final high energy day-night analysis. The exact reduction will depend upon the detailed split of the livetime distribution.

The first step in the data analysis is an analysis of the raw data without signal extraction. The raw data can be used to test the stability of the asymmetry independent of signal extraction. In order to extract the raw asymmetry, the cuts (as described in this thesis) are applied to the selected data, which is then further reduced by the application of a fiducial volume cut at 550cm and a kinetic energy threshold at 6.75MeV. In the usual

[4] The cutoff occurs at run 12168.

[5] Including data taken up to May 2001, that has not been analysed in this thesis.

way an upper energy cut is applied at 19MeV to remove atmospheric neutrinos. The data is then split into two bins corresponding to day and night. There are 288 events during the day and 482 during the night. As was stated in chapter 6 these data sets correspond to 64.6 and 93.1 days of data respectively. These numbers give the results:

$$\Delta_{ND} = 0.72 \pm 0.35 \text{ days}^{-1} \tag{7.8}$$

and

$$A_{ND} = 0.149 \pm 0.074 \tag{7.9}$$

The raw data shows a day-night asymmetry, that is significant from zero at the 2.05σ level. The significance is estimated from the Δ_{ND} parameter. The significance using the asymmetry and its error is 2.01σ showing that the asymmetry is small enough that the Gaussian approximation is only just begining to break down, meaning that the error on A_{ND} is still meaningful for this measurement of the asymmetry. These errors are statistical only however, the next step is to propagate the systematics into the result.

7.4.1 Systematic errors

So far the systematic errors have been estimated as an effect on a detector property. The results of these measurements must be propagated to give the systematic errors on the asymmetry. This is achieved by a perturbative approach. For example the day-night energy scale has been estimated to be the same within 0.5%. In order to turn this into a measurement of the asymmetry the energy scale is perturbed in this way, and the effect on the asymmetry gives an estimate of the systematic error.

For effects like the drift, the systematic error is found by the application of the alternative drift models to the data. For other effects however there is a infinite set of ways that an asymmetry can appear. Taking the 1σ upper bound on the day-night energy scale difference as an example, the energy scale could be more at night than the day in any number of ways, bounded by the case where the night is higher than the best value by 1σ and where the day is less by 1σ. To measure the systematic the mean deviation from the central value of the two extreme cases and the case where the deviation is split equally between day and night is used.

Table 7.6 shows a summary of the systematic errors on the measurement of the raw day-night asymmetry. The first thing that can be seen in this table is that the systematic errors derived by the perturbative approach are not symmetric. The reason for this to be the case in the energy scale drift case is relatively straight forward. The drift has been bound using two separate drift formulae that are designed to maximise the impact of the drift on the measurement of the day-night asymmetry. These two models

Error Source	Error in $\Delta_{DN}/days^{-1}$	Error in A_{DN}
Energy Scale Drift	+0.096/-0.034	+0.020/-0.004
Day-Night Energy Scale	+0.12/-0.076	+0.027/-0.015
Vertex Accuracy	+0.025/-0.021	+0.005/-0.004
Livetime	$\pm 4 \times 10^{-4}$	$\pm 10^{-4}$
Sacrifice	$\pm$0.004	$\pm$0.001
Combined	+0.16/-0.086	+0.034/-0.016
Statistical Error	$\pm$0.35	$\pm$0.074
Total Error	+0.38/-0.36	+0.081/-0.076

Table 7.6: Summary of errors in the raw asymmetry measurement.

are however relatively ad-hoc and they do not bound the drift in such a way that the primary drift formula lies exactly between the two. The reason that the other systematics are not symmetric despite being symmetric when bounding a detector property is due to the perturbative approach itself. The systematic error bars are derived by perturbing a measured event parameter, such as the reconstructed position, by an amount set by the systematic study. When this occurs a number of extra events are then either gained or lost when fiducial or energy cuts are applied to the data. The change in the number of events is directly related to the derived systematic error. The number of these events is however subject to statistical fluctuation. A difference in the statistical fluctuation when the day-night asymmetry is increased compared to when it is decreased leads to an asymmetric error. The asymmetric error reflects the data taken by the detector and is therefore the correct systematic error to use. The use of different drift models is also a perturbative method, thus giving a further argument for the asymmetric error bars observed in the propagation of the error in the drift.

Table 7.6 also shows the systematic errors combined in quadrature with the statistical error. The statistical error totally dominates the systematic errors, and will continue to do so even with the reduction of $\sqrt{2}$ expected for the entire D_2O data set. The inclusion of the systematic errors reduces the significance from zero of the measured day-night asymmetry to the 2.0σ level.

As well as the systematic errors on the data, there is also a systematic effect due to the null hypothesis of no day-night asymmetry. This arrises due to the orbital eccentricity of the Earth, which causes a small change in the flux of neutrinos through the year due to the $\frac{1}{r^2}$ correction. As the data set integrates only part of this effect, the expected day-night asymmetry can be non-zero. This means that the significance of the measurement may not give the correct significance of the measured effect.

The expected day-night asymmetry is measured using the charged current PDF discussed below. Using the number of events on the day and night bins as well as the raw livetime[6] during the day and night (65.9 and 94.9 days respectively) the day-night asymmetry can be measured:

$$A_{ND}^{MC} = 0.010 \pm 0.004 \tag{7.10}$$

This asymmetry can then be used to estimate the value of Δ that would be expected given the number of events in the open data set:

$$\Delta_{ND}^{MC} = 0.048 \pm 0.019 \text{ days}^{-1} \tag{7.11}$$

The error bar in Δ_{ND}^{MC} is dominated by the statistical error from the Monte Carlo. A comparison with Δ_{ND} measured from the data can made in a straightforward manner as the errors in the two values are therefore almost uncorrelated. The significance of the raw asymmetry is found to be 1.90σ.

7.4.2 Systematic Checks

The first and most important systematic check to make is the stability of the asymmetry as a function of time. To do this the dataset is split into monthly bins based upon the start date of the run. This gives 8 different data sets with various amounts of livetime. Figure 7.8 shows the day-night asymmetry for each month of the data set. The data shows that while there is some variation in the asymmetry, the variation is no more than would be expected by statistics. This provides evidence that the observed asymmetry is not due to a short timescale fluctuation as could be expected for some possible causes of the asymmetry. The consistency of the measured asymmetries also suggests that long term variations in the detector are not having a large impact on the measurement.

The data can also be split in several other ways to search for systematics. This approach is also used, although it is important to remember that any effect seen must be compared to the trials penalty that is incurred by the use of this approach. The data has been split into two independent data sets in five different ways:

- East/West. The detector volume was split into two depending upon which side of the detector the event occurred.
- Inner/Outer. The detector is split into two equal volume radial shells using a further cut at a radius of 430cm.
- Up/Down. The detector volume is split into two, depending upon whether the event occurred in the top or bottom half of the detector.

[6] The Monte Carlo is run over the entirety of each run, the burst and muon follower cuts are not taken into account.

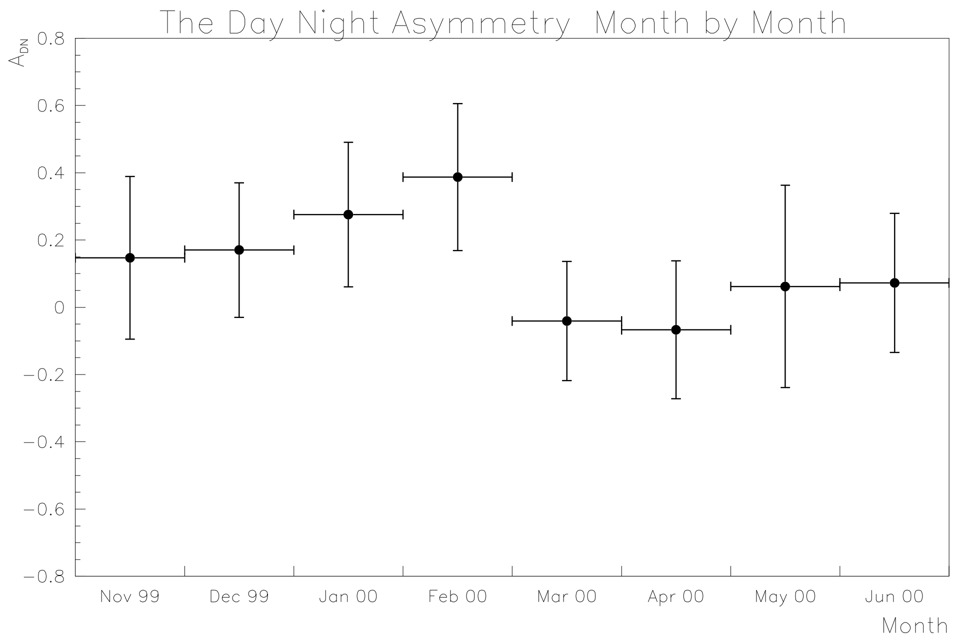

Figure 7.8: The day-night asymmetry for each month of the day-night open data set.

