The Photomultiplier Tube Calibration of the Sudbury Neutrino Observatory

James R. N. Cameron
Wolfson College, Oxford

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

The Photomultiplier Tube Calibration of the Sudbury Neutrino Observatory

James R. N. Cameron
Wolfson College, Oxford

Thesis submitted for the degree of Doctor of Philosophy
Trinity Term 2001

Abstract

The Sudbury Neutrino Observatory (SNO) is a solar neutrino experiment designed to search for neutrino flavour oscillations. It is located 2040 m underground in the INCO, Ltd. Creighton mine near Sudbury, Ontario. The detector consists of a kilotonne of heavy water (D_2O) viewed by approximately 9,500 Photomultiplier Tubes (PMTs). The PMTs detect Čerenkov photons produced following the interactions of ^8B solar neutrinos with the D_2O. These neutrinos can interact with D_2O in three different ways. This allows the electron neutrino flux and the total flux of all flavours of neutrino to be measured. In order to distinguish the different types of reaction, it is essential to reconstruct where the interactions took place in the detector. The reconstruction critically depends on an accurate knowledge of the relative times at which the PMTs detect the Čerenkov photons. The first results from SNO provide strong evidence for neutrino oscillations and have further constrained the possible values of the square of the neutrino mass differences and the neutrino mixing angles.

The topic of this thesis is the Photomultiplier Tube Calibration. The SNO electronics record a time and an integrated charge for each PMT that is triggered by a photon. The recorded time has an unknown offset and the time resolution of the PMTs is degraded by an effect arising from variations in the pulse sizes. The PMT calibration removes the time offsets, giving the relative timing of the PMTs, and restores the time resolution to the extent possible. The PMT calibration also converts the integrated charge into the same units for all PMTs. This thesis describes how the PMT calibration is implemented and presents the results of investigations into its performance. The PMT calibration is found to perform well and is crucial to the analysis of SNO data.

Acknowledgements

There are many people whom I wish to thank for their help in completing this thesis. Foremost is my advisor, Dave Wark. He has been a constant source of sound advice and guidance. His knowledge, enthusiasm and good humour has made working with him, not only a privilege, but a pleasure. My parents have provided constant encouragement throughout. Without their support, over the last year especially, the completion of this thesis would have remained a distant prospect.

Bill Frati pioneered much of the early work on the PMT calibration. I have also benefited from the advice of Steve Biller. Nick West has provided a constant model of good programming practice and near saintly patience. I thank Nick Jelley and George Doucas for their good nature and adding to the pleasant working environment. Kate Frame has given much encouragement and has already joined the select few who will read my thesis. Beverly Roger has been an enormous help in organising my various travel arrangements.

Gordon McGregor and Neil McCauley have both made working on SNO much more entertaining and enjoyable. Gordon is always full of ideas and we have shared many interests. Neil is always ready to provide helpful tips. I would also like to thank David Bailey who frequently lifted my spirits during the long writing process.

Derek Sweezey, John Roberts, Ha Nguyen and Araz Hamian all made life in Sudbury much more enjoyable. All the members of the Penn group have proved willing to explain the operation of the electronics and to consume beer and wings. I also thank Rich Helmer for encouraging me to spend my time in Sudbury wisely.

Contents

1 Solar Neutrinos **1**
- 1.1 The Neutrino . . . 1
- 1.2 The Standard Solar Model . . . 2
- 1.3 Solar Neutrino Experiments . . . 4
  - 1.3.1 Radiochemical Experiments . . . 5
  - 1.3.2 Water Čerenkov Experiments . . . 8
- 1.4 The Solar Neutrino Problem . . . 11
- 1.5 Neutrino Oscillations . . . 14
  - 1.5.1 Vacuum Oscillations . . . 14
  - 1.5.2 Matter Enhanced Oscillations . . . 16
  - 1.5.3 Possible Solutions . . . 17
- 1.6 The Sudbury Neutrino Observatory . . . 19
  - 1.6.1 The SNO Detector . . . 19
  - 1.6.2 Neutrino Interactions in SNO . . . 19
  - 1.6.3 Phases of the Experiment . . . 22
  - 1.6.4 Backgrounds . . . 23
  - 1.6.5 First Results from SNO . . . 24
- 1.7 The Layout of this Thesis . . . 28

2 Background Information **29**
- 2.1 The SNO Photomultiplier Tubes . . . 29
- 2.2 The SNO Electronics . . . 35
- 2.3 Electronics Calibration . . . 40
- 2.4 ECA Validation . . . 41
- 2.5 The Laserball . . . 42
- 2.6 The ^{16}N calibration source . . . 44
- 2.7 The ^{8}Li calibration source . . . 45
- 2.8 SNOMAN . . . 46

3 Introduction to the PCA **47**
- 3.1 The Walk Calibration . . . 47
- 3.2 The Gain Calibration . . . 50

4 The Implementation of the PCA **55**
4.1 The PCA Framework 55
4.2 The PCA Laserball Run 56
4.3 The Walk Calibration 57
4.4 The Gain Calibration 65
4.5 The Extended Walk Calibration 70
4.6 The Calibrator 75
4.6.1 Applying the Walk Calibration 76
4.6.2 Applying the Gain Calibration 78
4.7 Bad Pedestals 79
4.8 The Bad Channel Cuts 83
4.8.1 Setting the Level of the Bad Channel Cuts 87

5 The Performance of the PCA **94**
5.1 Testing the Walk Calibration 94
5.2 The Laserball Path Correction 108
5.2.1 Testing the Path Correction in the Monte Carlo 108
5.2.2 Testing the Path Correction with Data 110
5.2.3 The Reflection Fitter 116
5.3 The Effect of Occupancy 128
5.3.1 The Effect of Occupancy on the Timing 129
5.3.2 The Effect of Occupancy on the Gain 135
5.3.3 Rates and Loading 135
5.4 The Stability of the Constants 136

6 The High χ^2 Problem **145**
6.1 Overview 145
6.2 Global Time Since Last Hit Correction 147
6.3 Rate Correction 151
6.4 To Ping or Not To Twiddle. 153
6.5 Current Status 160

7 Conclusions **162**
7.1 ^{8}Li Reconstruction 162
7.2 Neutrino Reconstruction 166
7.3 Final Conclusions 171

Chapter 1

Solar Neutrinos

1.1 The Neutrino

The neutrino was first postulated by Pauli in 1930 to explain why the energy spectrum of electrons from nuclear beta decay was continuous, in contrast to the discrete energy spectra of alpha particles and gamma rays. It is now established that there are three flavours of neutrino, the electron, muon and tau neutrino (ν_e, ν_μ, ν_τ). Neutrinos are included in the Standard Model of particle physics, a very successful model of the fundamental particles and their interactions. Despite the success of the Standard Model it is incomplete and has a large number of free parameters.

Neutrinos are the only neutral fundamental fermions and only interact very weakly with ordinary matter. Their neutrality raises the possibility that they are their own antiparticles, referred to as Majorana neutrinos. Ordinary fermions are called Dirac particles, although both Dirac and Majorana neutrinos satisfy the Dirac equation. The neutrino masses are unknown but upper limits have been set from the kinematics of decays in which they are produced: $m_{\nu_e} < 2.2\,\mathrm{eV}$ [1][2], $m_{\nu_\mu} < 0.19\,\mathrm{MeV}$ [3] and $m_{\nu_\tau} < 15.5\,\mathrm{MeV}$ [4]. If neutrinos have mass then it is possible that their weak eigenstates, the states of definite flavour, are mixtures of their mass eigenstates, the eigenstates of the free particle Hamiltonian. Establishing these properties is an exciting area of research because of the limits they put on viable theories of particle physics as well as improving our understanding of the role of neutrinos in cosmology and astrophysics.

If neutrinos have mass, and the neutrino flavour eigenstates are mixtures of the mass eigenstates then it is possible for flavour oscillations to occur. Flavour oscillations cause the probability of detecting a given flavour of neutrino to vary with distance from the source. Neutrino oscillations will be discussed further in section 1.5. Experiments which search for neutrino oscillations provide a way to probe the neutrino mass differences and mixing angles. Neutrino oscillation experiments utilise four basic sources of neutrinos: reactors, accelerators, neutrinos produced in the upper atmosphere by cosmic rays, and neutrinos produced by the Sun. Solar neutrinos are particularly sensitive to small mass differences between the neutrinos because of the long baselines involved. The Sudbury Neutrino Observatory is a solar neutrino experiment and provides a very powerful way to search for neutrino oscillations and to test the Standard Solar Model.

1.2 The Standard Solar Model

The Sun is a main sequence star which is thought to have been steadily burning for 4.6 billion years. The only known processes capable of powering the Sun over such a long period are nuclear reactions. The overall reaction burns hydrogen nuclei to helium nuclei and can be represented as

$$4p \rightarrow \alpha + 2e^+ + 2\nu_e$$

which releases 26.7 MeV of energy, including the energy from positron annihilation. The two neutrinos carry away an average of about 0.6 MeV. 98.5% of the energy generated in the Sun is generated as a result of the pp chain, which is shown in figure 1.1. The remaining 1.5% of the energy is provided by another set of reactions called the CNO cycle.

A standard solar model (SSM) [5] is a model of the Sun which uses widely accepted physics and initial conditions. The Sun is modelled as a body in hydrostatic

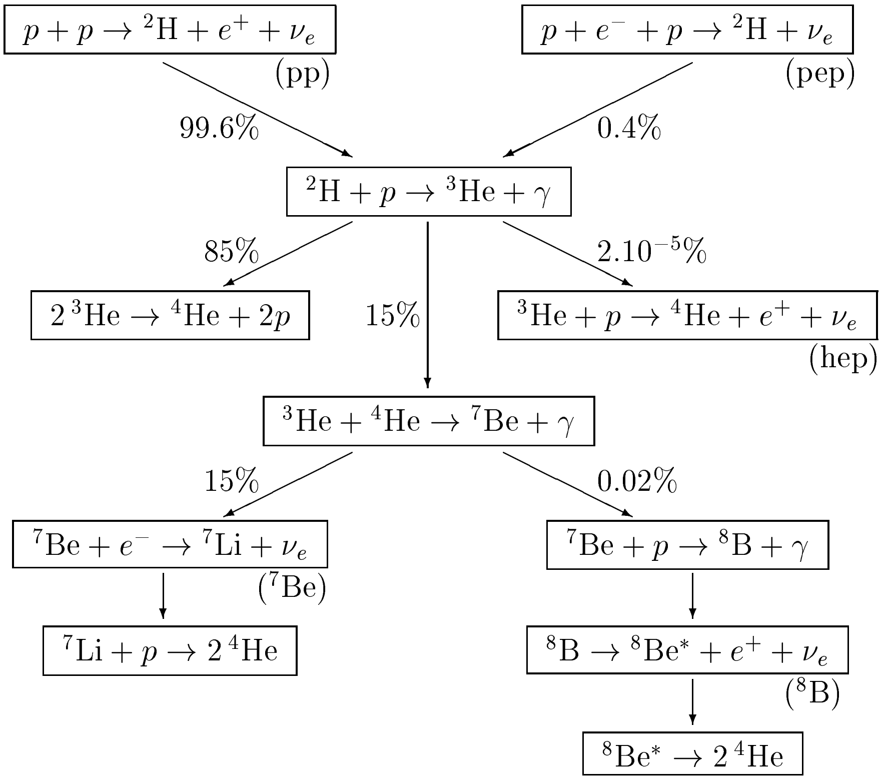

Figure 1.1: The pp chain [5].

equilibrium, with the radiative and particle pressures exactly balancing the force of gravity at all radii. This condition must be true as the time-scale for the gravitational collapse of the Sun is only an hour. Deep in the interior of the Sun the outward energy flux is mainly determined by the radiative opacity, while near the solar surface the dominant mode of energy transport is by convection. The time scale for a photon to randomly walk out of the solar interior is 10^4 years.

The evolution of the Sun is calculated by numerically solving a set of coupled differential equations. The current state of the Sun is found by quasistatically evolving the solar model from the time when the Sun first joined the main sequence and started hydrogen burning. The initial chemical composition is taken to be homogeneous. The mass fractions of hydrogen, helium and heavy elements (all elements heavier than helium) are usually denoted by X, Y and Z respectively. The initial abundance of elements heavier that helium (Z/X) is taken from meteorite data and agrees with that from the present day solar surface, which is isolated from the nuclear reactions in the core. As the model is evolved the chemical composition gradually changes due to nuclear burning and also from the small effects of gravity and the temperature gradient. Successive models accommodate these composition changes and the evolution of parameters such of the radius and luminosity. The initial conditions, like X, are varied until a model is found which correctly reproduces the current radius and luminosity of the Sun at the present age of 4.6 billion years. For the Bahcall Pinsonneault 2000 (BP2000) SSM [6] the hydrogen mass fraction starts at 0.708 and has reduced to 0.340 in the solar core by the current epoch. The neutrino flux predicted by the BP2000 SSM is shown in figure 1.2.

Helioseismology can also provide information about the conditions deep within the Sun [5][6]. The Sun acts like an spherical acoustic cavity and pressure mode oscillations can be measured on its surface. These oscillations can be used to deduce the sound speed and density as a function of radius. There is excellent agreement between the helioseismological measurements and the SSM calculations giving credence to the neutrino flux predictions.

1.3 Solar Neutrino Experiments

Solar neutrino experiments are interesting because they provide a window on the reactions happening deep in the core of the Sun. The neutrinos have no trouble escaping the Sun but, similarly, only very rarely interact in a detector. To get a measurable signal, detectors must have large target volumes and operate for long periods of time. Because of the feeble signal the suppression of backgrounds is particularly challenging. There are two basic types of solar neutrino experiment which are currently operating: radiochemical experiments and water Čerenkov experiments.

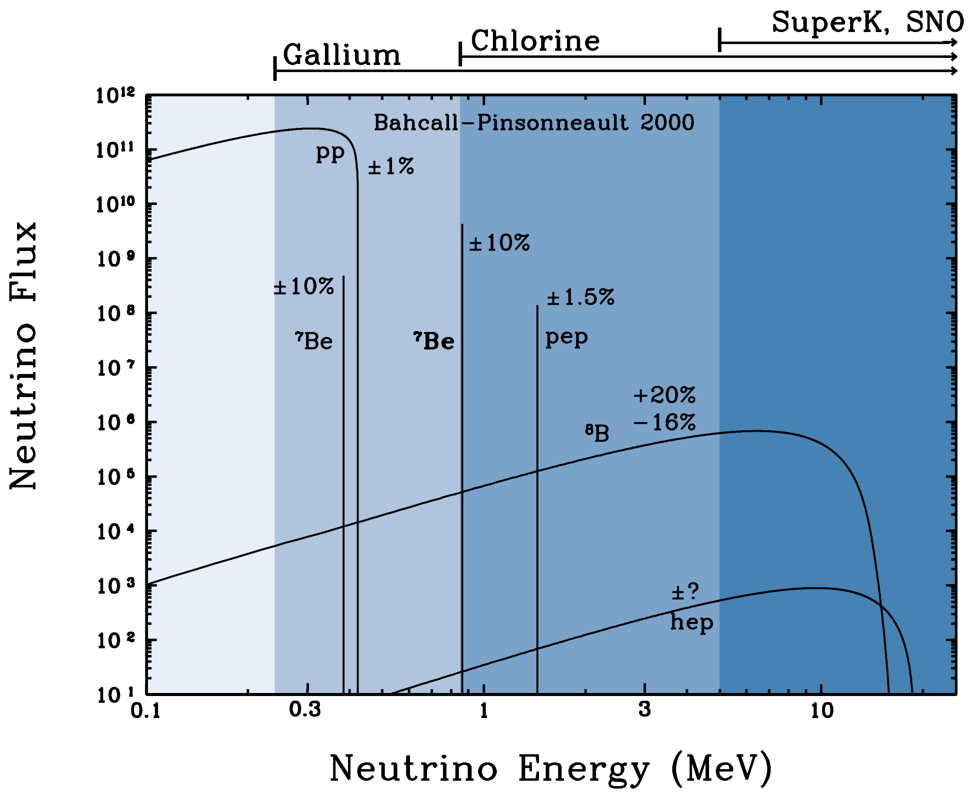

Figure 1.2: The energy spectrum of solar neutrinos as predicted by the BP2000 standard solar model with 1σ uncertainties. The fluxes from continuum sources are in units of number per cm^2 per second per MeV at one astronomical unit. The line fluxes are in units of number per cm^2 per second. The energy thresholds of solar neutrino experiments are also shown. Taken from [7].

1.3.1 Radiochemical Experiments

A radiochemical detector consists of a volume of material containing a large number of the chosen target nucleus. Neutrinos interact with the target nuclei to produce a different element. The target material is exposed for a period of many weeks and then the product nuclei are extracted and counted. The number of product nuclei gives a measure of the neutrino flux. The measured neutrino flux is therefore averaged over many weeks and contains no information on variations over shorter time scales. These experiments provide no directional or spectral information, other than that from the threshold of the reaction. In compensation the thresholds of radiochemical detectors are lower than that of water Čerenkov detectors. In the following descriptions of the radiochemical experiments the neutrino fluxes are presented in units of SNU, the Solar Neutrino Unit. One SNU is equal to 10^{-36} captures per target atom per second.