- High/Low energy. The data is split into two by the use of an further kinetic energy cut at 8MeV.
- Inwards/Outwards. The data is split by the use of a **u.r** cut to select inward or outward going events.

It should be noted that the energy test is not purely a systematic check, a difference could occur here due to neutrino oscillations via a day-night energy distortion.

The asymmetry parameters measured by splitting the data in these ways are shown in figure 7.9. The day-night asymmetries for each split should be compared to each other to look for any serious effects. Any other comparison should not be carried out as the measured asymmetries contain common data. Each data splitting produces asymmetries that are in agreement with their companion, showing that there is no systematic observable in this way.

7.4.3 A multibin analysis

Certain regions of MSW parameter space predict a significant enhancement of the ν_e flux for solar neutrinos that pass through the core of the Earth. In order to look for this effect a multibin analysis can be carried out. A multibin analysis is essentially the same

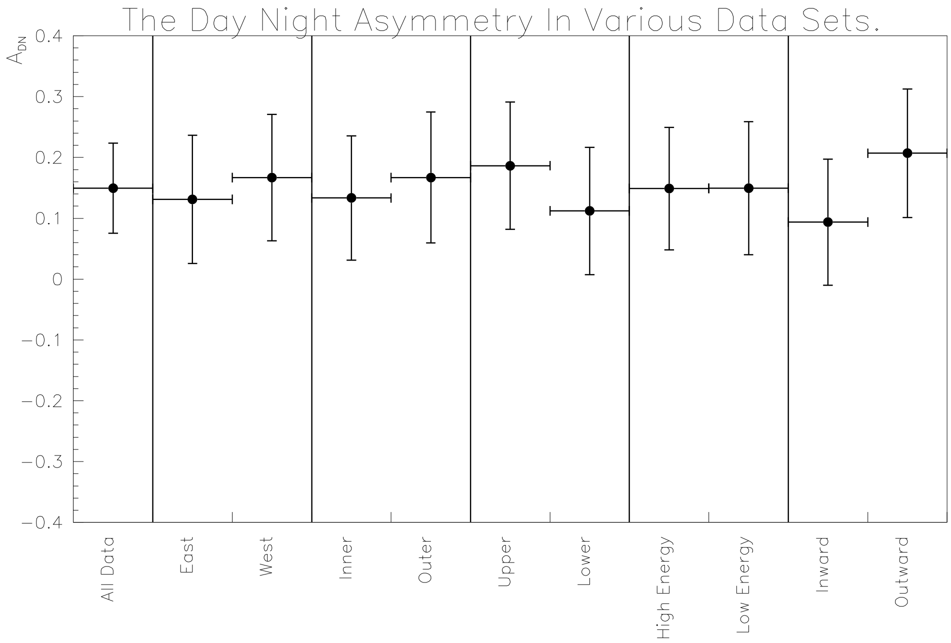

Figure 7.9: The day-night asymmetry measured by spitting the data up in various ways, see text for details.

as the measurement of the asymmetry except that instead of a cut on the solar zenith angle, the data is binned across this angle. Figure 7.10 shows the distribution when this is carried out. There is no evidence for any substantial enhancement in the core bin, but the statistical error bars are large, and it is not possible to make quantitative statements.

One thing of particular interest in the plot is the constancy of the asymmetry across zenith angle. Also shown in the plot is the result of a fit with separate day and night rates to the data. The best fit gives a measurement of the asymmetry that is essentially identical to that already discussed and gives an acceptable χ^2 for the fit. It must be stated however that a fit to a single rate also gives an acceptable (although slightly higher) χ^2 of 13.3 for 11 degrees of freedom.

7.5 Signal extraction

The next step in the analysis is to use the signal extraction technique to obtain the day-night asymmetry in the charged current flux. This is achieved by the use of a maximum likelihood technique using probability density functions (PDFs) based upon the energy,

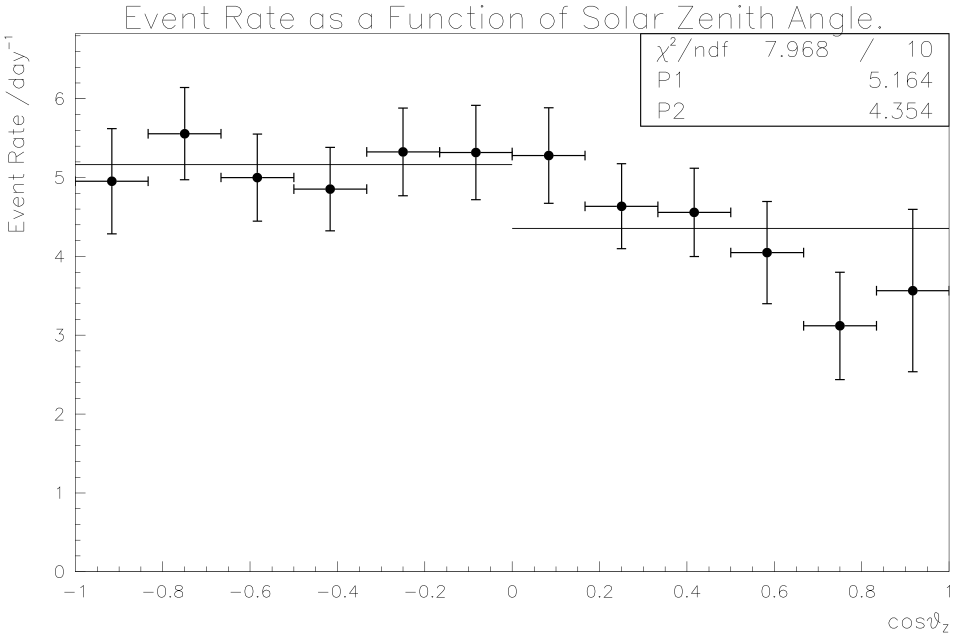

Figure 7.10: The event rate as a function of the cosine of the solar zenith angle. Also shown is a fit to a separate day and night event rate. Core data is contained in the first ($\cos(\theta_Z) < -0.83$) bin.

reconstructed position cubed[7] and the direction of the event with respect to the position of the sun ($\cos(\theta_{\odot})$). Signal extraction is carried out using code developed by Gordon McGregor [51].

The PDFs are generated by the simulation of each signal in each individual run in the data set. Each Monte Carlo is generated with 50$\times$ the number of events expected from the standard solar model. This allows the PDFs to be generated such that the statistical error in the data is dominant. The run-by-run approach also ensures that the signal is generated using the same conditions as applied during the actual run, including the number of PMTs online (from the DQxx banks) and the PMT noise rate measured by PulseGT data. Of particular importance for measurement of the day-night asymmetry measurement this approach also ensures that the integration over day and night periods is correct, with the exception of livetime corrections which can be accounted for. The orbital eccentricity of the Earth's orbit is also included in the Monte Carlo, to remove any effect this variation can have on the asymmetry. The combined day and night PDFs are shown in figure 7.11. The PDFs contain all the pertinent information about the three event classes in SNO. Above 6.75MeV the majority of the neutrons are removed from the data sample, this was indeed the prime reason for this choice of threshold for the CC analysis presented in[58]. Meanwhile inside 550cm ($(\frac{r}{600cm})^3$ =0.77) the radial PDFs for the CC and ES fluxes are essentially identical, so the $\cos(\theta_{\odot})$ distribution provides almost all the separating power for the ES signal. The small number of remaining neutrons are separated by the radial and energy distributions.

The signal extraction procedure does not allow events to be uniquely identified as coming from a particular reaction. Instead the most likely number of contributory events from each reaction is determined. Once the number of events in each reaction has been determined the day night asymmetry for each reaction can be found. Before the extraction is carried out, it is however informative to examine the three distributions used in the extraction in both the day and night data sets. The energy, radial and $\cos(\theta_{\odot})$ distributions are shown in figure 7.12. The shape of the energy and radial distributions agree between day and night, but this is not the case for the $\cos(\theta_{\odot})$ distribution. There is no apparent elastic scattering peak at night. This suggests that after signal extraction there will be a large elastic scattering asymmetry, causing the CC asymmetry to be increased over the asymmetry measured in the raw data. The results of the signal extraction procedure on application to the data are shown in table 7.7.

As well as providing the PDFs for the signal extraction, the Monte Carlos also give the prediction for the expected number of events. This can be used to find the ratio of

[7]The cubic radius is used as CC events are volume weighted. The actual variable used is the$(\frac{r}{600cm})^3$ variable used in chapter 5.

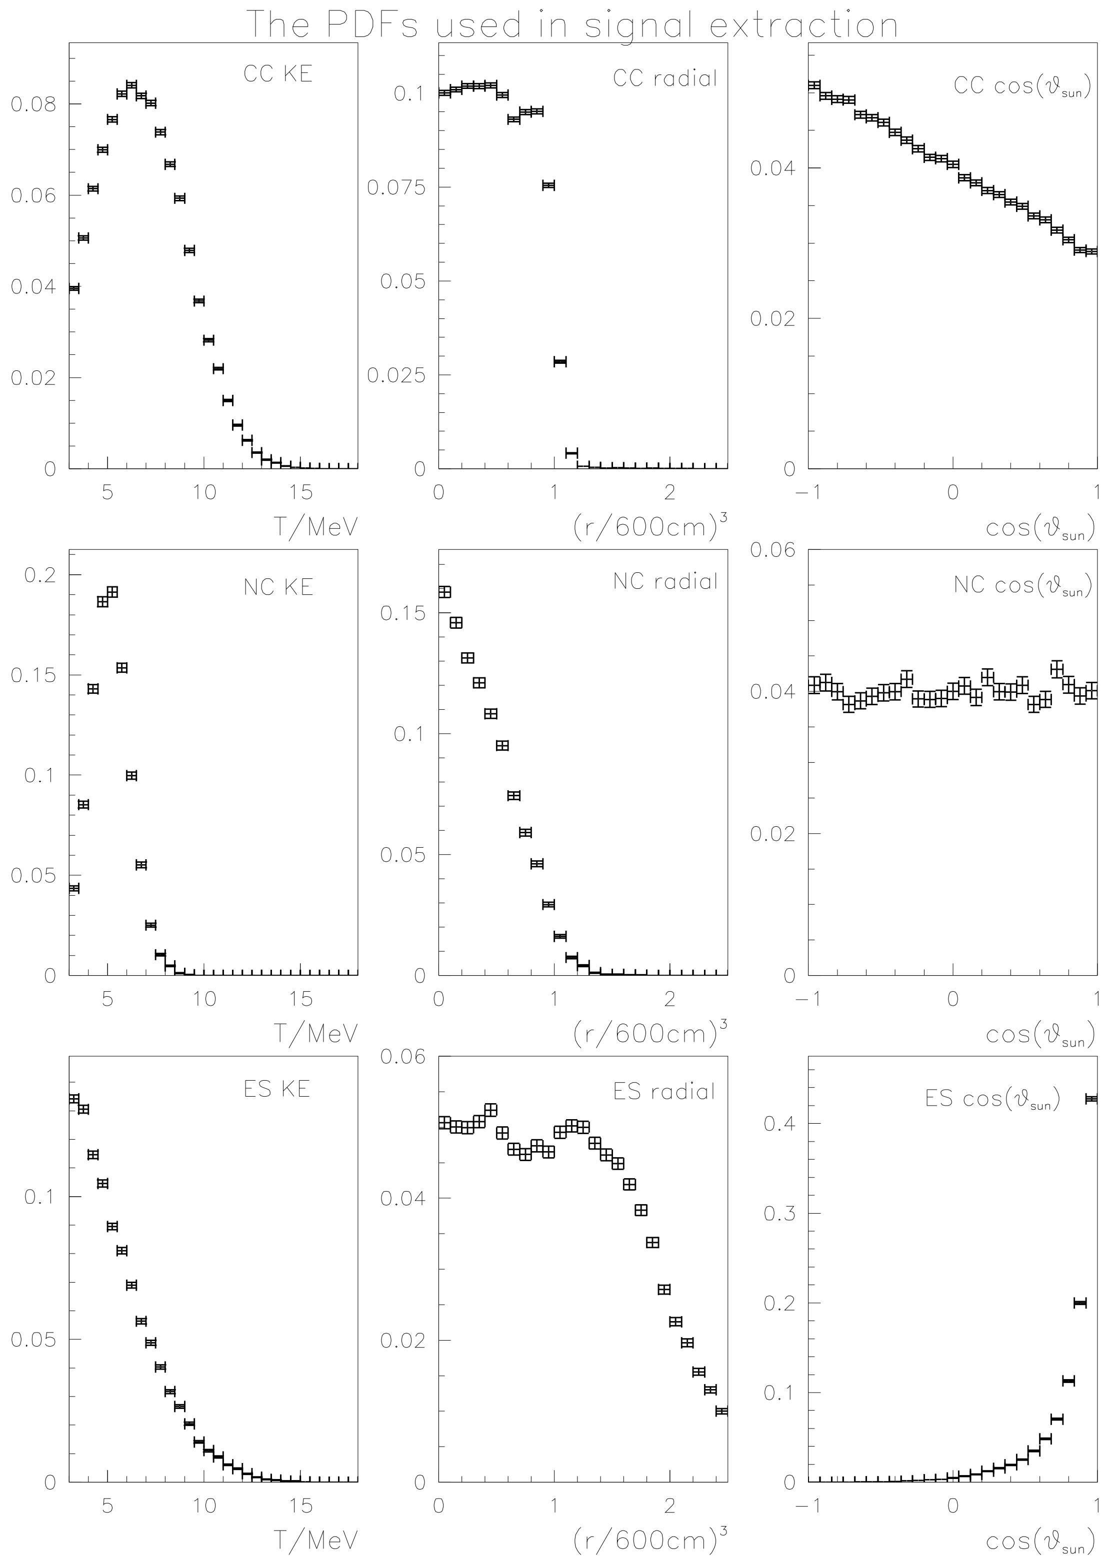

Figure 7.11: The combined day-night PDFs used in signal extraction

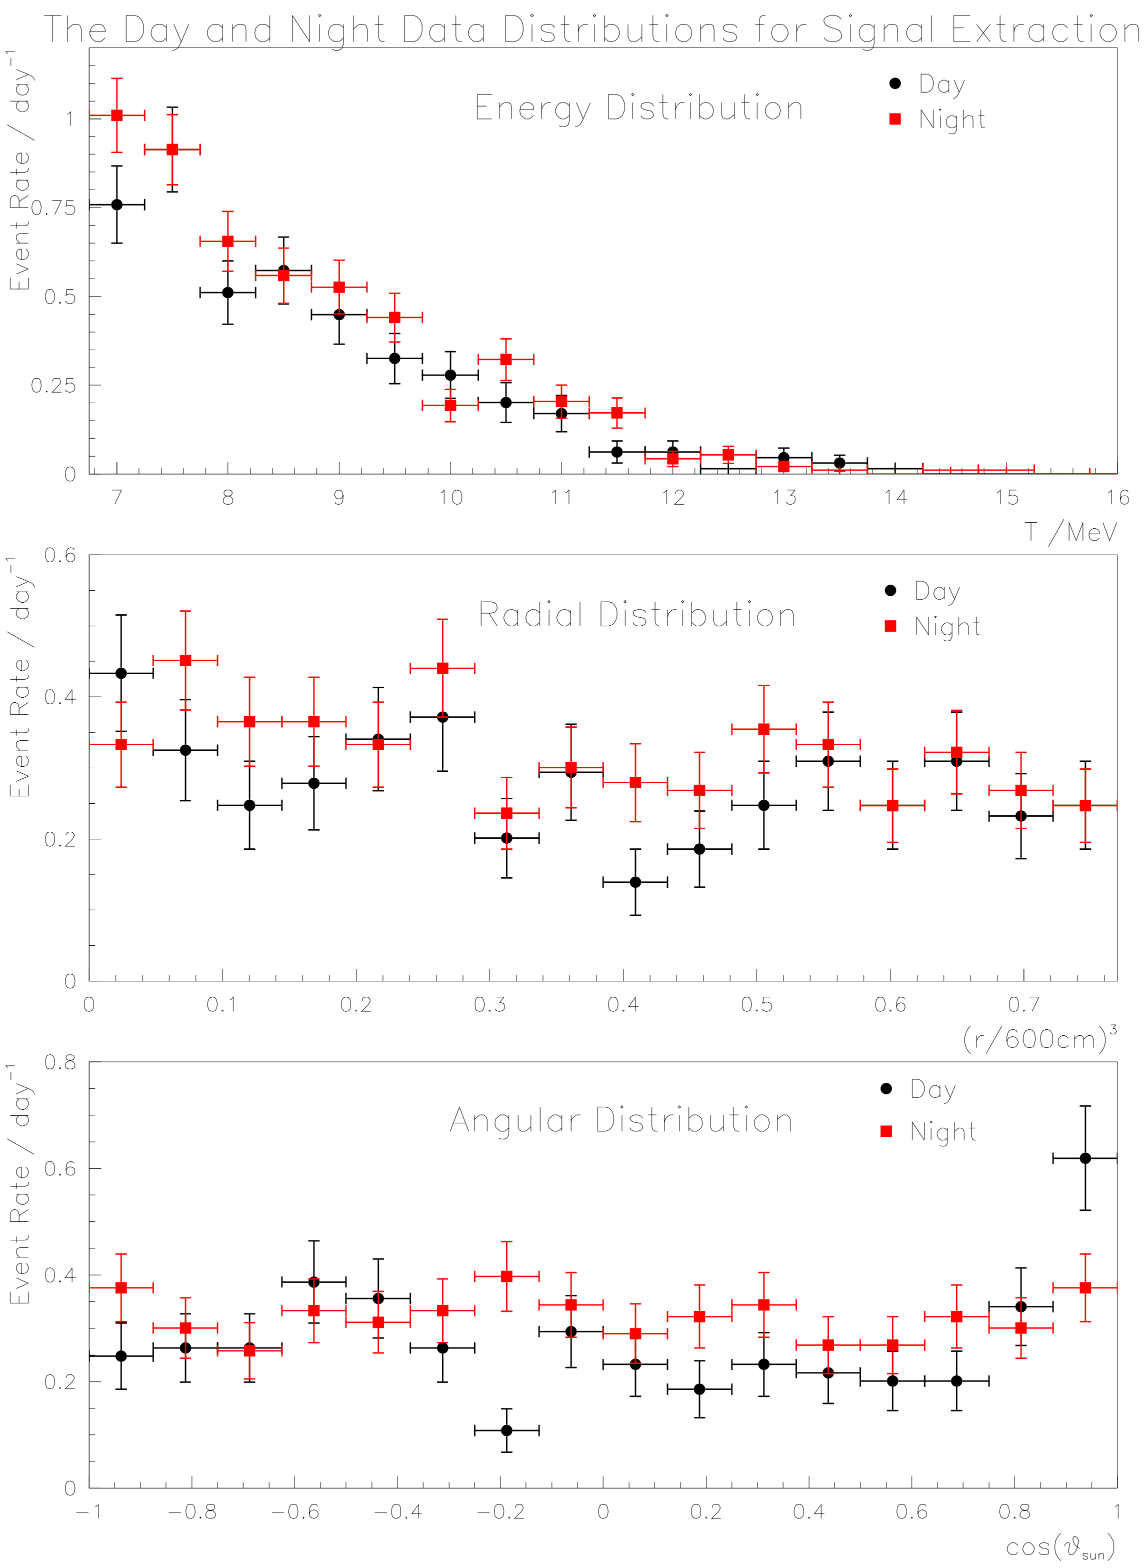

Figure 7.12: The day and night distributions used in signal extraction.

Reaction	Day Events	Night Events
CC	220.3 ± 19.1	424.3 ± 26.3
ES	43.2 ± 9.1	23.3 ± 9.1
NC (Neutrons)	24.5 ± 11.8	34.4 ± 16.6

Table 7.7: The extracted number of events of each class in the day and night data sets. Errors are statistical only.

Reaction	Day Rate	Night Rate	Δ_{DN}	A_{DN}
CC	0.299 ± 0.026	0.395 ± 0.025	0.096 ± 0.036	0.277 ± 0.105
ES	0.683 ± 0.144	0.268 ± 0.105	-0.415 ± 0.178	-0.872 ± 0.359
NC (Neutrons)	1.82 ± 0.87	1.70 ± 0.82	-0.12 ± 1.20	-0.07 ± 0.68

Table 7.8: The fluxes and the day-night asymmetry in the three reactions. Errors are statistical only.

the observed to predicted number of events, allowing the effect of the eccentricity of the Earth's orbit to be removed from the calculation. The day and night fluxes are shown in table 7.8 along with the measurement of the day-night asymmetry.

As expected from the $\cos(\theta_{\odot})$ distribution the CC asymmetry has been increased above the raw asymmetry due to the negative asymmetry observed in the ES reaction. The measurements the number of neutrons are extremely uncertain due to the choice of energy threshold and will not be considered further. Before the significance of the asymmetries is considered the systematic errors should be found. These are measured in a similar way as was used for the raw asymmetry, except the perturbed distributions are then sent to the signal extraction code. A summary of errors is shown in table 7.9.

As was the case for the raw asymmetry the measurement of the asymmetry parameters is dominated by the statistical error. The combined error can be used to estimate

Error Source	Δ_{DN}^{CC}	A_{DN}^{CC}	Δ_{DN}^{ES}	A_{DN}^{ES}
Energy Scale Drift	+0.002/-0.008	+0.003/-0.021	+0.005/-0.010	+0.035/-0.024
Day-Night Energy Scale	+0.005/-0.007	+0.012/-0.020	+0.011/-0.008	+0.029/-0.010
Vertex Accuracy	+0.001/-0.007	+0.004/-0.020	+0.015/-0.005	+0.046/-0.004
Livetime	$\pm 6 \times 10^{-5}$	$\pm 2 \times 10^{-4}$	$\pm 2 \times 10^{-4}$	$\pm 5 \times 10^{-4}$
Sacrifice	$\pm 6 \times 10^{-4}$	± 0.002	0.002%	± 0.005
Combined	+0.005/-0.013	+0.013/-0.035	+0.019/-0.014	+0.065/-0.024
Statistical Error	± 0.036	± 0.105	± 0.178	± 0.359
Total Error	+0.036/-0.038	+0.106/-0.111	± 0.179	+0.365/-0.360

Table 7.9: Summary of errors in the extracted asymmetries.