The Chlorine Experiment

The ^{37}Cl experiment was the first solar neutrino experiment. The reaction studied was

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

which has a threshold of 0.814 MeV and is therefore sensitive to all the types of solar neutrinos except pp. ^{37}Ar has a 4.98 MeV excited state which makes the experiment most sensitive to ^{8}B neutrinos. The target consisted of a tank of 615 tonnes of liquid C_2Cl_4 (tetrachloroethylene, a dry cleaning fluid), 4850 feet underground (equivalent to 4200 metres of water) in the Homestake Gold Mine, South Dakota. The experiment was deep underground to reduce the flux of cosmic rays, which can also react with ^{37}Cl to produce background ^{37}Ar.

The first step in the measurement cycle was to stir in a small, but known, amount of ^{36}Ar or ^{38}Ar carrier into the tank. The tank was then left for a period of 1–3 months. The argon, including any ^{37}Ar produced by solar neutrino interactions, was then extracted by bubbling helium through the tank and separating out the argon. The volume of extracted argon provides a measure of the extraction efficiency which was demonstrated to be more than 90%. ^{37}Ar decays by electron capture back to ^{37}Cl with a half-life of 35 days. The collected argon was put into a proportional counter to count the Auger electrons caused by the decay. The counters were studied for about a year to get a good measurement of the background. The experiment first started taking data in 1967 but could not measure the solar neutrino flux above the counter background. This situation changed in 1970 when an analysis of the energy and rise time characteristics of the counter pulses was introduced.

An analysis of 108 extractions made between April 1970 and May 1994 [8] gives a capture rate of

$$2.56 \pm 0.16(\text{stat.}) \pm 0.16(\text{sys.})\ \text{SNU}$$

The BP2000 standard solar model prediction [6] is

$$7.6^{+1.3}_{-1.1}\ \text{SNU}$$

The Gallium Experiments

Three gallium experiments have been built: SAGE, GALLEX and GNO. The basic method is very similar to the chlorine experiment. The reaction studied is

$$\nu_e + {}^{71}\text{Ga} \rightarrow e^- + {}^{71}\text{Ge}$$

which has a threshold of 0.2332 MeV and is therefore sensitive to all types of solar neutrino. The fact that the gallium experiments are sensitive to pp neutrinos is very significant because the pp flux is very tightly constrained by the luminosity of the Sun, whereas ^{8}B neutrinos are very sensitive to the core temperature. The gallium is exposed for a period of 2–4 weeks and then the germanium is extracted. ^{71}Ge decays back to ^{71}Ga by electron capture with a half-life of 11.43 days. Like the chlorine experiment the germanium decays result in Auger electrons which are detected in proportional counters.

SAGE is the Soviet American Gallium Experiment and is located in an adit under Mount Andyrchi in the Baksan Valley in the Caucasus Mountains. The mountain overhead provides shielding equivalent to 4700 metres of water. The gallium target is liquid metallic Ga and was increased from 30 tons to 57 tons in mid-1991. Using the metal form of gallium eases the problem of radioactive impurities. It also minimises the target volume which increases the reaction rate when calibrating with ^{51}Cr. The extraction process for metallic gallium is relatively complex and results in GeH_4 which is then counted. The overall result from 70 runs taken over the period January 1990 to October 1999 [9][10] is

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

The BP2000 standard solar model prediction for the gallium experiments [6] is

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

GALLEX (GALLium EXPeriment) is located in the Gran Sasso Laboratory in Italy. The target is 30.3 tons of ^{71}Ga in 101 tons of aqueous gallium-chloride solution. Germanium atoms produced by neutrinos form $GeCl_4$. A measured amount of inactive $GeCl_4$ is added at the beginning of a run to help determine the extraction efficiency. The $GeCl_4$ is extracted by bubbling nitrogen through the tank and is converted into GeH_4 for counting. The result from 65 runs taken from May 1991 to Jan 1997 [11] is

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

GNO (the Gallium Neutrino Observatory) is the successor of GALLEX after a major overhaul. The overall result of 19 runs taken from May 1998 to January 2000 [12][13] is

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

1.3.2 Water Čerenkov Experiments

Normal water (H_2O) Čerenkov experiments detect neutrinos by neutrino-electron scattering:

$$\nu_x + e \to \nu_x' + e'$$

The electron is scattered in the forward direction and, if it has enough energy, can be detected by its Čerenkov radiation. This reaction is most sensitive to electron neutrinos, ν_e, which can scatter by both neutral and charge current interactions. The cross section for ν_μ or ν_τ scattering, which can only occur by the neutral current interaction, is about one sixth of that for ν_e.

Čerenkov radiation is emitted when a charged particle travels faster than the speed of light in a medium. The kinetic energy of a particle is

$$E_K = mc^2(\frac{1}{\sqrt{1-\beta^2}} - 1)$$

where m is the particle mass, $\beta = v/c$ and v is the particle's velocity. Substituting $\frac{1}{n}$ for β gives the minimum kinetic energy a particle must have to emit Čerenkov radiation. For an electron in water this is 0.26 MeV. The light is emitted in a cone of half angle $\theta_C = \cos^{-1}(1/n\beta)$. For a relativistic electron in water this angle is about 42 degrees. The number of photons produced between wavelengths λ_1 and λ_2 per unit path length is

$$\frac{dN_\gamma}{dL} = 2\pi\alpha Z^2 \int_{\lambda_1}^{\lambda_2} (1 - \frac{1}{\beta^2 n^2}) d\lambda/\lambda^2$$

where α is the fine structure constant and Ze is the charge of the particle. This corresponds to about 220 Čerenkov photons, between 250 and 600 nm, per MeV for an electron in water with 5–15 MeV of kinetic energy.

Water Čerenkov experiments detect neutrino events in real time making them sensitive to day-night variations. The energy and direction of the scattered electrons carries information about the energy and direction of the incident neutrinos. Naturally occurring radioactivity masks the neutrino signal at low energy. Radioactive sources come from impurities in the water, the PMTs and the surrounding rock and also from the products of cosmic ray interactions. As a result, water Čerenkov detectors have thresholds of about 5 MeV. They are therefore only sensitive to ^{8}B and hep neutrinos.

Kamiokande

There have been a number of generations of water Čerenkov experiments located 1000 metres underground (2700 metres of water equivalent) in the Kamioka mine, in the Japanese Alps about 200 km west of Tokyo. The original experiment was the Kamioka <u>N</u>ucleon <u>D</u>ecay <u>E</u>xperiment or Kamiokande I. As the name suggests it was built to search for proton decay. The experiment was upgraded to become Kamiokande II in 1985 in order to make a measurement of the solar neutrino flux. It had better cosmic ray vetoing, improved water purification systems and better electronics. The detector was a cylindrical tank containing 3000 tons of H_2O. The inner volume was 2140 tons of water viewed by 948 20-inch photomultiplier tubes (PMTs). There was an outer volume, about 1.5 m thick, viewed by 123 PMTs to veto cosmic rays and also to provide shielding from gammas and neutrons coming from the surrounding rock. The distribution of scattered electrons showed a peak pointing away from the Sun, confirming that the Sun is a neutrino source. The experiment was again upgraded in 1990 to become Kamiokande III. Dead PMTs were replaced, light reflectors and new electronics were installed and the radon levels reduced. The combined ^{8}B flux measured from Kamiokande II and III [14], made with a threshold that was reduced from 9.3 MeV down to 7.0 MeV, is

$$2.80 \pm 0.19\,(\text{stat.}) \pm 0.33\,(\text{sys.}) \times 10^6\,\text{cm}^{-2}\text{s}^{-1}$$

The BP2000 standard solar model ^{8}B flux [6] is

$$5.05 \times 10^6 (1.00^{+0.20}_{-0.16}) cm^{-2} s^{-1}$$

Super-Kamiokande

Super-Kamiokande is a much larger version of Kamiokande. The detector consists of 50,000 tons of water in a cylindrical cavity, 41.4 m high and 39.3 m in diameter. The inner detector is viewed by 11,146 50 cm PMTs. The outer detector is 2.6 m thick and is viewed by 1,885 20 cm PMTs. The measured ^{8}B flux [15], made with a threshold that was reduced from 6.5 MeV to 5 MeV over the period May 1996 to September 2000, is

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

Water Čerenkov experiments measure events in real time and so the difference between day (Φ_d) and night (Φ_n) fluxes can be studied. The transition between day and night is defined as when the solar elevation passes through zero. The day-night asymmetry, defined as $\mathcal{A} = (\Phi_n - \Phi_d)/\frac{1}{2}(\Phi_n + \Phi_d)$, is

$$\mathcal{A} = 0.033 \pm 0.022\,(\text{stat.})\,^{+0.013}_{-0.012}\,(\text{sys.})$$

The energy spectrum of the scattered electrons is also useful for searching for spectral distortions. The fit to an undistorted spectrum is excellent with a χ^2/d.o.f. = 19.0/18. The seasonal variations due to the eccentricity of the Earth's orbit are also as expected with a χ^2/d.o.f. of 3.7/7. Super-Kamiokande also places an upper limit on the hep flux of 4.3 times the BP2000 prediction.

The Sudbury Neutrino Observatory

The Sudbury Neutrino Observatory (SNO) is also a water Čerenkov detector, except that it uses heavy water (D_2O), rather than normal water (H_2O). The use of heavy water enables the electron neutrino flux and the total flux of all flavours of neutrino to be measured separately. The Sudbury Neutrino Observatory and its latest results are described in section 1.6.

1.4 The Solar Neutrino Problem

The original solar neutrino problem was the difference between the solar neutrino flux measured by the chlorine experiment and the standard solar model prediction. The ^{8}B flux is very sensitive to the core temperature of the Sun so, by itself, the chlorine result could be attributed to errors in the solar model. When the results of the other solar neutrino experiments are added into the picture, much stronger statements can be made. Figure 1.3 shows the difference between the experimental measurements and predictions for all solar neutrino experiments except SNO. All experiments detect fewer neutrinos than expected and all of the experiments have been subject to rigorous testing.

The gallium experiments are particularly problematic because the flux of pp neutrinos is thought to be known to 1%. Both SAGE and GALLEX measure a rate

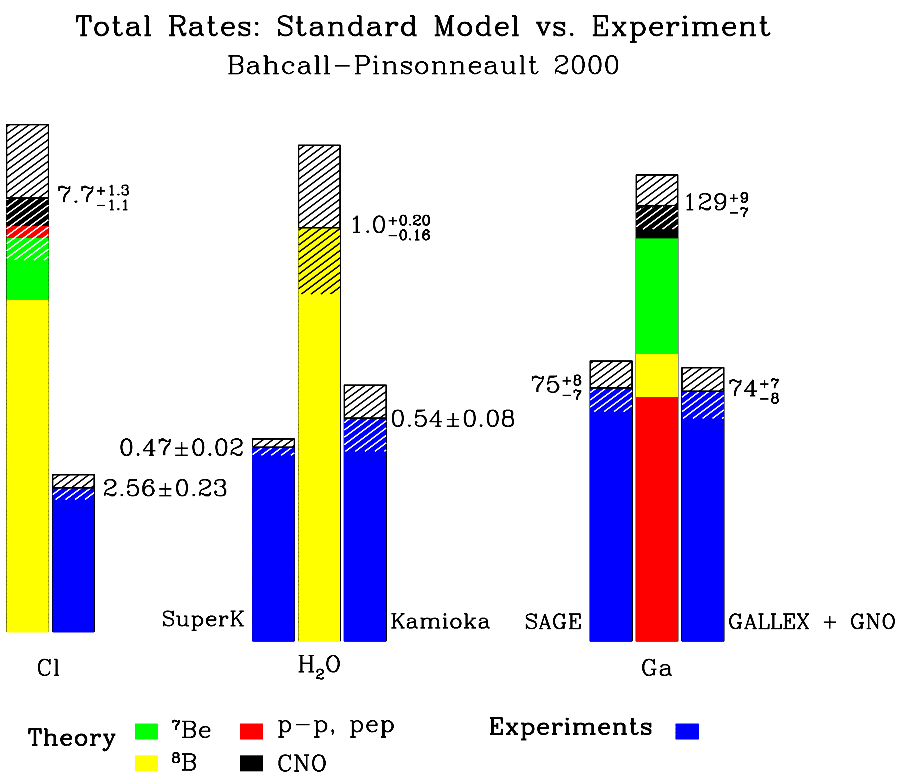

Figure 1.3: The solar neutrino problem. This plot compares the predictions of the BP2000 standard solar model to the measured rates of all solar neutrino experiments except SNO. The model predictions are colour coded to show the components from different nuclear reactions. The cross hatching indicates the 1σ uncertainties of both the experiments and predictions. Taken from [7].

that allows for the pp neutrinos but little else. It therefore appears that the ^{7}Be neutrino flux is almost non-existent. The Kamioka experiments show that the ^{8}B neutrino flux is about half the expected value. It is very hard to suppress the ^{7}Be neutrino flux and not the ^{8}B neutrino flux because ^{7}Be is needed to produce ^{8}B, see figure 1.1. The lack of ^{7}Be neutrinos is therefore very paradoxical.

Figure 1.4 shows the results of a fit, described in [16], to the neutrino fluxes using the results of the Cl, Ga and Kamioka experiments. The fit was done varying the fluxes of the pp, ^{7}Be, ^{8}B and CNO neutrinos as free parameters, within an overall solar luminosity constraint and with undistorted spectra. The best fit is

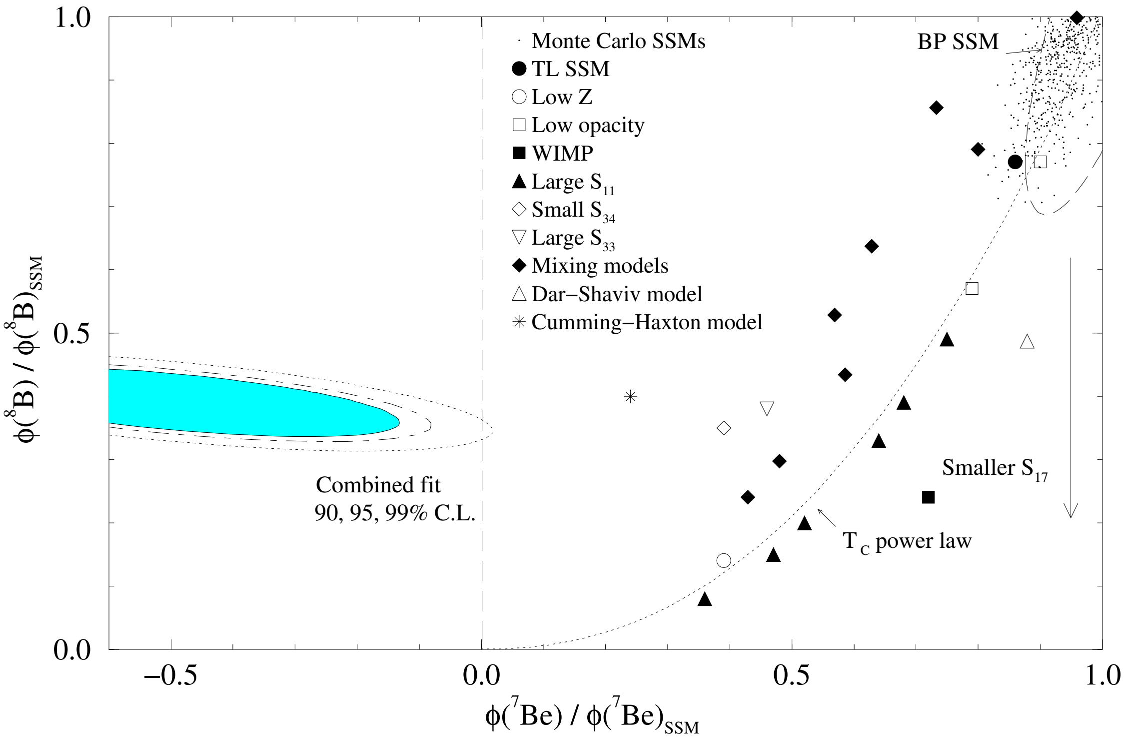

Figure 1.4: The fitted flux of ^{8}B and ^{7}Be neutrinos expressed as a fraction of the SSM prediction at 90, 95 and 99% confidence level using the results of the Cl, Ga and Kamioka experiments. The 90% confidence level Bahcall-Pinsonneault SSM allowed region is shown at the top right. The core temperature power law and various non-standard solar models are also displayed. Taken from [16].

with a negative (unphysical) flux of ^{7}Be neutrinos. No non-standard solar model yet discovered can explain the observed solar neutrino flux. The fit was repeated four times omitting either the Cl result, the Ga result, the Kamioka result or the luminosity constraint. Each of these four fits produced a result that is incompatible with astrophysical solutions.

This result shows that even if one type of experiment had a systematic error, no known alteration of the standard solar model could explain the other experimental results. How can this conundrum be resolved? A possibility is that the Sun produces neutrinos in the expected numbers but something happens to them on their passage to the Earth which reduces the observed flux. The most popular mechanism which can reduce the experimental rates is neutrino flavour oscillations.

Under this hypothesis the electron neutrinos produced by the Sun are converted into another flavour of neutrino, which the experiments are insensitive to, on route to the Earth. The following section describes neutrino oscillations in more detail.

1.5 Neutrino Oscillations

Neutrino flavour oscillations were first suggested as a solution to the solar neutrino problem by Pontecorvo and Gribov [17][18]. If neutrinos have non-degenerate masses and the flavour eigenstates are mixtures of the mass eigenstates then neutrino oscillations will occur. A neutrino which is created as an electron neutrino may, after propagating some distance, be detected as a different flavour of neutrino. The simplest form of neutrino oscillations are those which occur in the vacuum and these are discussed next.

1.5.1 Vacuum Oscillations

This discussion follows that in Bahcall chapter 9 [5]. To keep things simple, suppose that there is only mixing between the electron and muon type neutrinos. The weak and mass eigenstates will be related by the unitary transform

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

where θ_V is the vacuum mixing angle and ν_1 and ν_2 are the mass eigenstates with masses m_1 and m_2. The three flavour version of the above matrix is the analogue of the CKM matrix which specifies the mixing between the weak and strong eigenstates of the quarks. The mass eigenstates evolve in time as

$$|\nu_i\rangle_t = \exp(-iE_it/\hbar)|\nu_i\rangle \tag{1.2}$$

If an electron neutrino is created at t=0, then at time t its state will be

$$|\nu_e\rangle_t = \cos\theta_V \exp(-iE_1t/\hbar)|\nu_1\rangle + \sin\theta_V \exp(-iE_2t/\hbar)|\nu_2\rangle \tag{1.3}$$

The probability amplitude for an electron neutrino remaining an electron neutrino after time t is then

$$\langle \nu_e | \nu_e \rangle_t = \cos^2 \theta_V \exp(-iE_1 t/\hbar) + \sin^2 \theta_V \exp(-iE_2 t/\hbar) \tag{1.4}$$

and so the probability of detecting an electron neutrino after time t is

$$P(\nu_e \to \nu_e) = |\langle \nu_e | \nu_e \rangle_t|^2 = 1 - \sin^2 2\theta_V \sin^2[(E_2 - E_1)t/2\hbar] \tag{1.5}$$

We can assume the two mass eigenstates have the same momentum and therefore different energies (a discussion of the proper wavepacket approach may be found in [19]). For relativistic neutrinos $E \approx p + m^2/2p$ and therefore the energy difference is

$$E_2 - E_1 = \frac{(m_2^2 - m_1^2)c^4}{2E} \equiv \pm \frac{\Delta m^2 c^4}{2E} \tag{1.6}$$

where $\Delta m^2 \equiv |m_2^2 - m_1^2|$ and so the plus sign applies if $m_2 > m_1$ and vice-versa.

The probability that an electron neutrino remains an electron neutrino can then be re-written as

$$P(\nu_e \to \nu_e) = 1 - \sin^2 2\theta_V \sin^2[\pi R/L_V] \tag{1.7}$$

where R is the distance travelled after time t and L_V, the vacuum oscillation length, is

$$L_V \equiv 4\pi E\hbar/\Delta m^2 c^3 = 2.48\, E(\mathrm{MeV}) \,/\, \Delta m^2(\mathrm{eV}^2) \text{ metres} \tag{1.8}$$

From equation 1.7 it can be seen that the probability of detecting an electron neutrino oscillates from a maximum of 1 to a minimum of $1 - \sin^2 2\theta_V$. The maxima are at integer multiples of the oscillation length. Conversely the maximum probability of detecting ν_μ is $\sin^2 2\theta_V$. From equation 1.8 it can be seen that long baselines are needed to detect the effects of a very small Δm^2. Solar neutrino experiments use the longest baseline and are therefore sensitive to the smallest Δm^2.

Vacuum oscillations are not the only way in which neutrino oscillations can occur. Matter enhanced oscillations provide a more powerful way in which the flavour content of solar neutrinos could be altered.

1.5.2 Matter Enhanced Oscillations

Neutrino oscillations can be greatly enhanced when propagating through matter because of the presence of electrons. The effect of matter on neutrino oscillations was first derived by Wolfenstein [20]. Mikheyev and Smirnov extended this theory to address the varying density in the Sun [21][22]. The large amount of mixing that can be caused by matter in the Sun in called the MSW effect after its originators.

When neutrinos pass through matter they undergo coherent forward scatterings which produces an index of refraction. The scatterings can be divided into neutral current and charged current interactions. The neutral current scattering is the same for all flavours of neutrino and has no net effect on oscillations. Electron neutrinos can undergo coherent forward scattering off electrons in matter through the charged current interaction. Ordinary matter has zero muon and tau densities and so there is no analogue for mu and tau neutrinos. This asymmetry is the basis of the MSW effect.

A discussion of matter oscillations for two flavours of neutrino can be found in Bahcall chapter 9 [5]. The net result is that the oscillation length and mixing angle in matter is different to that in the vacuum. The mixing angle in matter, θ_M, is given by

$$\tan 2\theta_M = \frac{\tan 2\theta_V}{[1 \pm (L_V/L_e)\sec 2\theta_V]} \tag{1.9}$$

where the plus sign applies for $m_2 < m_1$ and vice-versa. L_e is the neutrino-electron interaction length and is

$$L_e = \frac{\sqrt{2}\pi\hbar c}{G_F n_e} \tag{1.10}$$

where G_F is the Fermi coupling constant and n_e is the electron density. The electron density as a function of the solar radius is a prediction of standard solar models.

From equation 1.9, for $m_2 \geq m_1$ the mixing in matter will be maximal ($\theta_M = \pi/4$) when $L_V/L_e = \cos 2\theta_V$. This is known as the 'resonance condition'. For the high electron densities in the centre of the Sun the denominator of equation 1.9

can be large, depending on the neutrino energy, and so $\theta_M \approx \pi/2$. In this case an electron neutrino will be created almost entirely in state ν_2. As the neutrino travels out through the Sun the electron density drops and θ_M will change towards the vacuum mixing angle, passing through resonance. If the rate of change of the mixing angle is small, the *adiabatic* condition, then the neutrino will remain in state ν_2. If the change is more rapid, the *non-adiabatic* condition, then the neutrino can switch to state ν_1. In this way, large energy-dependent alterations can be made to the flavour composition of neutrinos emerging from the Sun.

1.5.3 Possible Solutions

The number of neutrinos produced as a function of energy and the electron density as a function of the solar radius are predictions of the standard solar model. The amount of suppression of the experimental rates can therefore be calculated for different values of θ_V and Δm^2. A global analysis of all the experiments can therefore show the regions of the $\tan^2\theta_V$,Δm^2 plane which fit the experimental rates. One such global analysis may be found in [23]. This analysis uses the predictions of the BP2000 SSM. The input data is the total rates from the chlorine and gallium experiments, the day and night energy spectra from Super-Kamiokande and the CHOOZ spectral results [24] (CHOOZ is a reactor neutrino experiment). The day and night flux measurements help constrain the allowed oscillation parameters because some oscillations could occur when passing through matter in the Earth. Figure 1.5 shows the result of this analysis without SNO results.

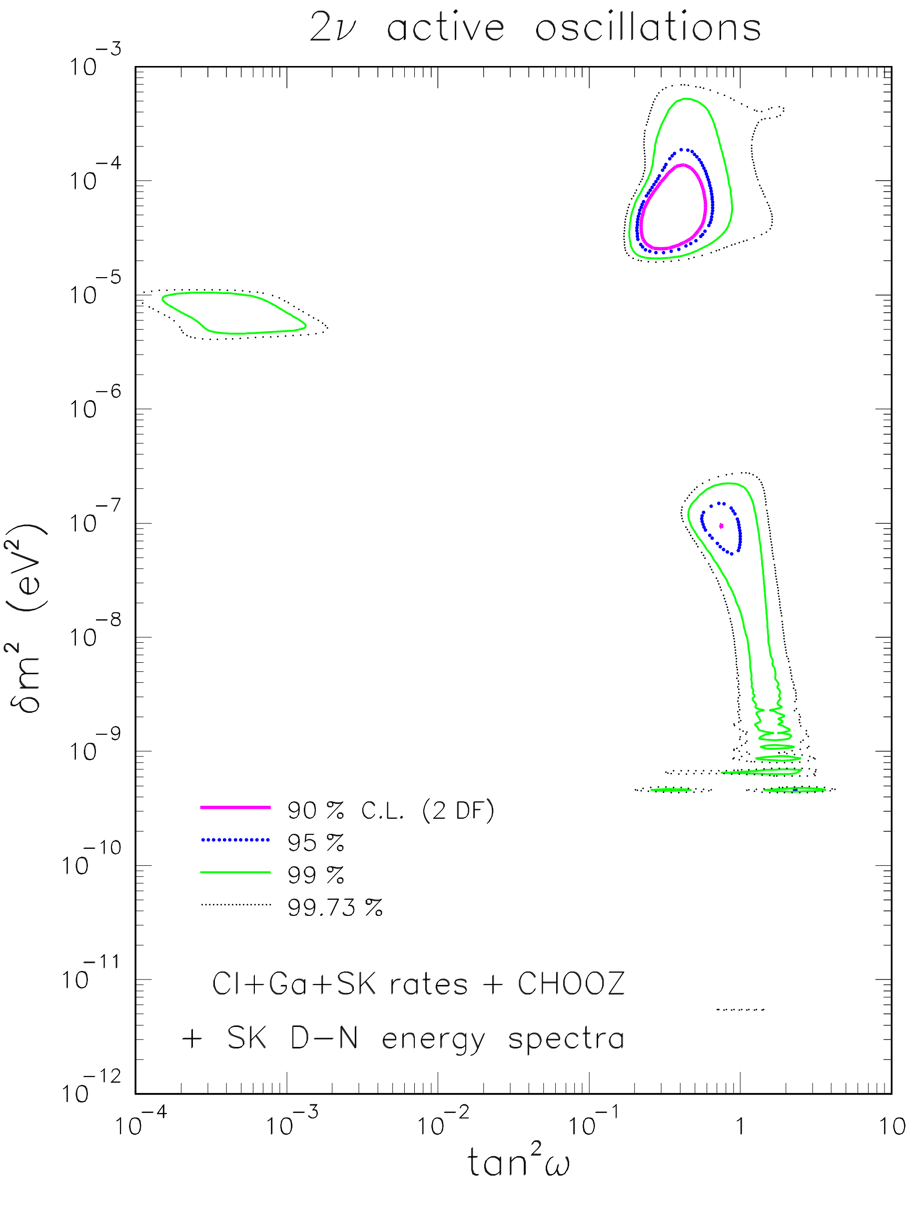

Figure 1.5: Pre-SNO global solutions for oscillations to active neutrinos. In the text ω is called θ_V. The global analysis used the BP2000 SSM, the total rates from the chlorine and gallium experiments, the day and night energy spectra from Super-Kamiokande and the CHOOZ spectral results. The confidence level contours are 90%, 95%, 99% and 99.73% (3σ). Taken from [23].

1.6 The Sudbury Neutrino Observatory

The Sudbury Neutrino Observatory (SNO) is a heavy water Čerenkov detector and is designed to test the neutrino oscillation hypothesis. All solar neutrino experiments so far have measured a short fall in the number of neutrinos detected. The favoured explanation is that some fraction of the electron neutrinos have oscillated to a different neutrino flavour. H. H. Chen proposed a heavy water Čerenkov detector as a way to test this hypothesis [25]. Neutrinos can interact with heavy water, D_2O, by both charged-current and neutral-current interactions. By measuring the rate of these interactions the flux of electron neutrinos and the total flux of all active flavours of neutrino can be determined.

1.6.1 The SNO Detector

SNO is located 2040 metres underground in the INCO, Ltd. Creighton mine, near Sudbury, Ontario, Canada. The granite rock provides the equivalent of 6000 metres of water shielding and reduces the number of cosmic ray muons passing through the detector to about 70 per day. A diagram of the detector is shown in figure 1.6. The detector is inside a large barrel shaped cavity, 22 m in diameter and 34 m in height. The heavy water is contained in a transparent acrylic flask, 12 m metres in diameter and 5.5 cm thick. Near the top of the cavity is a deck structure. The cavity below is filled with 'light' water (H_2O) and the acrylic vessel is suspended from the deck inside this water by 10 rope loops. The purified light water provides both support and shielding. Surrounding the acrylic vessel is a geodesic structure, 17.8 m in diameter, carrying 9456 inward-looking photomultiplier tubes (PMTs).

1.6.2 Neutrino Interactions in SNO

SNO can detect three types of neutrino interactions: the charged current, neutral current and elastic scattering interactions. There are two additional reactions which are only sensitive to anti-neutrinos.

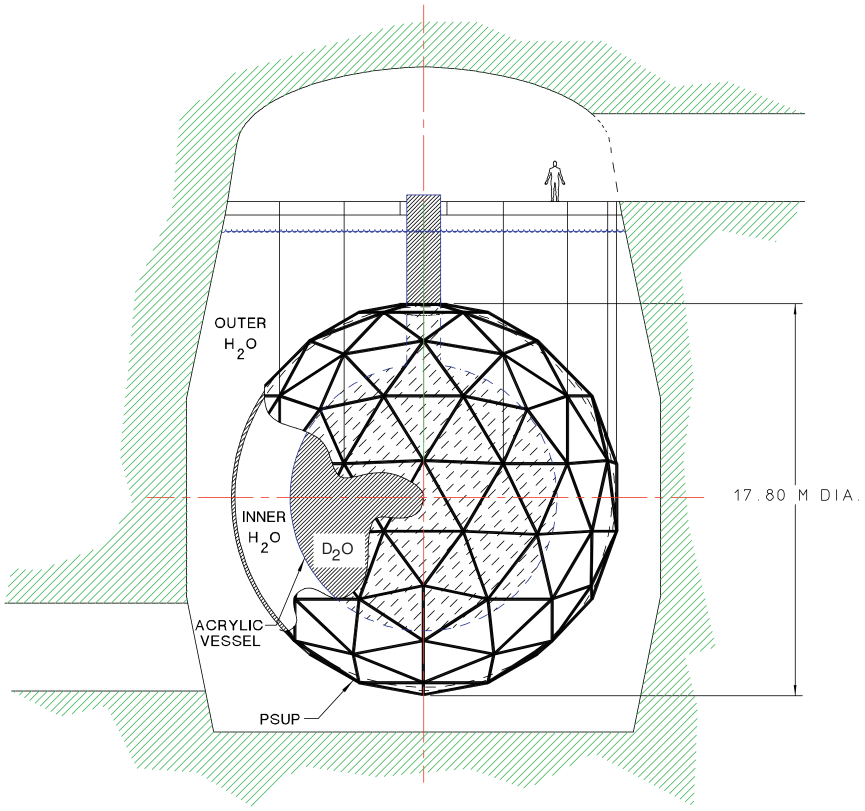

Figure 1.6: The SNO detector is located in a cavity 2040 metres underground. The target is a kilotonne of heavy water (D_2O) contained in an acrylic flask and submerged in H_2O. The heavy water is viewed by an array of photomultiplier tubes which are arranged on a geometrical support structure (marked PSUP). The PSUP divides the H_2O into inner and outer volumes. Taken from [26].

The Charged Current Interaction

The reaction

$$\nu_e + d \rightarrow p + p + e^-$$

proceeds purely by the charged current interaction and is referred to by that name in the context of SNO. The reaction has a threshold of 1.44 MeV. The reaction is detected by the Čerenkov light produced by the electron. The energy of the emitted electron is almost equal to the energy of the incident neutrino minus 1.44 MeV and is therefore a sensitive way to measure the neutrino energy spectrum. There is a weak tendency for the electrons to be emitted in the opposite direction to the incident neutrino, i.e. back towards the Sun. The angular distribution is expected to be $1 - 0.340 \cos\theta_{\odot}$ [27]. This reaction occurs uniformly throughout the D_2O volume. Approximately four of these events were detected per day inside 550 cm and above an effective electron kinetic energy threshold of 6.75 MeV, see section 1.6.5. This number might double for a 5 MeV threshold.

The Neutral Current Interaction

Neutrinos can disintegrate deuterium via a neutral current interaction:

$$\nu_x + d \rightarrow \nu_x + p + n$$

where ν_x symbolises any flavour of active neutrino. This reaction is equally sensitive to all flavours of neutrino and enables the measurement of the total neutrino flux from the Sun. The reaction has a threshold of 2.2 MeV. The signature of this reaction is the detection of a free neutron. The ways in which the neutron is detected during the three phases of running are described in section 1.6.3. The neutrons are produced uniformly in the D_2O volume. They then wander around before being detected and those near the acrylic vessel are likely to capture on H giving no signal. This results in a characteristic radial distribution. The neutral current interaction provides no directional or spectral information. Around 13 neutrons are expected to be produced per day with the detection efficiency depending on the phase of the experiment and the energy threshold of the analysis.