the significance of each asymmetry. In the case of the extracted fluxes the effect of the orbit eccentricity of the Earth has been included in the normalisation to the standard solar model. This means that the significance of the result should be calculated from zero asymmetry. The significance is as usual calculated from the Δ_{DN} parameter as the asymmetries measured are quite large. The charged current asymmetry is found to be 28% significant at the 2.53σ level. The asymmetry is higher than measured in the raw data because of the measured asymmetry in the elastic scattering signal of -87% with a significance of 2.39σ.

At this point the best way to reduce the error bars on the asymmetries would be to add more data to the analysis. As stated earlier approximately half the data is currently blind for the day-night analysis. Furthermore the reduction of the energy threshold can also add more data, at the expense of a more complicated signal extraction due to the addition of a significant number of neutrons. The addition of this extra data could potentially reduce the error bars on the charged current asymmetry by a factor of two.

The reduction of the analysis threshold introduces significant complications to the analysis that are avoided by the choice of the relatively high threshold used in this thesis. A lower threshold produces a neutral current measurement with significantly higher statistics, at which point it is important to understand the photodisintegration background. The contribution of the background wall due to radioactive decays in the detector also becomes much more significant with the reduced threshold. At the time of writing these analyses have not been completed and the day-night asymmetries they could introduce are not understood. The decision of the SNO collaboration is therefore that blind data must remain blind at this time, and the box will not be opened in this thesis.

7.6 Conclusion

Using the day-night open data set the raw asymmetry of the data is found to be 15%. This measurement is 1.90σ above the expected value for the data set. The data can then be split into the three components in the usual way using signal extraction to find the asymmetry in the these components. The charged current asymmetry is 28% significant at the 2.53σ level, while the asymmetry in the elastic scattering signal is -87% with a significance of 2.39σ. The asymmetry of the neutral current is not considered due to the choice of threshold which removes the majority of neutrons.

The asymmetry in the charged current and elastic scattering signals both occupy interesting regions of parameter space, but the error bars on both are rather large. The errors in these measurements are both dominated by statistics, and can ultimately be expected to reduced by a factor of at least $\sqrt{2}$ by the addition of blind data and more if the energy threshold is reduced. At this point it is likely that it will be possible to make quantitative statements about the asymmetry that are not possible with the more limited statistics of the open data set.