Elastic Scattering

The elastic scattering reaction is the same one used by normal water Čerenkov detectors:

$$\nu_x + e^- \rightarrow \nu_x' + e'^-$$

This reaction can occur by both charged and neutral current interactions for the electron neutrino but only by the neutral current interaction for the other flavours. The cross section for electron neutrinos is about six times that for the other flavours. There is no threshold for this interaction. The energy is shared between the neutrino and electron and so does not provide as much spectral information as the charged current reaction above. The direction of the electron is strongly peaked away from the Sun. This reaction occurs uniformly throughout the D_2O and H_2O volumes. About 0.44 events were detected per day inside 550 cm and above a threshold of 6.75 MeV, see section 1.6.5. This might rise to about one a day for a 5 MeV threshold.

Anti-Neutrino Interactions

These reactions are unlikely to be relevant to the solar neutrino problem but will be of interest if there is a supernova. Antineutrinos can be detected by the neutral current and the elastic scattering interactions. In addition there are two reactions based on the inverse beta decay of the proton:

$$\bar{\nu}_e + d \rightarrow n + n + e^+$$
$$\bar{\nu}_e + p \rightarrow n + e^+$$

The first reaction will occur in the D_2O and the second in the H_2O. These reactions will be detectable from the positron and the neutron capture.

1.6.3 Phases of the Experiment

SNO will have three phases of running. The purpose of these phases is to pin down the neutral current (NC) rate. The first phase is with pure D_2O. In this phase neutrons from the NC interaction were detected by the 6.25 MeV gamma ray emitted from neutron captures on deuterium. The gamma ray Compton scatters resulting in Čerenkov light from the relativistic electron. This phase of running has been completed and the first results presented, see section 1.6.5.

For the second phase of running, sodium chloride was added to the D_2O. This second phase started around the beginning of June 2001. The ^{35}Cl greatly enhances the neutron signal. Capture on ^{35}Cl results in a gamma cascade totalling 8.6 MeV. ^{35}Cl also has a larger neutron capture cross section than deuterium and most of the neutrons produced will capture on it. The multiplicity of the gamma decay results in events with a hit pattern which is more isotropic than events from a single Čerenkov electron. Although the isotropy difference is not enough to reliably identify individual NC events it is enough to estimate the contribution of the NC signal to an overall hit pattern distribution.

In third phase of SNO an array of ^{3}He filled proportional counters will be positioned inside the D_2O. The counters will be arranged in a square lattice with 1 m spacing and will have a total length of 770 m. These counters can identify individual neutrons by the reaction

$$n +^3 \mathrm{He} \rightarrow t + p + 0.764\,\mathrm{MeV}$$

Pulse shape discrimination can separate neutron events from βs and αs. The counters have to be ultra pure to limit the number of neutrons generated by radioactivity in the counters.

1.6.4 Backgrounds

Cosmic rays are a fairly benign source of background. Only about 70 cosmic ray muons pass through the detector every day. These are easily identifiable by the vast amount of light they produce. There are 91 PMTs mounted on the outside of the PMT support structure that view the outer volume of H_2O. These PMTs are used to veto the comic rays. To reject events caused by spallation products all events immediately following a cosmic ray (within 20 s) are rejected.

High energy gamma rays from the surrounding rock are attenuated by the outer volume of H_2O and the numbers reaching the inside of the PMT sphere are negligible. The amount of ^{222}Ra diffusing into the H_2O outer volume is greatly reduced by an 8mm coating of Urylon on the cavity walls. The H_2O is separated from the radon-rich laboratory air by a nitrogen cover gas.

The PMT support structure and PMTs themselves are also a source of radioactive background. The materials which were used to construct the PMTs and the support structure were carefully selected to minimise radioactive contaminants. Gamma rays above 2.2 MeV can photodisintegrate a deuteron so the D_2O is shielded from the PMTs by the 2.4 m thick H_2O inner volume. The acrylic vessel was constructed by bonding together 122 acrylic panels. Acrylic is a simple hydrocarbon and can be manufactured with very low intrinsic radioactivity.

The radiopurity of the D_2O and H_2O is also of great importance. The number of neutrons produced by photodisintegration must be well below the rate expected from solar neutrinos. The level of internal radioactive background also sets the energy threshold below which analysis becomes impossible. The highest energy internal backgrounds come from the decays of ^{208}Tl and ^{214}Bi which come from the ^{232}Th and ^{238}U decay chains respectively. Both decays can produce gamma rays which are capable of dissociating deuterium. To keep the neutron background below one tenth of the SSM rate requires stringent limits of 3.7×10^{-15} g/g of ^{232}Th and $45. \times 10^{-15}$ g/g of ^{238}U in the D_2O. Complex water purifying and assaying systems were developed to successfully achieve these targets.

1.6.5 First Results from SNO

The SNO collaboration announced its first results on 18th June, 2001 [28]. The data presented was recorded between 2nd November 1999 and 15th January 2001, and has a live time of 240.95 days. First various instrumental backgrounds were removed from the data. The remaining events had their positions and directions in the detector reconstructed using the calibrated PMT times. The effective energy of each event was determined from the number of PMTs which fired and the position and direction of the event.

As was described in section 1.6.2 the charged current (CC), neutral current (NC) and elastic scattering (ES) interactions have different energy, angular and radial dependences. The contribution of each of these components to the total signal inside a fiducial volume of radius 550 cm was determined using an extended maximum likelihood method. The number of neutron captures cannot be converted into the total active neutrino flux without first subtracting the number of neutrons produced by radioactive backgrounds. The neutral current flux will be the topic of future SNO publications. The ^{8}B flux measured by the CC and ES reactions, with a kinetic energy threshold of $T_{eff} = 6.75\,\mathrm{MeV}$, and assuming no oscillations are:

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

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

where the theoretical uncertainty is the CC cross section uncertainty. The energy spectrum of the CC events shows no evidence for distortions. The SNO ES measurement is consistent with the high precision Super-Kamiokande result:

$$\Phi^{\mathrm{ES}}_{\mathrm{SK}} = 2.32 \pm 0.03\,(\mathrm{stat.})\,^{+0.08}_{-0.07}\,(\mathrm{sys.}) \times 10^6\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$$

The SNO CC flux is only sensitive to the electron flavour component of the ^{8}B flux. The Super-Kamiokande ES flux is partially sensitive to the flux of mu or tau flavours. Symbolically we can write

$$\Phi^{CC}_{SNO} = \Phi_e$$

$$\Phi^{ES}_{SK} = \Phi_e + \epsilon\Phi_{\mu\tau}$$

where $\Phi_{\mu\tau}$ is the combined flux of mu and tau flavoured neutrinos and $\epsilon = 1/6.481$. The electron neutrino flux Φ_e, mu and tau neutrino flux $\Phi_{\mu\tau}$ and the total active neutrino flux Φ_x trivially follow:

$$\Phi_e = \Phi^{CC}_{SNO}$$

$$\Phi_{\mu\tau} = (\Phi^{ES}_{SK} - \Phi^{CC}_{SNO})/\epsilon$$

$$\Phi^{SK+SNO}_x = \Phi_e + \Phi_{\mu\tau} = (1 - 1/\epsilon)\Phi^{CC}_{SNO} + \Phi^{ES}_{SK}/\epsilon$$

This is more clearly shown on the $\Phi_{\mu\tau}$,Φ_e plane, see figure 1.7. The SNO result is only sensitive to Φ_e and so is a vertical bar. The SK result depends somewhat on

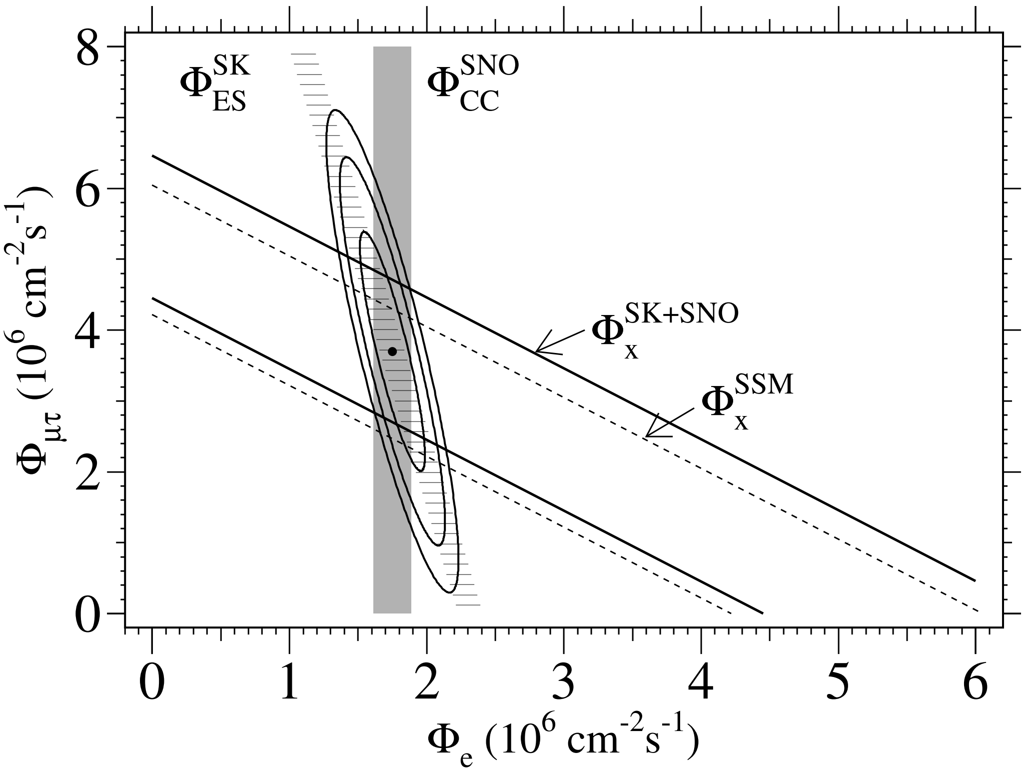

Figure 1.7: Flux of mu or tau flavoured ^{8}B neutrinos against the flux of electron flavoured ^{8}B neutrinos. See the text for details. Taken from [28][29].

$\Phi_{\mu\tau}$ and so is a slanted bar. The diagram shows the best fit point and the 68.27%, 95% and 99% confidence levels. This comparison shows at the 99.96% confidence level that there is a non-electron flavour active neutrino component in the solar neutrino flux. The best fit total flux of active ^{8}B neutrinos is

$$\Phi_{\mathrm{x}}^{\mathrm{SK+SNO}} = 5.44 \pm 0.99 \times 10^6 \mathrm{cm}^{-2}\mathrm{s}^{-1}$$

in close agreement with standard solar model predictions [28][29]. This result and the predicted SSM flux with 1σ errors are shown as diagonal lines.

Figure 1.8 shows allowed values of $\tan^2\theta_V$ and Δm^2 from a global analysis of all solar neutrino experimental data including the SNO CC measurement. The result of this analysis without the SNO CC measurement was given in figure 1.5

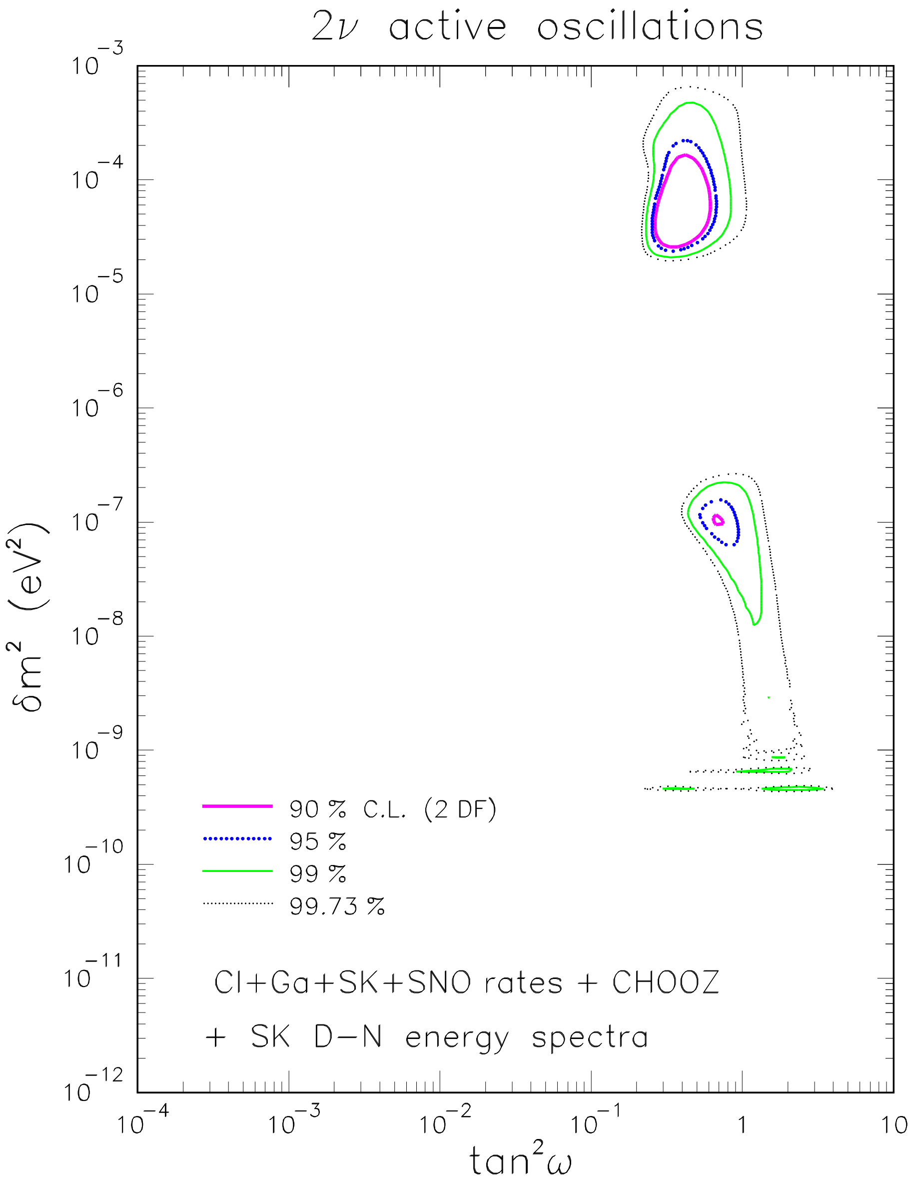

Figure 1.8: Global solutions for oscillations to active neutrinos including the first SNO CC measurement. In the text ω is called θ_V. The global analysis used the BP2000 SSM, the total rates from the chlorine and gallium experiments, the day and night energy spectra from Super-Kamiokande, the CHOOZ spectral results and the SNO CC measurement. The confidence level contours are 90%, 95%, 99% and 99.73% (3σ). Taken from [23].

on page 18. It can be seen that even the first results from SNO have affected the allowed regions. The small mixing angle solution disappears at $> 3\sigma$. The most favoured solutions all involve near maximal mixing.

SNO will provide much more in the way of exciting data. The neutral current measurement will allow the non-electron neutrino flux to be determined with much higher accuracy. The precision of the charged current measurement will be increased and the systematic error reduced. The energy threshold with which the charged current and elastic scattering response are analysed will be reduced which will allow searches for spectral distortions. These measurements will put much tighter constraints on the properties of neutrinos.

1.7 The Layout of this Thesis

The topic of this thesis is the Photomultiplier Tube Calibration of the Sudbury Neutrino Observatory. The thesis is organised as follows. Chapter 2 provides background information on the elements of the SNO detector which are relevant to the work contained in this thesis. Items discussed include the photomultiplier tubes themselves, the SNO electronics and various calibration sources. The Photomultiplier Tube CAlibration (PCA) has two main parts: the walk calibration and the gain calibration. Chapter 3 states the goals of these calibrations and the methods by which these goals may be achieved. Chapter 4 gives a detailed account of how these calibrations are implemented in the SNO collaboration's software package. It describes how the calibration constants are generated and applied and also explains how badly behaving channels are cut. Chapter 5 investigates the performance of the PCA and explores various possible systematics. Chapter 6 contains some studies on the rate sensitivity of the electronics. Finally, in chapter 7, the beneficial impact of the PCA on event reconstruction is demonstrated using data from the ^{8}Li source and neutrino data.

Chapter 2

Background Information

This chapter describes the features of SNO which are relevant to the photomultiplier calibration. The SNO photomultiplier tubes are described in section 2.1. Section 2.2 describes the arrangement of the SNO electronics and the data they collect from the photomultipliers. Section 2.3 details the first stage in the calibration of this data, known as the electronics calibration. The validation of the electronics calibration is briefly commented on in section 2.4. The light source used for the photomultiplier calibration, the laserball, is described in section 2.5. Two other calibration sources, the ^{16}N and ^{8}Li sources, are discussed in sections 2.6 and 2.7. Finally SNOMAN, the Monte Carlo and analysis software package developed and used by the SNO collaboration, is described in section 2.8.

2.1 The SNO Photomultiplier Tubes

Photomultiplier Tubes (PMTs) are used to measure the arrival time, and the approximate number, of photons incident upon them. The front of a PMT consists of a glass window with a thin photocathode layer deposited on its inside surface. Photons can eject electrons, called photoelectrons, from this layer by the photoelectric effect. The inside of a PMT is evacuated and photoelectrons are focused by electric fields down to the dynode stack. The dynode stack is a series of electrodes which produce multiple secondary electrons for each incident electron. For inter-dynode

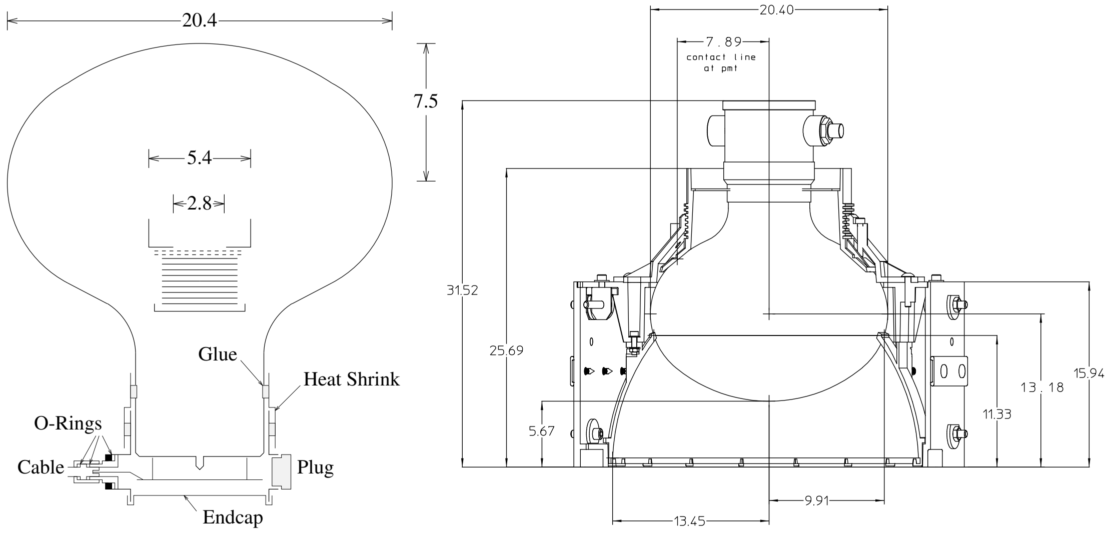

Figure 2.1: The left hand plot is a schematic of the Hamamatsu R1408 PMT. The nine dynodes are shown as horizontal lines. The right hand plot shows a section through a PMT mounted in its hexagonal housing. The curved faces around the edge of the front face hold aluminium reflectors. All dimensions are in centimetres. Taken from [26].

potentials of 100–200 V the number of electrons is increased by a factor of 3 to 5 at each stage. The dynode stack therefore greatly amplifies the number of electrons and produces a detectable pulse.

The photomultipliers used in SNO are the Hamamatsu R1408 PMT [26]. They have a 20 cm diameter, mushroom shaped glass envelope. The potential difference between the photocathode and the anode is in the range 1700 to 2100 V. There are 9 dynodes and this results in an amplification factor of about 10^7. The left hand drawing in figure 2.1 is a schematic of a SNO PMT. The photocathode is sensitive in the range 250 nm to 550 nm. The long wavelength cutoff is set by the photocathode material while the short wavelength response is limited by the transparency of the glass bulb. The quantum efficiency of the PMTs (the fraction of incident photons which produce a photoelectron) has a maximum of nearly 30% [30]. The total transit time inside the photomultiplier from the cathode to the anode is about 30 ns. The spread in this transit time through the photomultiplier (the 'time jitter') has an RMS of about 1.5 ns. This spread is mostly due to the different trajectories from

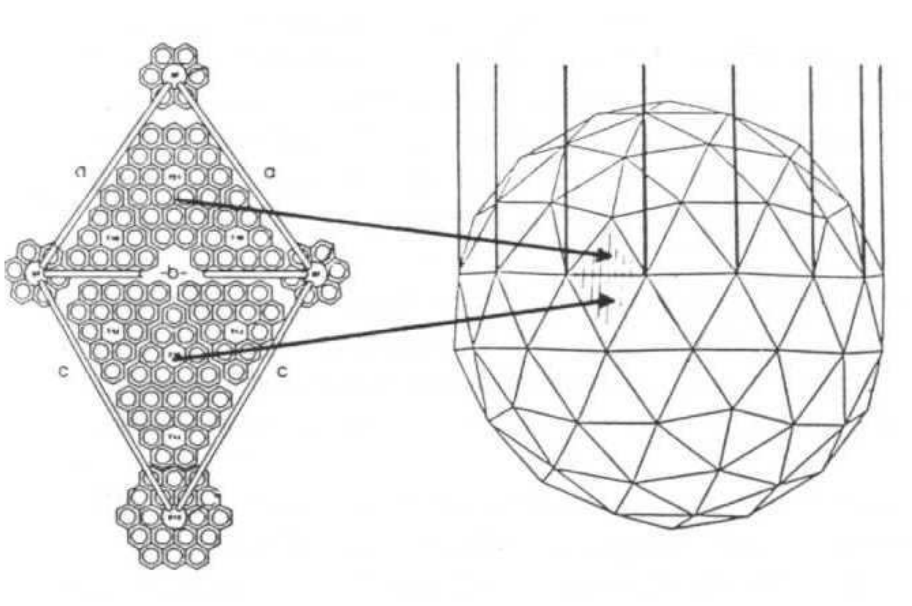

Figure 2.2: Rear view of a portion of the PMT support structure. The hexagonal PMT housings tessellate to form panels which are mounted as shown.

the cathode to the first dynode. There is also a small contribution, about 0.2 ns, from the different energies with which electrons are ejected from the photocathode. The dominant contribution to the radioactive content of the PMTs comes from the glass. A special, low activity, borosilicate glass was developed by Oxford and Schott Glaswerke for SNO. The total U and Th content of the PMTs is a factor of 14 below specifications.

The PMTs, by themselves, would have a photocathode coverage of 31%. Each PMT is surrounded by a 27 cm diameter light concentrator. This increases the photocathode coverage to 59%. The reflectivity of the concentrators reduces the effective coverage to 54%. The concentrators accept light up to a critical angle of 56.4°. They are constructed from 18 segments of dielectric coated aluminium sheet and are mounted on a plastic housing. The right hand plot of figure 2.1 shows a PMT in its housing. The plastic housings are hexagonal and tessellate to form panels of PMTs. Figure 2.2 shows how these panels are arranged on the PMT support structure. When the panels were being mounted, a diode laser was

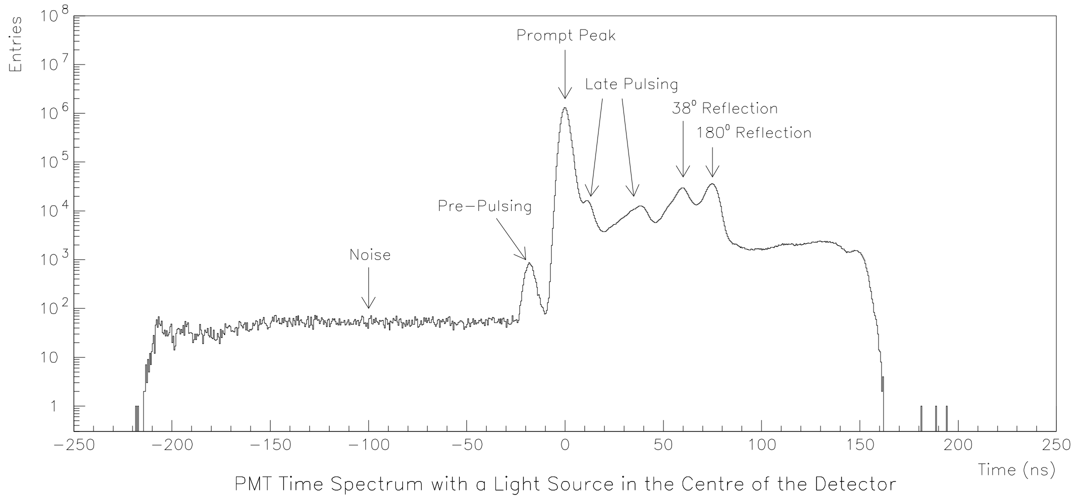

Figure 2.3: Histogram of the PMT firing times with a pulsed light source in the centre of the detector. Note the logarithmic scale. See the text for details.

attached perpendicular to the panel. Telescopic mounts were adjusted until the laser illuminated a 3” spherical target suspended at the centre of the PMT support structure, ensuring that all PMTs were correctly aligned.

PMTs are optimised to work in the absence of magnetic fields. The Earth’s magnetic field in SNO has a magnitude of $55\,\mu$T and is angled about 15° off the vertical. This reduces the PMT efficiency to 82% of the zero field value. The vertical component of the Earths magnetic field is cancelled by 14 horizontal coils embedded in the cavity walls. The residual magnetic field in the PMT region is $19\,\mu$T which reduces the PMT efficiency by only 2.5% from the zero field value.

Figure 2.3 shows a histogram of the PMT firing times with a pulsed light source at the centre of the detector. Light which travels directly from the source to the PMTs produces a large peak at zero nanoseconds. There are some other interesting features, some of which are due to PMT effects and some of which are from detector effects. The PMT based effects are:

Noise All PMTs produce noise pulses, even in the absence of incident light. The dominant cause is the thermal emission of electrons from the photocathode and dynodes. These thermal electrons are amplified by the dynode stack and are indistinguishable from genuine pulses. The water in SNO is cooled to 10° which minimises the thermal noise without risking freezing the D_2O. Radioactivity also causes noise pulses but this is a lesser effect. The ambient firing rate of the SNO PMTs is about 500 Hz.

Pre-pulsing Electrons lose energy when travelling through the photocathode because of collisions with other electrons and with the lattice. If the photocathode is made too thick then the electrons will not have enough energy to escape the surface barrier potential. The photocathode is therefore quite thin and is semi-transparent to photons. Pre-pulsing is due to photons passing through the photocathode, without being absorbed, striking the first dynode and ejecting an electron. These pulses come about 18 ns earlier than normal pulses because of the difference in time between an electron being accelerated from the photocathode to the first dynode and the light crossing time.

Late Pulsing Late pulsing is due to optical feedback inside the PMT. Light can be produced by electrons cascading down the dynode stack. If a photon travels back up to the photocathode it can trigger a second cascade. If the first cascade was too small to be recorded as a pulse (this is discussed below) then the time of the second pulse may be recorded. These pulses can come up to about 40 ns after the prompt pulse. There is also a late pulsing peak about 12 ns after the prompt light. The cause of this peak is under investigation and may be due to electrical pickup between PMT cables (this peak was not present when the PMTs were tested, see figure 6.9 of [30]).

After Pulsing Photoelectrons can ionise residual gas inside the PMT. Positive ions will drift back up towards the photocathode where they may trigger a second pulse. These pulses occur on the time scale of microseconds afterwards and so are not included in the figure.

Figure 2.3 also shows some detector based timing features. Direct light can be reflected from the PMTs and get detected on the other side of the detector. These 180° reflections come about 75 ns late, the light crossing time of the detector. The PMT/concentrator assembly also produces a strong reflection at 38°. Scattering also affects the time distribution, adding a late tail to all peaks.

In practice, a limit must be placed on the size of pulse that the electronics respond to, or the electronics, data acquisition and analysis would be swamped by noise. Small pulses could arise from thermal noise amplified over a small number of dynodes or electrical noise. The limit is imposed by a discriminator which only fires in response to a signal above a set threshold. In SNO, this threshold is about a quarter of the charge produced by a typical photoelectron [31].

There are five categories of PMT in SNO. The primary type of PMT are the 'normal' PMTs, numbering 9453. Neutrino events are detected using these PMTs and the data they provide is essential to neutrino analysis. Less than 2% of these tubes are currently dead [31]. There are 49 'low gain' PMTs which have a dynode tap to reduce the gain by a factor of about 100. These PMTs look inwards and increase the dynamic range of the detector response. There are 91 outward looking PMTs, or 'OWL' tubes, which view the outer volume of H_2O to veto cosmic rays. At the top of the neck of the acrylic vessel are 4 PMT which view the water pipes. These are known as 'neck' tubes and are used to veto background light from the water systems. Finally come the 22 Berkeley Underwater Test Tubes, or BUTTs, located in the light water. These PMTs are easily accessible and can be extracted to look for the possible effects of long term operation in ultra-pure water.

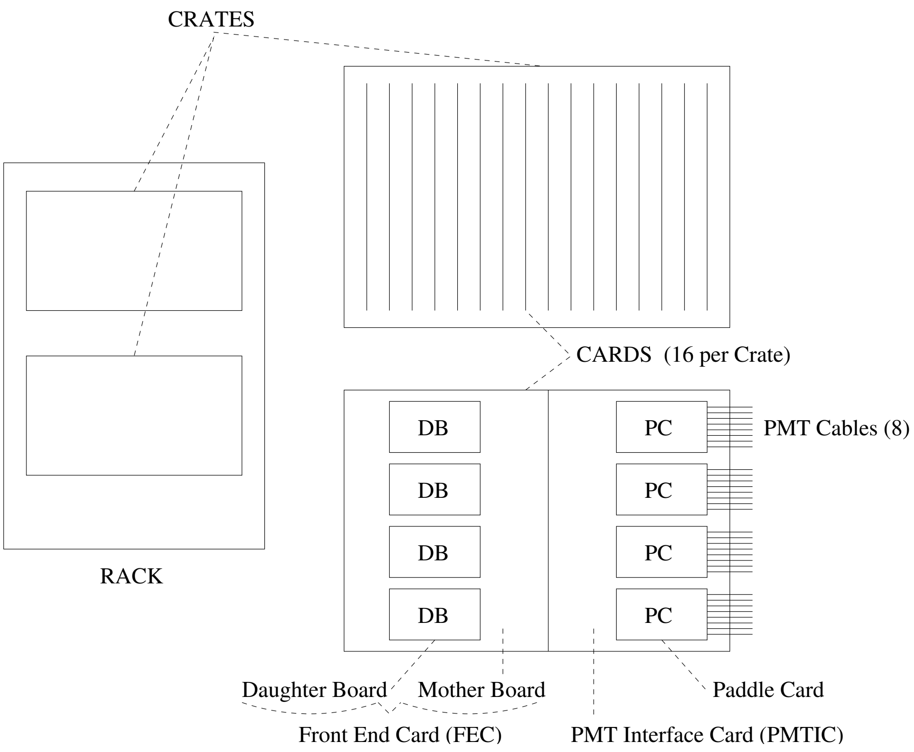

Figure 2.4: A schematic of the SNO electronics.

2.2 The SNO Electronics

The job of the front end electronics is to trigger on interesting events in the detector and to collect the firing times and integrated charges of all the PMTs involved. Once a trigger condition has been met, the data from all the PMTs involved is read out, grouped together and written to disk and tape for later analysis. The layout and the most common numbering scheme of the electronics will be described prior to explaining their basic operation.

The SNO electronics [32][33][34] are situated on top of the deck support structure, beneath which the detector is suspended. High voltage cables run from the electronics racks down to the PMTs below. A schematic of the SNO electronics is given in Figure 2.4. The electronics are divided into nineteen *crates*. Each crate occupies one level of a two level electronics rack. Within each crate are sixteen *cards*. A total of 32 PMT cables can be connected to a card. To process the signals coming from each cable a card contains 32 separate electronics *channels*. A card is a combination of the Front End Card (FEC), which does the signal processing, and the PMT Interface Card (PMTIC), which supplies the high voltage. Mounted on each PMTIC are four paddle cards and mounted on each FEC mother board are four daughter boards. The paddle cards and daughter boards deal with the same eight signal cables. This subdivision of 8 channels will be identified by its *paddle card* number. The electronics channels within a card are numbered from 0 to 31 from bottom to top. The paddle cards are numbered from 0 to 3 with paddle card 0 dealing with channels 0–7. The cards are numbered from 0 to 15 within a crate and the crates from 0 to 19. An electronics channel is therefore uniquely identified by its crate, card and channel number. These three numbers can be combined into one *logical channel number* or *lcn* using the formula

$$\mathrm{lcn} = 512 \times \mathrm{crate} + 32 \times \mathrm{card} + \mathrm{channel} \tag{2.1}$$

The lcn therefore runs from 0 to 9727. The electronics racks are arranged in a circle around the deck and the PMTs for which a given crate is responsible form an approximate slice in phi of the detector.