Chapter 8

Conclusions

One of the biggest obstacles to the analysis of data from SNO comes from the observation of the high rate of instrumental backgrounds recorded by the detector. These backgrounds overwhelm the solar neutrino signal and must be removed before data analysis can proceed. This is achieved through the use of data selection cuts that remove instrumental backgrounds by the use of five different characteristics:

- The PMT time distribution.
- The charge deposited on the PMTs.
- The geometry of the PMT hits.
- The presence of veto tubes
- The time correlation of the events.

These characteristics are used in a variety of ways, often in combination with each other in the construction of two different sets of cuts: the instrumental cuts and the high level cuts. The design of the cuts in these two sets differ in philosophy. The instrumental cuts are designed to remove backgrounds, while the high level cuts accept signal events. The entire process of removal of instrumental backgrounds is known as data cleaning.

As in any situation where cuts are applied to a data set, they can also remove signal events. The fraction of signal events removed is known as the sacrifice. The sacrifice of the cuts comes in two distinct forms, the first of these arises from cuts that consider individual events, the second from livetime corrections. The sacrifice of the former set of cuts is measured by the use of calibration sources, in particular the ^{16}N source. The biggest systematic in the measurement of the sacrifice is the stability of the cuts. This is monitored using calibration data taken periodically throughout the year, data from the pulsed global trigger and the data itself. The sacrifice of the cuts for electrons inside the analysis region for the charge current reaction is $1.46^{+0.54}_{-0.45}\%$ while, for neutrons, the sacrifice is slightly higher at $1.94^{+0.58}_{-0.50}\%$. Sacrifice is also imposed by the cuts that remove events based upon time correlations. These cuts operate by the removal of blocks of time in which the data fails some criteria. The sacrifice of these cuts therefore takes the form of a livetime correction. These cuts remove 1.86% of the livetime, a figure dominated by the muon follower cut.

As well as removing neutrino events, the cuts can also fail to remove background events. The fraction of instrumental backgrounds left in the data sample after the application of the cut is known as the contamination. By the use of a bifurcated analysis, which measures the contamination by comparing the results of two independent sets of cuts on the data, a limit can be set on the contamination of 0.1%. The results of the sacrifice and contamination analyses show that the instrumental backgrounds can be removed in a controlled way which leaves the majority of the solar neutrino signal intact. In other words they show that the data cleaning process works. The systematic errors on data cleaning are also small enough that the measurement of the flux is dominated by the errors on the energy scale and reconstruction as well as the statistical error, showing that data cleaning is sufficient for the pure D_2O phase of operation.

One use for the raw (uncut) data taken by the detector is to measure the livetime. This is achieved by the use of the time of the first and last event of each run. Using this approach the livetime of the selected data set for data taken between the 2$^{\text{nd}}$ of November 1999 and the 15$^{\text{th}}$ of January 2001 is 241.0 days after correction for the applied cuts. The livetime must also be split into day and night bins for measurement of the day-night asymmetry. The day-night data analysed in this thesis was taken between the 2$^{\text{nd}}$ of November 1999 and the 1$^{\text{st}}$ of July 2000 and provides 64.6 days of day data and 93.1 days of night data.

Once the instrumental backgrounds have been removed it is then possible to analyse this data to find the day-night asymmetry. The systematic errors on the asymmetry come from the same sources as for a flux analysis. But for the day-night analysis the emphasis is on the accuracy with which it can be shown that the detector behaves in the same way during the day and during the night. The systematics are measured using a combination of calibration data and identifiable contiguous Čerenkov backgrounds that occur in the detector. The systematics of the asymmetry are found to be much smaller than the statistical errors.

The asymmetry of the raw data is 15% and is different from the case of no asymmetry at the 1.9σ level. The raw data is passed through a signal extraction procedure in order to find the asymmetries in the charge current and elastic scattering signals. The charge current asymmetry is 28% and is significant at the 2.53σ level. The elastic scattering signal on the other hand shows an asymmetry of -87% with 2.28σ significance.

The central values of the asymmetries lie in interesting regions, but the statistical errors are large. Without the additional data that will be obtained when the day-night blindness box is opened it is not possible to make quantitative statements about the measured asymmetries. This thesis however lays out the day-night analysis that can be undertaken for the full data set once the decision has been made to open the blindness box.

Appendix A

DAMN banks

Due to the large number of instrumental cuts it is desirable to be able to store the results of those cuts. This avoids the situation where the cuts which are CPU intensive to apply and require the entire data set have to be processed time after time. To achieve this, the results of instrumental cuts are stored in special titles banks, known as DAMN banks.

DAMN (Data Analysis Mask Numbers) banks contain entries for every event in the data stream, consisting of three data words. The first word if the event GTID, required to uniquely identify each event. The second and third words are thirty-two bit bit-masks, the applied mask and the result mask. Each of the thirty-two bits then corresponds to a different cut of tag. Having two bits-masks per bit allows each bit to have four settings, three of which are used, see table A.1. Each of the three settings has a specific meaning when used to store the results of the instrumental cuts.

- True: The event has been identified by the cut/tag as an event that fits the criteria of the cut/tag. There is no merit value when a bit is set specifying if an event is 'good' or 'bad'.
- False: The event does not meet the criteria of the cut/tag.
- Undefined: The test corresponding to this bit has not been applied, or this bit does not yet have a corresponding cut or tag.

Applied mask bit	Result mask bit	Result
1	0	False
1	1	True
0	0 or 1	Undefined

Table A.1: The possible setting for each bit in a DAMN bank.

DAMN banks are associated with individual runs. To keep this clear the bank number if set equal to the run number. DAMN banks are also unique in that they are not loaded through titles. If the DAMN bank is not being loaded from the SNO database, it has to be loaded from a cache directory. Each file containing the DAMN bank from a particular run has a unique file name, which allows the DMM processor to load the required bank. DMM only has to be pointed to the cache directory and it will find the required bank.

Having three words per event can obviously lead to some very large files. To try and reduce the amount of data in a DAMN bank some compression is used. This takes the form of a default applied and result mask. When these are equal to the default they are omitted. The fact that they are omitted is then specified in the leading (and unused) digits of the GTID. The space saving that comes from this is rather large, particularly as all the events have the same cuts applied to them. This of course means that the applied mask is not usually saved event by event. The saving from the result mask is not so good, but is still considerable. This is because the majority of SNO data is due to low energy backgrounds the majority of which pass all the cuts.

However DAMN banks for larger runs are still very large and after run 10614 DAMN banks are stored on a subrun by subrun basis. To achieve this the subrun number was added to the zdab in the builder. This means that the change is not backward compatible and that there are still some very large banks associated to runs before 10614. These DAMN banks require special handling.

As well as DAMN banks which are designed to hold the results of different tests for the entire collaboration, there are also DARN (Data Analysis Request Numbers) banks. These banks can be used to store the results of test, by an individual user.

DAMN banks are accessed in three ways.

- Via the DMM processor: This is most useful when setting bits in a DAMN bank, but can also be used to test bits to filter events.

- Via DQFs: These allow bits to be tested to allow more complicated filtering that can be achieved using the DMM processor. This also allows DAMN bank results to be written to ntuples for further study.

- Via user code: This allows more complicate analysis of the results held in the DAMN banks.

At the current time only DAMN bank 0 is defined for the use of the collaboration. This bank holds mostly the results of the instrumental cuts, but does also hold the results of a few other useful tags that are checked at the same time. The layout of the bits of DAMN bank 0 is shown in table A.2

Bit	Test	Bit	Test
0	Retrigger	16	QCluster
1	First Event	17	Muon
2	Burst	18	Muon Follower Short
3	Charge	19	Muon Follower Long
4	Cluster	20	ITC
5	QvT	21	FGC
6	Q/NHIT	22	NHIT Burst
7	Crate Isotropy	23	OWLEHi
8	AMB	24	Subrun Boundary
9	FTS	25	Zero-Zero
10	OWL	26	Salt Blindness
11	Junk	27	Undefined
12	Ring of Fire	28	Undefined
13	Neck	29	Undefined
14	ESUM	30	Undefined
15	HIB	31	Undefined

Table A.2: The layout of the DAMN 0 bank.

Appendix B

The Runlogger

The runlogger is a system used in SNO as a book-keeping system. A record is kept of a variety of statistics and histograms for each run and stored in the database. Understanding of this system provides a useful complement to the livetime discussions in chapter 6.