The electronics crates are shown in the SNO event display. Figure 2.5 shows the event display for a neutrino event. The upper left window shows the distribution of hit PMTs on the geodesic PMT support structure. The upper right window is a projection of the front and back hemispheres of the detector. The lower left window shows a colour coded histogram of the PMT hit times. The lower right window is the electronics display. The boxes labelled 0–19 are the electronics crates. The cards are arranged in columns and are numbered from left to right. The channels are arranged in rows and are numbered from bottom to top. The coloured squares in the electronics display show the electronics channels to which the firing PMTs are connected.

The SNO electronics use two oscillators: one keeps track of absolute time and the other gives high precision relative time. Absolute time is kept by a commer-

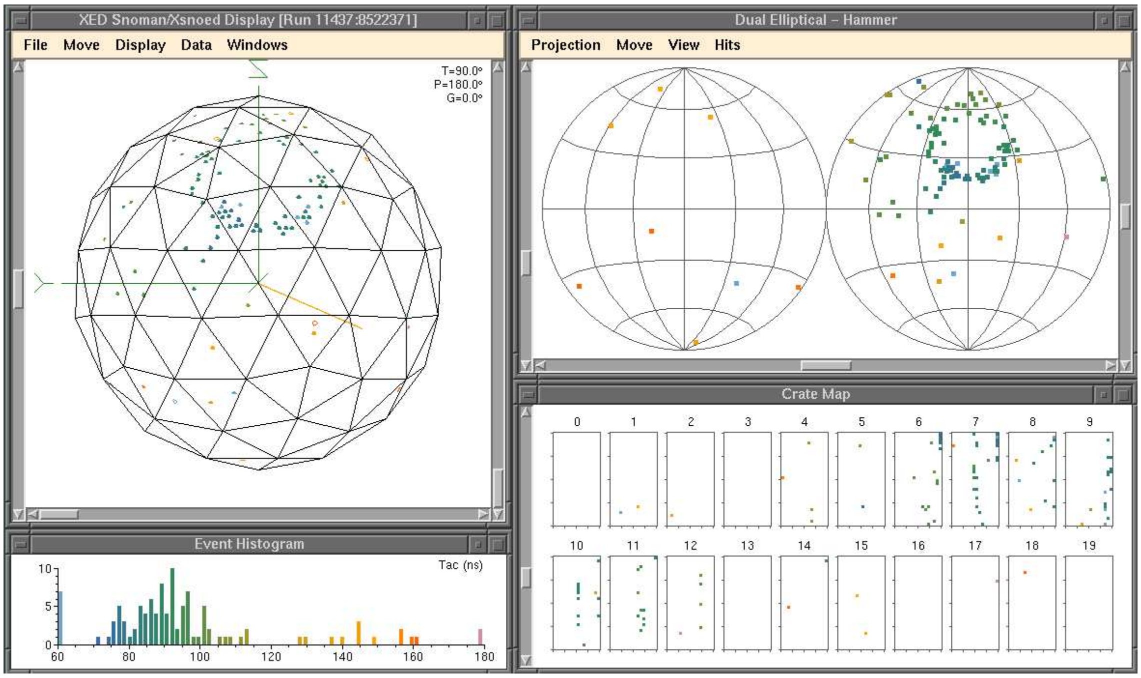

Figure 2.5: A neutrino event (Run 11437 GTID 8522371).

cial GPS receiver which provides a 10 MHz signal and precision time markers at requested Universal Times. The clock is accurate to 100 ns and is necessary for astronomical calculations and comparison with other detectors. The GPS receiver is on the surface and is connected to the detector via a 4 km cable; the resulting cable delay is corrected for. High precision inter-event timing is provided by a 50 MHz quartz oscillator. Pulses from this clock are stored on a 43-bit register which therefore rolls over every 2.036 days.

SNO is designed to detect the Čerenkov photons produced by neutrino interactions. The spread in the arrival times of these photons at the PMT array is determined by the different path lengths from the interaction point to the PMTs. A powerful hardware trigger is therefore to count the number of PMTs which register a photon, within a space of time that allows for the different path delays, and to issue a trigger when the count exceeds a pre-defined threshold. The PMT array is a 17 m diameter sphere so the arrival time of direct photons can vary by up to 77 ns. Whenever a PMT detects a photon the corresponding electronics channel issues an 100 ns long current pulse. An analog sum of these current pulses is therefore proportional to the number of PMTs which have fired in the previous 100 ns. The output of this analog sum is compared to a programmable threshold to look for contemporaneous hits above some number. This trigger is known as the NHIT100 trigger.

The electronics feature 26 different types of trigger, including the NHIT100 trigger. The operator can set a trigger mask which specifies which trigger types will be allowed to generate a global trigger. When an allowed trigger condition is met, a global trigger signal is sent to all front end crates on the next rising edge of the 50 MHz clock. If the crates contain any channels which have fired in the last 400 ns, their time and charge information is then read out and stored. All trigger conditions that are satisfied within a window of $\pm$10ns, whether they are masked out or not, cause bits to be set in a trigger word which gets written out with the other event level information. Once a global trigger has been generated there is a lockout period during which the generation of other global triggers is suppressed. If a trigger condition in the trigger mask is met during the lockout period then the inaptly named 'missed trigger' bit is set in the trigger word (a better name would be 'late trigger'). If any trigger is still above threshold after the lockout period then a new global trigger is immediately generated so no trigger conditions are ever truly missed.

To understand the nature of the PMT time and charge information which is recorded it is important to understand what happens in the electronics at the channel level. Figure 2.6 illustrates the state of the channel electronics following a PMT signal. Each channel has a discriminator which fires in response to a PMT signal above a chosen threshold. Once the discriminator has fired a voltage starts to ramp in a time to amplitude converter (TAC), labelled 'Time' in the diagram, and the signal from the PMT is integrated to get the total charge in the pulse, labelled as Q. If a global trigger is generated then the time and charge signals get stored on capacitors. If no global trigger is generated, then after about 400 ns, the

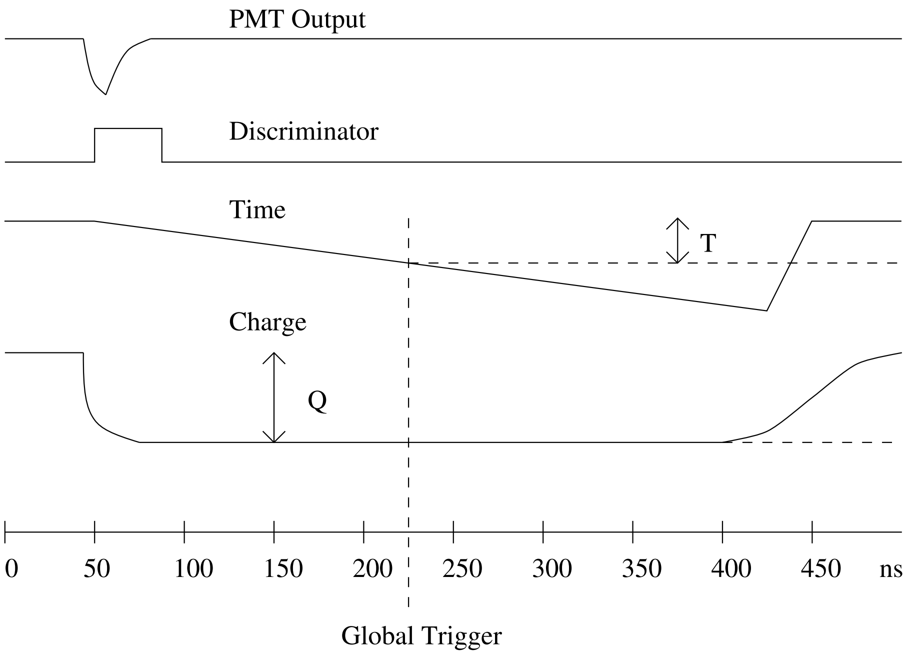

Figure 2.6: Single channel timing cycle. If no global trigger gets set then a channel resets after approximately 400 ns.

channel electronics reset and wait for the next firing of the discriminator. To be able to handle high PMT firing rates there are in fact 16 sets of storage capacitors, known as *cells*. These cells are numbered from 0 to 15 and are used sequentially. There are two TACs per channel which are used alternately and so, because the cells are filled sequentially, the data from one TAC is always stored on the even cells while the data from the other TAC always stored on the odd cells. This ordering is strictly maintained so that the electronics calibration for a particular cell is not complicated by having two TACs. There are actually three different measures of the charge which are recorded along with the TAC: the high gain short integrate charge (QHS), the high gain long integrate charge (QHL) and the low gain charge (QLX). The low gain charge could be set to either long or short integrate but is now always set to long. The short and long integration times are 70 ns and 400 ns respectively. The ratio of the high to low gain charge is approximately 12:1. The high gain short integrate charge (QHS) is the best measure of the size of the pulse that fired the discriminator while the long integrate charge includes all the reflected and scattered light. The time and charge information stored on the cells are read out by 12 bit (0–4095) analogue to digital converters (ADCs). There are four ADCs on each card, one for the time and one for each of the charges. At very high light levels QHS and QHL can reach the limit of their dynamic range and QLX is the only measure of the total amount of light. These ADC values constitute the raw time and charge data which gets written out to tape.

2.3 Electronics Calibration

The first step in the calibration of the raw ADC times and charges is the electronics calibration (ECA) [35]. All channels can be artificially fired using special 'pedestal' pulses. Once the channels have been fired, a global trigger is sent out and the data collected. The duration of the pulses and the delay between the pulse and the global trigger can be varied. Charge can be injected into a channel when it is fired, with the total amount of charge injected being proportional to the pedestal pulse duration. There are three types of calibration handled by the ECA: the pedestal calibration, the charge slope calibration and the time slope calibration. These are discussed below. The results of each calibration are written out to 'titles banks', a kind of data file, for later use in calibrating neutrino data.

The purpose of the pedestal calibration is to find the number of ADC counts that corresponds to zero charge for every cell of every channel. A time 'pedestal' is also measured but is only used for diagnostic purposes. No charge is injected into the channel during a pedestal calibration and no time delay is necessary. Every cell is fired 10 times and the median number of ADC counts is taken as the zero, a typical pedestal value is 500 counts.

In a charge slope calibration, different amounts of charge are injected into the channels by turning on the charge injectors and varying the pedestal pulse duration. The amount of recorded charge, in units of ADC counts, can then be plotted against the amount of injected charge and the slope found in arbitrary units. For the same pedestal pulse duration, different charge injectors supply different amounts of charge. Also the amount of charge that a typical photo-electron produces varies from PMT to PMT. It so happens that the charge slope of different cells of the same channel are almost identical. The channel to channel variations and the lack of variation from cell to cell means that in practice the charge slope calibration provides no more information than the number of ADC counts above pedestal. Consequently not many charge slope calibrations have been taken.

The time slope calibration uses the adjustable delay of the global trigger from the time of the pedestal pulse. 10 samples are taken at 31 (previously 19) well defined delays spanning the range of the TAC. Figure 2.7 shows, for one particular cell, the median of the 10 samples for each of the delays. These points do not form a perfect straight line and so an interpolation scheme is used to convert ADC counts into nano-seconds. The TAC turns over and goes double valued at the end of its range. All double valued regions are ambiguous and therefore do not have a unique calibration. This region only concerns very early hits and therefore has a negligible impact on analysis.

2.4 ECA Validation

A small fraction of electronics channels exhibit some degree of bad behaviour. A great deal of progress has been made in understanding these channels since the PMT calibration was first developed. A variety of cuts can now be applied to the data to remove rogue channels from the analysis. The calibration constants can be tested to see if they lie in an acceptable range. Channels can be cut if their constants have changed since the previous calibration. Non-linearities in the ADCs were a common problem at the time production data taking first got under way. There are now a range of tests for identifying and removing bad ADCs. A comprehensive description of all these cuts may be found in [35]. Unwelcome channel properties

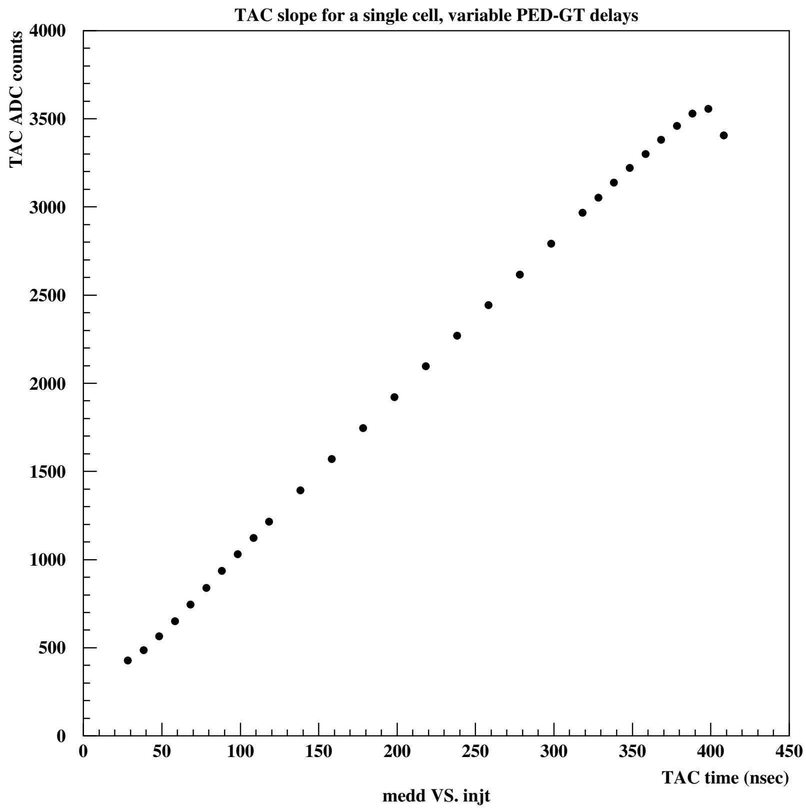

Figure 2.7: The time slope for a single cell showing the 31 different PED–GT delays. Note the increased sampling of the two non-linear regions at either end of the TAC range. Taken from [35].

will be described as they crop up throughout this thesis so it is not necessary to list them all here.

2.5 The Laserball

The laserball [36] is the light source used for the PMT calibration and is also the principle source for the optical calibration. The light is generated by a pulsed nitrogen/dye laser located in the deck clean room. The laser produces light pulses with a wavelength of 337.1 nm and a duration of 0.6 ns at a pulse rate of up to 50 Hz. Light can be produced at five wavelengths in the range 337–650 nm by using the laser directly or by pumping one of four dye cells. The intensity of the beam can be adjusted using neutral density filters mounted on remotely controlled selector

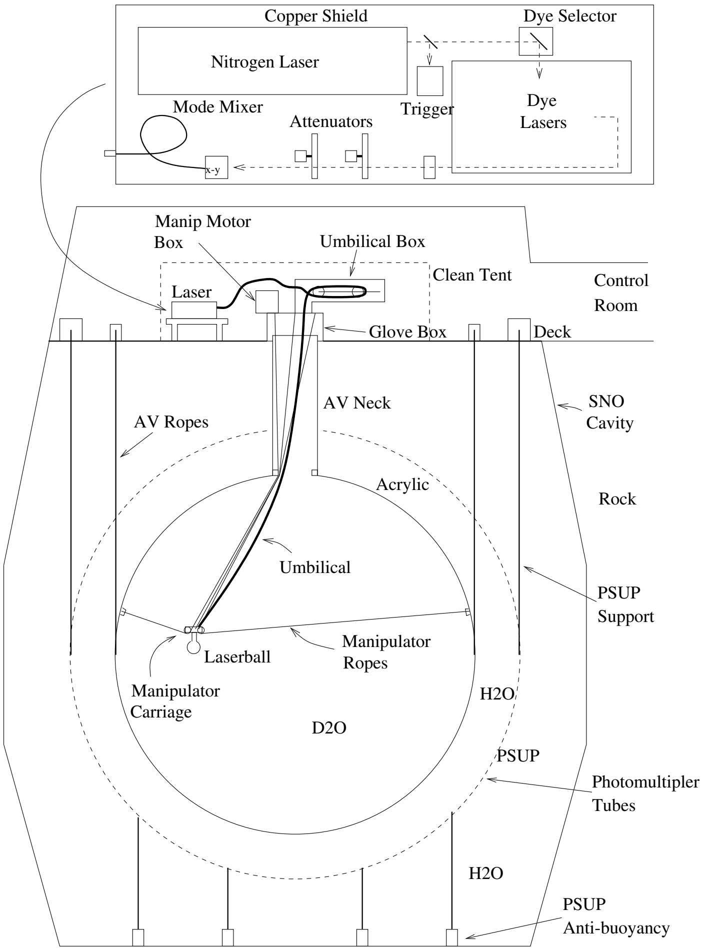

Figure 2.8: The laserball and manipulator system. Taken from [36].

wheels. The beam is focused into a bundle of optical fibres which lead to a 10 cm diameter diffuser ball. The light from the diffuser ball has a duration of less than a nano-second. It is isotropic to about 10% [37] although the intensity drops off towards the top of the ball.

The laserball can be moved to a wide variety of positions inside the detector courtesy of a 'manipulator' system. The manipulator consists of a carriage, on which sources can be mounted, suspended by ropes. By drawing in or letting out the ropes, the carriage can be moved to the desired location in the detector. Figure 2.8 illustrates the laserball and manipulator system.

As previously described, a global trigger is not normally issued until the next rising edge of the 50 MHz clock, giving rise to a variable delay of up to 20 ns. This arrangement enables very precise inter-event timing because there is always an integer number of clock pulses between global triggers. If there is a 20 ns variation between the initiating trigger and the global trigger there is also a 20 ns variation in the typical PMT times relative to the initiating trigger. For most analyses this is unimportant as the time of the event is found by reconstruction. This does not work when calibrating the PMT times because the reconstruction depends on the calibration. The PMT calibration requires collecting the PMT timing distributions from the laserball, accumulated over many events, without the 20 ns spread. This problem is solved by a special trigger, called the external asynchronous trigger, which immediately issues a global trigger, regardless of the state of the 50 MHz clock. The very first stage in the laser box is to split some of the beam to a photo-diode. The signal from this diode is used as the input to the external asynchronous (EXTASY) trigger after an adjustable delay. In this way the sharp timing of the laserball is preserved when summing the timing distribution over many events.

2.6 The ^{16}N calibration source

^{16}N is the principle energy calibration source used in SNO [38]. ^{16}N beta decays to ^{16}O with a half life of 7.13 seconds. The most common decay mode produces a beta particle and a 6.1 MeV gamma ray. The beta is used to generate a trigger while the gamma escapes the source and produces Čerenkov light in the detector. There is also a decay mode which goes straight to the ground state, producing an energetic beta particle but no gamma.

The ^{16}N is produced by the reaction ^{16}O(n,p)^{16}N. The neutrons are generated by firing a deuterium beam at a tritium target (D-T Generator). The oxygen is present in CO_2 gas surrounding the target chamber. A gas transport system moves the resulting ^{16}N into a decay chamber in the detector. The β is detected using a combination of scintillator and a photomultiplier tube.

In this thesis most of the ^{16}N plots were made selecting events where the number of PMTs which fire was between 20 and 150. The vast majority of ^{16}N data is in this range to begin with. The gammas from ^{16}N typically convert $\sim$60 cm from the source so the event vertex must be fitted in order to study the time distribution.

2.7 The ^{8}Li calibration source

In the Sun, ^{8}B β^{+}-decays to ^{8}Be. The isospin mirror nucleus of ^{8}B is ^{8}Li which also decays to ^{8}Be, this time by β^{-}-decay. Apart from the slightly lower Q value, the energy spectrum of the emitted electron is very similar to the energy spectrum of the ^{8}B neutrino. ^{8}Be breaks up into two alpha particles in 0.07 fs. The half life of ^{8}Li is 0.84 s.

The ^{8}Li for the calibration source [39] is produced by the reaction ^{11}B(n,α)^{8}Li, where the neutrons are generated by the D-T generator. The ^{8}Li is carried by a gas transport system to a decay chamber in the detector. Betas from the decay chamber escape into the D_2O where they produce Čerenkov light. The ^{8}Be breaks into two alpha particles which scintillate in the helium used in the transport gas. The scintillation light is detected by a PMT observing the decay volume and provides a tag for ^{8}Li decays.

The D-T generator also produces a small amount of ^{16}N background. For an event to feature in a ^{8}Li analysis the detector must trigger and the tag PMT must fire. The beta from ^{16}N decay can fire the tag PMT and the gamma can trigger the detector, producing a background. ^{16}N events can be distinguished from ^{8}Li event because the beta triggers the tag PMT faster than the alphas and the PMT collects less charge. All ^{8}Li analyses in this thesis are filtered to only contain the ^{8}Li beta decays. The data sets used are the same ones generated for a data cleaning analysis [40]. The electrons from the ^{8}Li source generate Čerenkov light as soon as they enter the water. The event location is therefore well known and this is exploited in chapter 7 to take a brief look at how much the PCA improves reconstruction.

2.8 SNOMAN

SNOMAN is the SNO Monte Carlo and ANalysis software package [41][42][43]. The Monte Carlo part of SNOMAN models the generation and propagation of all relevant particles though the geometry and media of the detector. The generation of Čerenkov light and its detection by the PMTs is modelled as are the resulting detector triggers. The analysis part of SNOMAN can perform specified analyses on both Monte Carlo data and real SNO data files. Results are usually stored in an HBOOK file for inspection with PAW [44]. SNOMAN is primarily written in FORTRAN, with dynamic memory allocation being provided by ZEBRA [45].

SNOMAN consists of a collection of largely autonomous processors working on a common data structure. SNOMAN is controlled by a combination of command files and data files. The command files, amongst other things, specify which processors are to be run and in what order. Processors are identified by a three letter name. The processor which generates the Electronics CAlibration constants is called the ECA. The processor which generates the PMT CAlibration constants is called the PCA. These calibration constants are written to SNOMAN data files, known as 'titles banks'. These titles banks may be stored in the SNOMAN database (SNODB) [46]. The processor which handles the calibration of SNO data is called CAL. When an event is being calibrated, the pertinent calibration constants for that date are retrieved from the database. Part of the job of the calibration processor is to apply the PCA constants to the data.

Chapter 3

Introduction to the PCA

The topic of this thesis is the PMT CAlibration (PCA). The PCA has two main components. One part is concerned with providing the best possible timing calibration of the PMTs and the other part improves the charge calibration of the PMTs. The part concerned with the timing calibration, the walk calibration, is discussed in section 3.1. The part dealing with the charge calibration, the gain calibration, is described in section 3.2.

3.1 The Walk Calibration

The task of the walk calibration is to give the relative times at which PMTs fire in an event to well within their 1.5 ns time jitter. The walk calibration is very important to neutrino analysis because it directly affects event reconstruction (this, in turn, affects the energy assigned to each event because the energy estimator uses the reconstructed position and direction). The purpose of SNO is to infer the makeup of the neutrino flux through the detector by measuring the rates of the various neutrino interactions discussed in section 1.6.2. The different interactions have different directional, angular, radial and energy characteristics which means that the contribution of each interaction to the total signal can be determined. In order to study these characteristics it is essential to be able to reconstruct where each interaction took place in the detector. The interaction vertex can be fitted

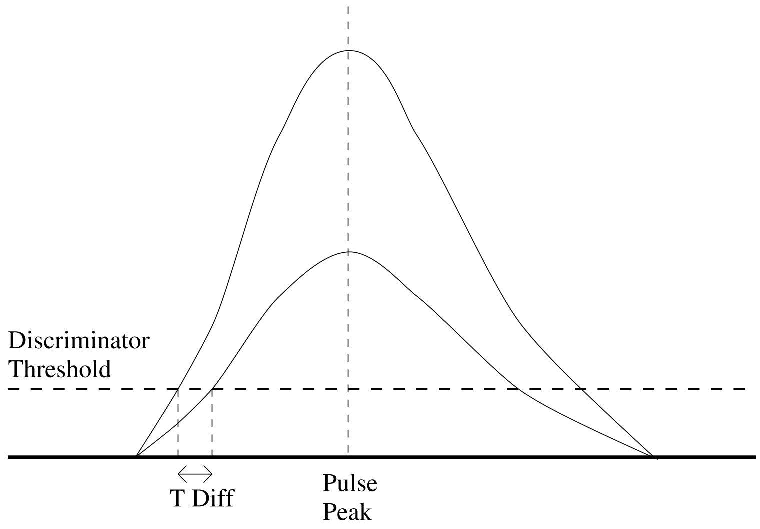

Figure 3.1: Small pulses take longer to cross the discriminator threshold with respect to the centre of the peak than large pulses.

using the relative firing times and locations of all the struck PMTs. A knowledge of the relative timing is therefore essential to analysis.

As was discussed in section 2.2, the raw charge and time data is recorded in ADC counts. The ECA calibration, described in section 2.3, converts the time in ADC counts to time in nanoseconds and the charge in ADC counts to charge in ADC counts above pedestal (cap). Although the ECA gives the time at which the discriminator fires in nanoseconds, there are unknown offsets between the ECA times of different PMTs, introduced by their different cable delays. The reconstruction of events in the detector is based on the relative times of PMT hits and so the cable delays must be corrected for to within a global offset.

There is a second effect which the walk calibration must address; the time at which the discriminator fires due to a PMT pulse is not independent of the size of that pulse. The size of a pulse produced by a single photoelectron is subject to statistical fluctuations. Small pulses take longer to cross the discriminator threshold with respect to the centre of the peak than large pulses, this effect is illustrated in figure 3.1. There is, therefore, a systematic change with pulse size of the time at

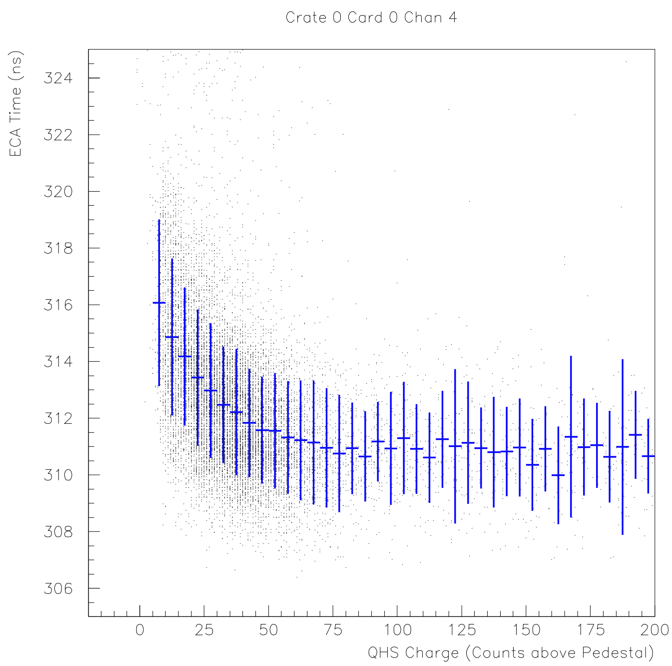

Figure 3.2: An example of discriminator walk. It shows ECA time against QHS charge for data collected from an external asynchronously triggered laserball run for crate 0, card 0, channel 4. The blue crosses are a profile histogram of the data with the vertical bars showing the RMS spread. The size of the walk is about 5 ns.

which the discriminator fires relative to the time at which the photon struck the photocathode. This effect is known as discriminator 'walk' and it must be corrected for in order to achieve the maximum possible timing resolution. Discriminator walk can be seen by plotting ECA calibrated time against QHS for an external asynchronously (EXTASY) triggered laserball, see figure 3.2. QHS is the most relevant charge for the walk calibration because it is the size of the pulse which fired the discriminator which is important and not the size of any subsequent photon pulses.

Both the cable delays and the discriminator walk can be removed by a single calibration. With the laserball placed near the centre of the detector, the delay between the time of the laserball pulse and the time the PMTs are struck is easily calculated. The location of the laserball is given by the manipulator system. The positions of all the PMTs was recorded during construction. Knowing these positions and the group velocity of light in D_2O and H_2O the light travel time from the laserball to each PMT can be calculated. Recording when the PMTs fire in response to the laserball pulse and knowing the light travel times provides the information needed to line up all the PMT times to within a global offset. The discriminator walk can be simultaneously removed by making this cable delay correction at the same charge as the charge of the data to be calibrated. The PCA calibrated time of a PMT hit in neutrino data, $\mathrm{TPCA}_{\mathrm{data}}$, is therefore given by:

$$\mathrm{TPCA}_{\mathrm{data}} = \mathrm{TECA}_{\mathrm{data}} - \mathrm{TECA}_{\mathrm{laserball}}(\mathrm{QHS}_{\mathrm{data}}) + \mathrm{path\_delay} \tag{3.1}$$

where TECA is the ECA calibrated time and path_delay is the light travel time from the laserball to the PMT in question. Chapter 4 goes into some detail about how the laserball time as a function of charge is recorded. The performance of the walk calibration is investigated in chapter 5 and, in particular, the systematics of the optical path correction is considered in section 5.2.

3.2 The Gain Calibration

The ECA calibration gives the charge in units of ADC counts above pedestal. The number of ADC counts above pedestal that corresponds to one photoelectron varies from PMT to PMT. The purpose of the gain calibration is to convert the charge into units proportional to the number of photoelectrons. Solar neutrinos produce hits at the single photoelectron level so the integrated charge is of little interest to reconstruction or separation of different event classes. The gain calibration will be important for estimating the effect of changing the discriminator threshold on the relative phototube efficiency and therefore on the energy scale of the detector (The energy scale is normally defined as the typical number of hits produced by an electron per MeV of kinetic energy). The gain calibration will also be useful when studying higher energy events like cosmic rays and atmospheric neutrinos.

The channel to channel variation in the number of counts above pedestal that correspond to one photoelectron (pe) comes from two sources, variation in the true gains of the PMTs and variations in the components of the electronics (the charge integrators and ADC converters). The gain of individual PMTs could change over time because of ageing effects like an increase in the amount residual gas in the PMT. It is also possible that the voltage across the PMT could drift, leading to an alteration of its true gain. The apparent gain of a channel could also change if some of its electronics cards are swapped out for repair. To limit change over time the voltages and electronics are kept as stable as possible.

The most obvious way to find the number of counts above pedestal (cap) that corresponds to 1 pe is to collect a charge spectrum for each PMT from a laserball run, taken at a low light level, and look for a peak. Unfortunately the charge spectra are cut off at low charge by the discriminator and, as will be shown, this interferes with finding the peak. There is a method for collecting the full charge spectrum of the PMT without the discriminator threshold interfering: when the laser pulses, and the trigger photodiode in the laser box gives a signal, a desired set of channels are all artificially fired regardless of whether they detect a signal above the discriminator threshold. This kind of laserball run is known as a 'gated' charge run. Because every channel is read out after every event, gated charge runs are very intensive for the data acquisition system. To collect reasonable statistics on a single channel either the event rate must be reduced and the laserball run for a very long time or the number of gated channels must be restricted. Figure 3.3 shows the gated charge spectrum for crate 0, card 0, channel 4. The expected features for a PMT charge spectrum can be identified. There is a spike at low charge where no photoelectrons were produced and so either no charge was collected or only a small amount from noise. There is a dip between the noise spike and the single photoelectron peak. The single photoelectron peak has a long tail to high charge due to the poor charge resolution of the PMT and some multi photoelectron events. Gated charge runs

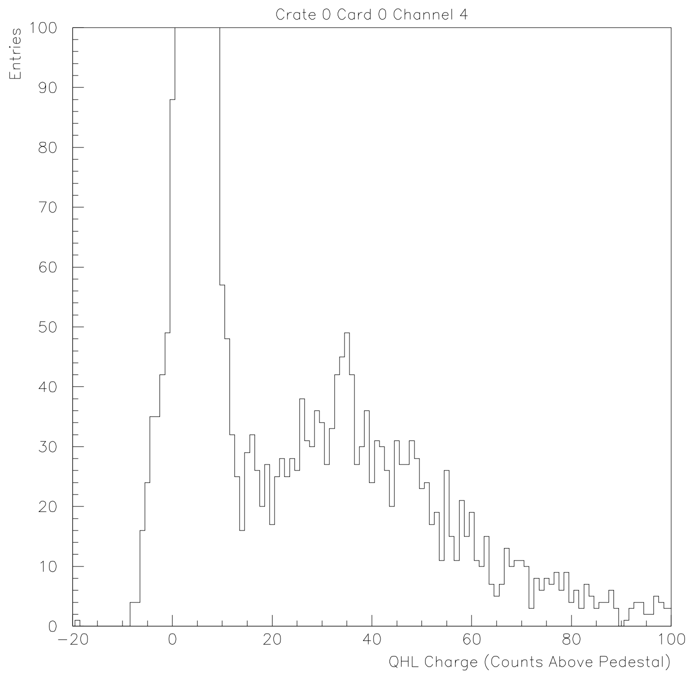

Figure 3.3: The gated charge spectrum from a laserball run for crate 0, card 0, channel 4.

are not a practical way of regularly monitoring the gains of the PMTs owing to the length of time they take. A method is required which can determine the gain using discriminated data and which be checked against gated charge data.