B.1 The original runlogger

The runlogger was originally designed as a single pass process that would collect statistics and histograms on each run in order to store them in the database. This data can be retrieved at a later time in order to carry out tasks such as run selection without the need to resort to processing the raw data.

There is a large amount of information stored in the runlogger for each run. The highlights are:

- The Run Header Bank: RHDR

  The run header bank is written by the DAQ at the start of each run and contains run start information including:

  - The Start Date. The MAC time at the beginning of the run.
  - The Run Mask. A bit mask specifying the conditions of the run. It is here that tags like the unusual circumstances or ECA run are set.
  - The Source Mask. A bit mask recording which calibration sources are deployed.

- The Trigger Bank: TRIG

  This bank records which triggers are masked in for the run and the trigger thresholds.

- The Runlogger Bank: RLOG

This bank hold the derived statistics for the run. A large amount of information is held here including:

  - Generation statistics. The date, time and machine where the runlogger information was generated.
  - The start and end time of the run. The end time is taken from the 10MHz clock. However the start time was taken from the RHDR bank. This unfortunately leads to incorrect livetime information in the database, when the RHDR time is incorrect, which can happen as this is linked to the MAC. This problem is most noticeable at daylight savings changes.
  - Run/event statistics. The number of events in total and sorted by trigger bit. The number of channels online, the number of subruns.

- Run Histograms

  There are also a large number of histograms created and stored by the runlogger. These are used to store hit information run by run.

B.2 Additions to the runlogger

The runlogger is very CPU intensive to create, so rather than recreate the RLOG structure each time it is desirable to add something to the runlogger a separate system has been created. This works along different lines to the original runlogger. The additional system creates a master bank, to support supplementary banks containing the information to be added to the runlogger.

The master bank is the runlog analysis information bank, or RLAI. This is also written to the database and stored along side the RLOG banks. When additional information needs to be added to the runlogger, the RLAI bank is taken from the database the new bank, or banks are added to RLAI and filled. The RLAI bank is then written back into the database. This approach avoids the recreation of any of the data already in the database. The generation details of the latest updates to the RLAI system are stored in the RLAI bank.

At the current time there are three supplementary banks attached to RLAI.

- The Runlog Dead times Bank: RLDT

  The RLDT banks is used to store live and dead time data on a run by run basis. The run times from both the 10 and 50MHz clocks are stored here. The bank also stores the dead time introduced by the cuts that introduce dead time, such as the burst or muon follower cuts. The combined dead time due to all the data cleaning cuts and the muon follower short cut is also stored in this bank split into hourly and solar zenith angle bins.

- The Runlog Electronics Status Bank: RLES

  The RLES bank stores some simple derived quantities from the DQxx banks, that can be used to track these properties as a function of time. As well as the number of channels on-line for each class of PMT, data concerning the number of triggers enabled and the average threshold of the PMTs is also stored.

- The Runlog Clock Synchronisation Bank: RLCS

  The RLCS bank stored the 10 and 50MHz time of the first event of each run. This information is used to rebuild the 10MHz clock in the situation where the this information is missing. This then allows the data to correctly obtain the banks form the database.

B.3 Access to runlogger data

The runlogger data is retrieved from the database by the use of the runlog reader processor. This is a SNOMAN processor, which allows the data from the data to be processed in the usual way, with each run counting as an event. Data can be retrieved in three ways, via a run or date range, or by providing a run-list. The last approach is the most useful when collecting data about the selected neutrino runs for example.

For each run the RLOG bank is first retrieved from the database an placed in memory to be treated as an event. The RLAI tree is then retrieved and attached to the RLOG banks to make it appear as though the two were one entity. The final step is to set the current context to the centre of the run. This allows banks to be read in from the database that are valid for the current run, in particular to allow trends to analysed.

Appendix C

Data Selection

Before any analysis of neutrino data can take place the data must undergo a selection process. The goal of this selection process is to provide data to the next stage of the analysis that satisfies a few basic criteria. These criteria are selected to provide a basic level of data quality so that any downstream analysis does not have to cope with certain situations.

The criteria that are used in data selection look specifically for problems in the hardware of the detector that have been deemed to produce suspect data. Data selection can also remove data that has been affected by some software problem, such as a problem in the builder.

In SNO the data is quite naturally split into chunks of varying size, called runs. As a first step in data selection a run selection process can be implemented. A specific list of requirements can then be made to select the run based on all available information without looking at the data itself. Data about runs can be found in a number of different places in the data stream and each is checked to provide a provisional run list. This list is then checked against the shift reports to find any further problems that could compromise the data. The detailed list of run criteria is as follows.

- The run should be a neutrino run.

  This is checked in the run header (RHDR) and it requires that specific bits be either set or unset. The RHDR bank contains a number of bits set at the begining of the run, describing what type of run the current run is.

  This allows calibration runs of any sort, to be removed without the need for a long task of shift report checking to go ahead. In a similar way experimental, maintenence and supernova runs are also be removed at this stage. The system also allows the quick removal of runs with problems in the magnetic compensation coils. Runs which have the unusual circumstances (UC) bit set are passed to an alternative data stream for consideration about the inclusion of the run, rather than the exclusion.

- The trigger settings should be as for a normal neutrino run.

  This is a check of the trigger bank (written to the data stream at the start of the run) to make sure that the run looks like a neutrino run. This is a minimum requirement for which trigger bits are masked in as well as a maximum threshold for the NHIT100med trigger at twenty hits. The NHIT100med, NHIT100hi, NHIT20, ESUMHi, PulseGT and OWLn triggers should be masked in.

- All 19 electronic crates should be online

  This requirement is check via the DQxx banks[1] and helps to ensure that a sufficient number of channels are online and distributed randomly throughout the detector.

- No more then 16X32 channels offline.

  Another DQxx check, carried out to ensure that a sufficient number of channels are online. This check is complementary with the electronic crate check.

- A sufficient number of nonstandard channels are on line.

  The NECK, OWL and BUTT tubes are used extensively instrumental cuts. A requirement that a sufficient number are on, ensures that the cuts that rely on these tubes will continue to operate. The requirement can be broken down as follows.

  - All four NECK tubes online.
  - More than sixty OWLs + BUTTs online, with an additional requirement that there and more than nine BUTTs and forty-eight OWLs online at any one time.

  As before this information can be gleaned from the DQxx banks.

- There be no more than 6 channels that are potential blind flashers. In this case a potential blind flasher is a channel that has high voltage on but has its sequencers disabled, so will not be read out. This is again check via the DQxx banks. Runs should not occur that fail this criterion , but due to operator error they have been known to occur just after maintenance periods.

[1] DQxx banks are written out by the DAQ at the start of the run and record the electronic status of each channel at that time.

- The run should be greater than 30 minutes in length

  There is usually a very good reason why a run is shorter than 30 minutes in length. Most of these reasons result for a condition that would cause the run to be undesirable. UC bit runs have a cut imposed at 1 hour to reduce the size of the list.

- The imposed deadtime should be acceptable.

  If one of the burst cuts removes more than 25% of the run, then the remainder of the data is suspect.

- If an orphan burst occurs it should be understood.

  Some orphan burst have occurred which are not understood by the DAQ group. The data is suspect if the reason for the orphan burst is not understood.

- Shift report verification.

  The final step in run selection is shift report[2] verification. The shift report is read for each of the prospective neutrino runs and if one of a number of criteria are met then the run is discarded. These same criteria are used to rescue runs where the UC bit has been set. These criteria include the following:

  - The run had one of the bits such as the UC bit set incorrectly in the run header.
  - The run contained a D_2O assay and is full of light.
  - There was a hardware failure during the run and part of the detector was disabled for part of the run. This has been known to happen during thunderstorms for example, or during an episode of WEB. If the error is noticed immediately and a new run started then the run is deemed OK, however, this type of problem can go unnoticed for some time, and in these situations the run is cut. A more sophisticated procedure than run selection can be employed to retrieve this data.
  - The report notes something unusual happening with the compensation coils.
  - The run occurred just after extensive hardware changes and there are not yet any channel calibration constants.
  - The is a calibration source deployed in the detector. On occasion the source has been left in the glove box [3], if the source has been left here and it is not a radio-active source , such as the acrylic source, then the runs are used.

[2]The shift report is a detailed log of all events during a shift, filled in by the shift operator during the shift.

[3]The glovebox is a region of the detector above the neck at z=+1300cm.

– The humidity or temperature in the DECK is much higher than normal due to problems with the chiller.

By carrying out this process a list of runs is produced known as the selected runs. The data from these runs, after the application of the instrumental cuts is known as neutrino data, and is used for all studies of the neutrino signals.

Glossary

ADC Analogue to Digital Converter.

AMB Analogue Measurement Board. The basis of one of the instrumental cuts.

AV The Acrylic Vessel.

Berkeley Blob A patch of enhanced radioactivity on the acrylic vessel.

BUTT Berkeley Underwater Tube Testing PMT.

Builder A program that combines Event Headers with PMT Bundles to form events.

Contamination The residual background content of a data set after the application of cuts to remove that background.

CC The Charged Current reaction.

Crosstalk The creation of a signal of a PMT by the signal in a neighbouring PMT.

DAQ Data AcQuisition system.

ECA Electronics Calibration of PMT signals.

ES The Elastic Scattering reaction.

ESUM A trigger based upon the charge deposited on inward-looking PMTs.

Event Header Data output by the detector containing information about an event.

FGC Flasher Geometry Cut. An instrumental cut.

Fitter A process that uses PMT time and position information to reconstruct the time and position of the event.

FTS Fitterless Timespread Cut. An instrumental cut.

GTID Global Trigger ID. A unique number for an event in a run.

HCA History Calibration of PMT signals.

ITC In-Time Channel Cut. An instrumental cut.

ITR In-Time Ratio. A high level cut.

Just-So^2 A region of MSW parameter space that can explain the solar neutrino problem.

LMA Large Mixing Angle solution. A region of MSW parameter space that can explain the solar neutrino problem.

LOW A region of MSW parameter space that can explain the solar neutrino problem.

MSW effect. Mikheyev, Smirnov and Wolfenstein effect. The theory of the modification of neutrino oscillations by the presence of matter and its application to the solar neutrino problem.

NC The Neutral Current reaction.

NHITs Th number of hit inward looking PMTs. An approximate measure of the energy of an event.

NHIT trigger A trigger based upon the number of hit PMTs inside a given time window.

Orphan A PMT bundle that has not been correctly combined with an event header.

OWL OutWard Looking PMT.

PCA PMT calibration.

PulseGT The unbiased trigger. The state of the detector is recorded every 200ms by this trigger.

PMT PhotoMultiplier Tube.

PMT bundle Data output from the detector containing the details of an individual PMT hit.

pp chain The pp chain is the dominant series of nuclear reactions by which energy is generated in the sun.

PSUP The PMT SUPport structure.

QHL High gain, long integration time charge. A PMT measurement.

QHS High gain, short integration time charge.A PMT measurement.

QLX Low gain charge. The integration time is set by the DAQ system. A PMT measurement.

QvT Charge vs Time. An instrumental cut.

Railed In SNO a measurement is said to be railed if there are insufficient bits to store the digital result of a measurement.

RLDT RunLogger Dead Times. A bank that stores livetime information, part of the runlogger system.

RSP Response Processor. The principle energy calibration in SNO.

Run A conveniently sized period of data taking.

Runlogger A system used for storing information about individual runs.

Sacrifice The removal of signal by cuts designed to remove background.

SMA Small Mixing Angle solution. A region of MSW parameter space that can explain the solar neutrino problem.

SNO The Sudbury Neutrino Observatory.

SNOMAN SNO Monte carlo and ANalysis. The principle Monte Carlo and analysis package for the SNO experiment.

TAC Time to Amplitude Converter. Also the name of the uncalibrated time measurement for a PMT.

Bibliography

[1] C.L. Cowan et al. Detection of the free neutrino. *Science*, 124:103, 1956.

[2] B.T. Cleveland et al. Measurement of the solar electron neutrino flux with the Homestake chlorine detector. *The Astrophysical Journal*, 496:505, 1998.

[3] LEP. Results from lep concerning the width of the Z^0 from L3, ALEPH, OPAL and DELHI. *Physics Letters B*, pages 509, 519, 530 and 539 respectivly, 1989.

[4] J. Bonn et al. The Mainz neutrino mass experiment. In R.W. Ollerhead J. Law and J.J Simpson, editors, *Neutrino 2000, Proceedings of the XIXth International Conference on Neutrino Physics and Astrophysics*, volume 91 of *Nuclear Physics B (Proc. Suppl.)*, page 273, 2001.

[5] V.M. Lobashev et al. Direct search for neutrino mass and anomaly in the tritium beta-spectrum: Status of the "Troisk Neutrino Mass" experiment. In R.W. Ollerhead J. Law and J.J Simpson, editors, *Neutrino 2000, Proceedings of the XIXth International Conference on Neutrino Physics and Astrophysics*, volume 91 of *Nuclear Physics B (Proc. Suppl.)*, page 280, 2001.

[6] D.E. Groom et al. Review of particle physics. *The European Physical Journal*, C15:1, 2000.

[7] J.M. Roney. Review of tau neutrino mass. In R.W. Ollerhead J. Law and J.J Simpson, editors, *Neutrino 2000, Proceedings of the XIXth International Conference on Neutrino Physics and Astrophysics*, volume 91 of *Nuclear Physics B (Proc. Suppl.)*, page 287, 2001.

[8] H. Ejiri. Double beta decay and neutrinos. In R.W. Ollerhead J. Law and J.J Simpson, editors, *Neutrino 2000, Proceedings of the XIXth International Conference on Neutrino Physics and Astrophysics*, volume 91 of *Nuclear Physics B (Proc. Suppl.)*, page 255, 2001.

[9] B. Pontecorvo. Neutrino experiments and the problem of conservation of leptonic charge. *Soviet Physics JETP*, 26(5):984, 1968.

[10] V. Gribov and B. Pontecorvo. Neutrino astronomy and lepton charge. *Physics Letters B*, 28(7):493, 1969.

[11] L. Wolfenstein. Neutrino oscillations in matter. *Physical Review D*, 17(9):2369, 1978.

[12] S. Mikheyev and A. Smirnov. Resonance enhancement of oscillations in matter and solar neutrino spectroscopy. *Soviet Journal of Nuclear Physics*, 42:913, 1985.

[13] S. Mikheyev and A. Smirnov. Resonant amplifications of ν-oscillations in matter and solar-neutrino spectroscopy. *Il Nuovo Cimento*, 9C:17, 1986.

[14] J.N. Bachall. *Neutrino Astrophysics.* Cambridge University Press, 1989.

[15] J.N. Bachall P.I. Krastev A.Y. Smirnov. Solar neutrinos: global analysis and implications for SNO. hep-ph/0103179, 2001.

[16] John N. Bahcall M. H. Pinsonneault and Sarbani Basu. Solar models: current epoch and time dependences, neutrinos, and helioseismological properties. *The Astrophysical Journal*, 555:990–1012, 2001. astro-ph/0010346.

[17] J.N. Bachall. http:://apollo.sns.ias.edu/ jnb/.

[18] Y. Fukanda et al. Solar neutrino data covering solar cycle 22. *Physical Review Letters*, 77(1683), 1996.

[19] S. Fukuda et al. Solar ^{8}B and *hep* neutrino measurements from 12258 days of Super-Kamiokande data. *Physical Review Letters*, 86, 2001. hep-ex/0103032.

[20] J.N. Abdurashitov et al. Measurement of the solar neutrino capture rate with gallium metal. *Physical Review C*, 60(055801), 1999.

[21] W. Hampel et al. GALLEX solar neutrino observations: Results from GALLEX IV. *Physics Letters B*, 447:127, 1999.

[22] W. Altmann et al. GNO solar neutrino observations: Results from GNO I. *Physics Letters B*, 490:16, 2000.

[23] E. Bellotti. First results from GNO. In R.W. Ollerhead J. Law and J.J Simpson, editors, *Neutrino 2000, Proceedings of the XIXth International Conference on Neutrino Physics and Astrophysics*, volume 91 of *Nuclear Physics B (Proc. Suppl.)*, page 44, 2001.

[24] V.N. Gavrin. Solar neutrino results from SAGE. In R.W. Ollerhead J. Law and J.J Simpson, editors, *Neutrino 2000, Proceedings of the XIXth International Conference on Neutrino Physics and Astrophysics*, volume 91 of *Nuclear Physics B (Proc. Suppl.)*, page 36, 2001.

[25] N. Hata P. Langacker. Solutions to the solar neutrino anomaly. *Physical Review D*, 56:6107, 1997.

[26] M. Apollonio et al. Limits on neutrino oscillations from the CHOOZ experiment. *Physics Letters B*, 466:415, 1999.

[27] G.L. Fogli et al. Model-dependant and independant implications of the first Sudbury Neutrino Observatory results. *Physical Review D*, 64, 2001. hep-ph/0106247.

[28] Herbert H. Chen. Direct approach to resolve the solar-neutrino problem. *Physical Review Letters*, 55(14):1534–1536, September 1985.

[29] S. Nakamura T.Sato V. Gudkov and K. Kubodera. Neutrino reactions on deuteron. *Physical Review C*, 63(034617), 2001.

[30] M. Butler J.W. Chen and X. Kong. Neutrino-deuteron scattering in effective field theory at next-to-next-to leading order. *Physical Review C*, 63(035501), 2001.

[31] F. Duncan. Calculation of the solar neutrino yield in sno (updated). Technical report, Queen's University, 2001.

[32] J. Bodger et al. The Sudbury Neutrino Obsevatory. *Nuclear Instruments and Methods in Physics Research*, A(449), 2000.

[33] M.R Dragowsky. Pure D_2O neutron calibration results. Technical report, Los Alamos National Laboratory, 2001.

[34] M. Moorhead. *Reflectors in Čerenkov Detectors*. PhD thesis, Univerity of Oxford, 1992.

[35] E.W. Beier et al. Performance of the electronics for the Sudbury Neutrino Observatory. Technical report, University of Pennsylvania, 1998.

[36] D.F. Cowen et al. The Sudbury Neutrino Observatory electronics chain. Technical report, The Sudbury Neutrino Observatory, 2000.

[37] J. Klein et al. The SNO trigger system. Technical report, University of Pennsylvania, 1997.

[38] M. Neubauer. *Evidence for Electron Neutrino flavour change through Measurement of the ^8B Solar Neutrino Flavour Content at the Sudbury Neutrino Observatory.* PhD thesis, University of Pennsylvania, 2001.

[39] H. Heron. *Techniques to Measure the NC Background in the SNO Experiment.* PhD thesis, Universiry of Oxford, 1998.

[40] A.B. McDonald. First neutrino observations from the Sudbury Neutrino Observatory. In R.W. Ollerhead J. Law and J.J Simpson, editors, *Neutrino 2000, Proceedings of the XIXth International Conference on Neutrino Physics and Astrophysics*, volume 91 of *Nuclear Physics B (Proc. Suppl.)*, page 21, 2001.

[41] The SNO Collaboration. *The SNOMAN User's Manual*, 4.01 edition.

[42] The SNO Collaboration. *The SNOMAN Programmer's Manual*, 4.01 edition.

[43] The SNO Collaboration. *The SNOMAN Companion*, 4.01 edition.

[44] Zebra. CERN Program Library Long Writeups Q100/Q101, 1995.

[45] Hideo Hirayama Walter R. Nelson and David W.O. Rogers. The EGS4 code system. SLAC report 265., 1985.

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

[47] S.J. Brice. *Monte Carlo and Analysis Techniques for the Sudbury Neutrino Observatory.* PhD thesis, University of Oxford, 1996.

[48] M.J. Lyon. *Neutron Transport in the Sudbury Neutrino Detector.* PhD thesis, University of Oxford, 1996.

[49] M.D. Lay. *Creation and Detection of Čerenkov Light in the Sudbury Neutrino Observatory.* PhD thesis, University of Oxford, 1994.

[50] S. Biller et al. SNO electronic calibration constants. Technical report, The Sudbury Neutrino Observatory, 2001.

[51] G. McGregor. Private Communication.

[52] J.R.N. Cameron. *The Photomultipler Tube Calibration of the Sudbury Neutrino Observatory.* PhD thesis, University of Oxford, 2001.

[53] B.A. Moffat. *The Optical Calibration of the Sudbury Neutrino Observatory*. PhD thesis, Queen's University, 2001.

[54] S. Biller et al. Optics6b. Technical report, The Sudbury Neutrino Observatory, 2001.

[55] M.G. Boulay. *Direct Evidence for Weak Flavour Mixing with the Sudbury Neutrino Observatory*. PhD thesis, Queen's Univesity, 2001.

[56] M.G. Boulay. An energy calibration of SNO: The RSP processor. Technical report, Queen's University, 2001.

[57] The SNO Collaboration. Analysis for the first solar neutrino paper. Technical report, The Sudbury Neutrino Observatory, 2001.

[58] Q. R. Ahmed et al. Measurement of the charged current interactions produced by ^{8}B solar neutrinos at the Sudbury Neutrino Observatory. *Physical Review Letters*, 87, 2001. nucl-ex/0106015.

[59] J. Cameron et al. The gain and pain: A study of detector response and instrumental backgrounds after the high voltage raise. Technical report, The Sudbury Neutrino Observatory, 1999.

[60] J. Orrel K. Shaffer and J. Wilkerson. The Junk cut. Technical report, University of Washington, 2000.

[61] J. Klein. Data Cleaning: An Analog Measurement Board cut. Technical report, University of Pennsylvania, 1999.

[62] J. Klein. Addendum to Analog Measurement Board cut. Technical report, University of Pennsylvania, 2000.

[63] P. Wittich. Data Cleaning: The Cluster cut. Technical report, University of Pennsylvania, 1999.

[64] J. Klein and V. Rusu. Data Cleaning: The QCluster cut. Technical report, University of Pennsylvania, 2000.

[65] X. Chen. Flasher Geometry cut. Technical report, Lawrence Berkeley Laboratory, 2000.

[66] J. Klein. Data Cleaning: The QvT cut. Technical report, University of Pennsylvania, 1999.

[67] J. Klein. Addendum to the QvT cut: Monitoring the sacrifice. Technical report, University of Pennsylvania, 2001.

[68] M. Neubauer. Data Cleaning: The Fitterless Timespread cut. Technical report, University of Pennsylvania, 1999.

[69] M. Neubauer. Addendum to the FTS cut report. Technical report, University of Pennsylvania, 2000.

[70] R. Van de Water. Data Cleaning: The Q/NHIT cut. Technical report, University of Pennsylvania, 2000.

[71] M. Neubauer. Data Cleaning: The In-Time-Channel (ITC) cut. Technical report, University of Pennsylvania, 2000.

[72] X. Chen. *Monte Carlo Simulation and Analysis of Backgrounds in the Sudbury Neutrino Observatory*. PhD thesis, University of Oxford, 1997.

[73] A.S. Hamer. *Energy Calibration of SNO for Measurement of the Charged-Current Reaction*. PhD thesis, Queen's University, 1999.

[74] N. Tagg. *The ^{8}Li Calibration Source and Through-Going Muon Analysis in the Sudbury Neutrino Observatory*. PhD thesis, University of Guelph, 2001.

[75] N. McCauley. Sacrifice update to include recent ^{8}Li data. Technical report, University of Oxford, 2001.

[76] R.J. Ford. *Calibration of SNO for the Detection of ^{8}B Neutrinos*. PhD thesis, Queen's University, 1998.

[77] A.W.P. Poon. *Energy Calibration of the Sudbury Neutrino Observatory using Monoenergetic Gamma-Ray Sources*. PhD thesis, The University of British Columbia, 1998.

[78] S. Brice et al. Sacrifice and contamination in SNO Data Cleaning cuts. Technical report, The Sudbury Neutrino Observatory, 2000.

[79] J. Klein. Time dependant systematic uncertainties of AMB cut. Technical report, University of Pennsylvania, 2001.

[80] M. Chen. Bifrucated analysis 1.0. Technical report, Queen's University, 2000.

[81] M.K. Ochi. *Applied Probability and Stochastic Processes*. Wiley, 1990.

[82] N. Tagg. Data Cleaning: The Muon tag. Technical report, University of Guelph, 1999.

[83] A.D. Marino. Muon-induced spallation products in the SNO detector. Technical report, Lawrence Berkeley Laboratory, 2000.

[84] K. Schaffer et al. The Orphan document. Technical report, The Sudbury Neutrino Observatory, 2001.

[85] N.S.Oblath and A.W. Poon. Notes on the anomalous event cluster on the acrylic vessel. Technical report, Lawrence Berkeley Laboratory, 2000.

[86] A.W. Poon. Calibrating the radioactivity of the AV hot spot with the weak Th source I. Technical report, Lawrence Berkeley Laboratory, 2001.

[87] S. Oser. Private Communication.