Figure 3.4 shows the discriminated charge spectra for nine selected channels collected from a laserball run. The blue lines in the figure indicate the approximate level of the discriminator (known from the electronics settings). The discriminator threshold is actually set in DAQ counts and not ADC counts but, as will be demonstrated in section 4.7, it happens that the two very similar. The spike at zero charge is an unexpected additional feature and varies in size from PMT to PMT. It is almost certainly 'pickup' noise, caused by the channel acting like an antenna and generating a bipolar signal. This signal can be large enough in amplitude to fire the discriminator but will integrate to very little. Discriminator firings due to pickup seem to be particularly prevalent during laserball runs. The gas in the laser

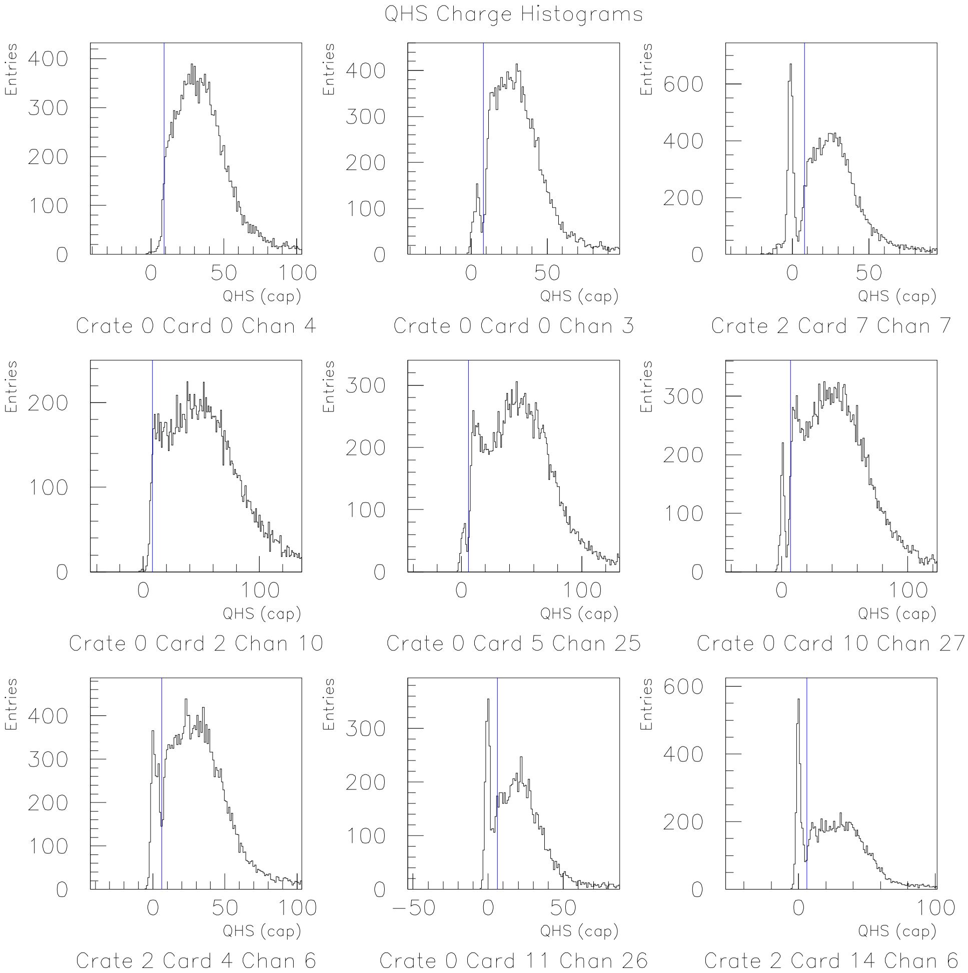

Figure 3.4: Discriminated charge spectra from laserball run 11326 for nine selected channels. The blue lines indicate the level of the discriminator threshold in DAQ counts. See the main text for a discussion.

is excited by high voltage discharges and so is an obvious candidate for causing the noise, even although the laser is shielded by a copper box. The true cause of the noise has not yet been firmly established. As will be seen in the next chapter this low charge noise also has to be taken into account in the walk calibration.

The top three channels of figure 3.4 show typical charge spectra with the amount of zero charge noise increasing from left to right. The middle three channels are high enough gain to clearly show the single photoelectron peak. (These channels all have a blown 75 Ω resistor in their bases and this raises their gain a little. These resistors match the impedance of the PMT cable and are there to prevent reflections. The resistors were blown during the period when the detector was first filled with water. At this time the PMT connectors suffered from high voltage breakdown, due to a marginal design. The connectors have a small air gap at the connection. When submerged in degassed water, the air around the connection tends to diffuse out leaving a lower pressure in the connector and therefore a lower breakdown voltage. This hypothesis, and the solution of regassing the water with nitrogen gas was suggested by David Sinclair—saving the experiment.) In the bottom three channels the noise and the discriminated charge spectra are not well separated. These channels mimic the form of the gated charge spectrum but it is important to realise that the dip between the noise and the spectrum is determined by the discriminator.

On many channels, whether they suffer badly from noise or not, it can be difficult to tell whether the low end of the single photoelectron peak is part of a genuine decrease or is caused by the discriminator. One solution to this problem, identified by Bill Frati and pursued in chapter 4, is to use the 'high half' point. The high half point is found by starting at the peak of the charge spectrum and searching to the right hand side for the point which falls to half the height of the peak. This point is very robust to the level of the discriminator as the discriminator would have to be set very high to greatly affect the size of the highest bin. It will be shown in section 5.3 that this point is fairly insensitive to the intensity of the laserball.

One of the future uses of the gain calibration will be to estimate the effect of changing the discriminator threshold on the relative efficiency of the PMT. The first stage in this will be to derive an idealised single photoelectron charge spectrum from a gated charge run, excluding the noise spike at low charge. It will then be possible to derive the fraction of photoelectrons which are detected as a function of discriminator threshold (expressed as a fraction of the high half point). This is one of the factors which affect the energy scale of the detector.

Chapter 4

The Implementation of the PCA

This chapter explains how the PMT calibration is implemented in SNOMAN. The SNOMAN processor which generates the PMT CAlibration constants is named the PCA. The PCA provides a general framework in which calibrations derived from laserball data can be performed. There are currently three calibrations supported by the PCA: the walk calibration, the gain calibration and a second, improved version of the walk calibration. Section 4.1 describes how the PCA is controlled and outlines its general structure. The nature of the laserball data which the PCA uses is reviewed in section 4.2. The walk, gain and extended walk calibrations are described in sections 4.3, 4.4 and 4.5. These sections are quite detailed and are intended to double as documentation of the method used. How the resulting walk and gain constants are applied by the calibrator is discussed in section 4.6. Sections 4.7 and 4.8 are concerned with the problems of bad pedestals and bad ADCs.

4.1 The PCA Framework

The PCA is controlled by a command file and by specific titles banks. Example command files for running walk and gain calibrations are available with SNOMAN. Titles banks are identified by a four character name along with a number called the bank number. The set of titles banks which control the PCA have the name TPCA, where the 'T' stands for titles. TPCA bank number zero (TPCA 0) is the overall PCA control bank and specifies things of a general nature like which type of calibration is to be performed. Titles banks TPCA 1, TPCA 2 and TPCA 3 contain parameters specific to the running of the gain, walk and extended walk calibrations respectively. Once a walk or gain calibration is completed the resulting constants are written out to another set of titles banks for later use by the calibration processor, CAL. More detail on the contents of these banks is given in sections 4.3, 4.4 and 4.5. How these constants are finally used by the calibrator will be explained in section 4.6.

SNOMAN has three main phases of running: initialisation, execution and termination. When the PCA receives an initialisation call it reads the PCA control titles bank, TPCA 0, and then initialises the specified calibration. During the execution phase the PCA maintains statistics on the date and time range spanned, the number of extasy triggers and keeps a running average of the position of the laserball. For each event, the PCA calls the top level execution routine of the current calibration type so that it can collect the information that it needs from the data structure. On termination, the PCA calls the termination routine of the current calibration type and then prints out a brief summary of the data processed. Most of the work of a calibration is carried out within the specific routines of the calibration being performed. What these routines do is the topic of the following sections. First the type of laserball run which the PCA uses will be discussed along with the types of PMT that can be calibrated using this data.

4.2 The PCA Laserball Run

Currently the same laserball run is used to derive the walk and the gain calibration constants. The laserball is moved close to the centre of the detector. This simplifies the calculation of the time it takes for light to travel directly to the PMTs. If the laserball were moved off centre then the refraction that occurs at the D_2O/acrylic/H_2O boundaries might become important. The laserball is isotropic to about 10% [37] so placing the laserball near the centre also means that all PMT will receive similar statistics. The intensity of the laserball is adjusted so that the fraction of PMTs that fire in response to a single laserball pulse is between 5 and 10 percent. The intensity must not be raised above this level to avoid PMTs receiving a high fraction of multiple hits. The effect of multiple hits will be considered in section 5.3. A wavelength of 500nm is used for reasons that will be given in section 5.1. To satisfactorily sample the tail of the charge spectrum the aim is to collect around 10,000 hits per channel. To do this the laserball is typically run at 40 Hz for over an hour.

As was described in section 2.1, there are five types of PMT in SNO: normal, OWL, low gain, BUTT and neck. The principles of walk and gain apply to all these types of PMT but the only type of tube which the laserball adequately samples are the normal tubes. The low gain tubes are inside the detector but require very high levels of light before they fire. The OWL, BUTT and neck tubes are not illuminated by the laserball. The outward looking PMTs are mounted on the outside of the PSUP, the butt tubes are outside the detector on the cavity wall and the neck tubes are up near the pipes at the top of the neck of the acrylic vessel and are not well illuminated by a source from inside the detector. Normal tubes constitute the vast majority of PMTs and are the only type of tube which is relevant to the reconstruction of neutrino data.

4.3 The Walk Calibration

The task of the walk calibration is to collect and parameterise the form of the discriminator walk for all normal channels. Initial attempts to parameterise the shape of the walk focused on trying to fit some functional form to the data. An exponential decay worked to first order:

$$t(q) = p_1 + p_2\, e^{\mathrm{QHS}/p_3} \tag{4.1}$$

where p_1, p_2 and p_3 are to be determined. This method of parameterising the form of the walk was abandoned because a single functional form failed to mould itself to the variety of shapes thrown up by the 9453 normal PMTs. A similar problem arose in the time slope calibration of the ECA. Steve Biller, when writing the time slope calibration, got around the problem by using a set of interpolation points. It was decided that interpolation points would also be a good way to parameterise the discriminator walk. In order to adequately follow the shape of the walk function, the PCA uses 10 interpolation points. There is a huge amount of data in a PCA laserball run and simply reading it in takes a considerable amount of time. It is therefore desirable that any calibration should only require a single pass over the data. The way in which these 10 interpolation points are chosen for every normal PMT in the detector on a single pass will now be described.

SNOMAN has a considerable amount of memory available for use by its processors. As none of the other processors which are used during a walk calibration need very much in the way of memory the PCA can reserve the bulk of this memory for itself. On initialisation the walk calibration reserves three blocks of memory for use in the execution phase. One block contains channel level statistics, another block contains the interpolation point level statistics and the last block is a buffer for storing time and charge data. The time and charge data buffer is by far the largest of these blocks of memory with 1001 words of memory for every channel in the detector.

To begin with, the time and charge data buffer (tq buffer) is filled with zeros. The 1001 words of memory allocated to a given channel are blocked together with the blocks being ordered by logical channel number (lcn, defined in equation 2.1). The first word of the channel memory block is a flag indicating the current state of the block and is set to zero to begin with. The second word is an integer counter which records how many time/charge points have been stored in the rest of the buffer so far. To be able to buffer as many points as possible, instead of recording

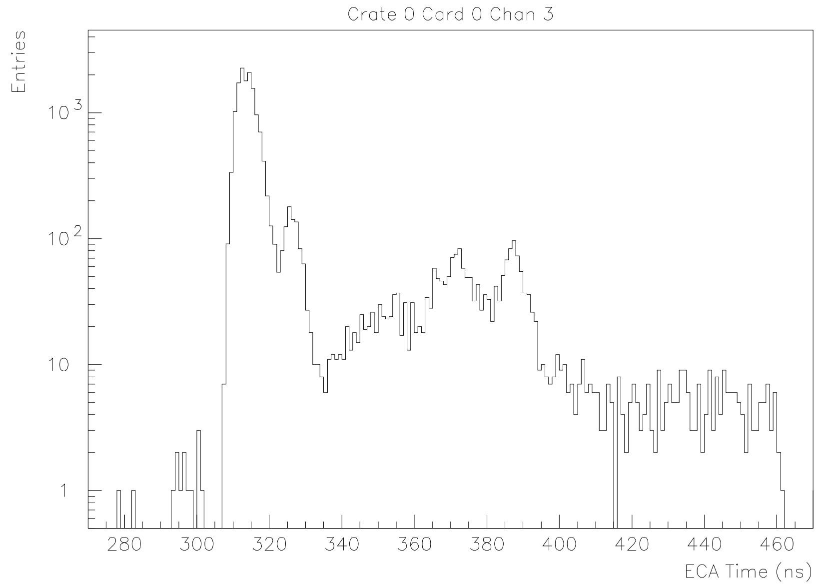

Figure 4.1: Plot of the ECA calibrated time collected by a single channel.

the ECA calibrated time and charge, which would take two words, the cell (4 bits), ADC time (12 bits) and QHS ADC charge (12 bits) are packed into a single word (32 bits) which can be calibrated as required. The memory ordering for a channel in state zero can therefore be summarised as

```
lcn-1:                                            ...TQDATA(999),
lcn  : STATE, N_TQDATA, TQDATA(1), TQDATA(2), ...TQDATA(999),
lcn+1: STATE,...
```

Once 999 charge time points have been collected, and the buffer is full, the contents are examined.

Figure 4.1 shows a typical ECA time spectrum collected by a channel. The largest peak is due to light travelling directly to the tube. Early hits are from random background and the tail to late times is due to scattered and reflected light, c.f. figure 2.3. Note that the peak immediately following the prompt peak is much more prominent, this is due to pickup noise. The prompt peak is clearly the best peak to use to parameterise the walk and must be isolated. The time of the prompt

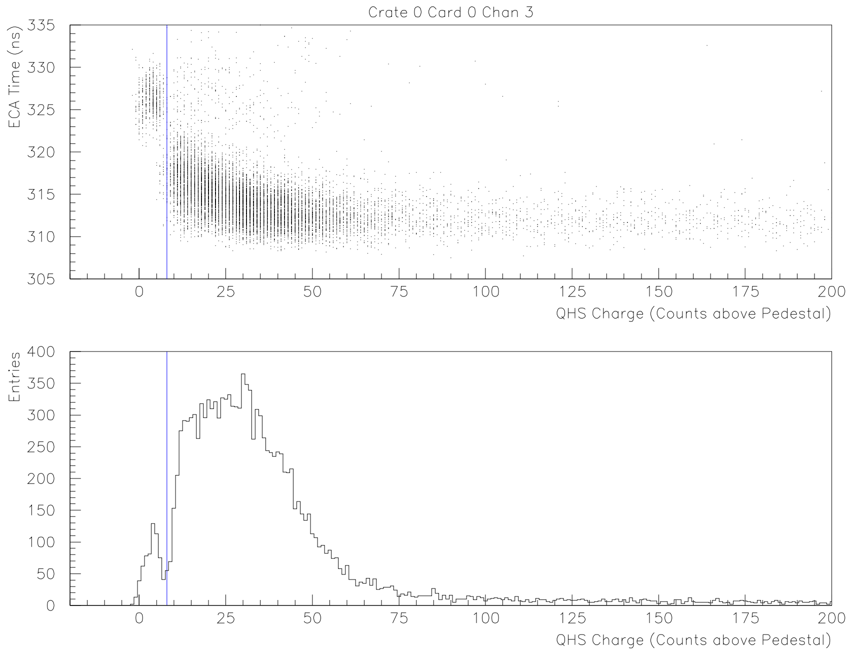

Figure 4.2: The upper plot shows the walk of a single channel. The lower plot shows the charge spectrum for the same channel. It can be seen that most hits are on the steep part of the walk. The blue line shows the position of the discriminator threshold in DAQ counts.

peak is found using a sliding window of length 8.5 ns and finding the average of the window with the most entries (8.5 ns was used to encompass the 5 ns walk and allow for the ~2 ns width of the laserball pulse). The ECA time of the prompt peak is then stored along with the rest of the channel level statistics.

The charge bins which will be used to form the interpolation points also need to be chosen. Figure 4.2 shows the walk of the prompt peak and its charge spectrum. It can be seen that the peak of the charge spectrum needs to be well sampled as it is both where the majority of hits come and where the walk is the steepest. The extent of the walk is strongly correlated to the width of the charge peak for all channels so the width of the charge peak can be used to choose the charge bins. The use of the charge spectrum is complicated by the noise near zero charge. It was felt that it would be better to avoid relying on the pedestal value to cut this noise. To meet these criteria the following method to choose the charge bins was adopted. First the shortest charge range containing 70% of the data is found to get the approximate position of the charge peak. This range could be biased by the ratio of the number of hits in the peak to number of hits in the tail. To solve this problem, any hits more than 50 counts above the upper range of the shortest span are dropped. Using the remaining data the shortest span containing 90% of the data is found and this is taken as an estimate of the range of the peak. Because the algorithm looks for the most clustered data, the lower limit has a strong tendency to exclude the noise near zero charge provided that there is a dip between the peak and the noise and the noise peak is not too strong. It is desirable, but not absolutely necessary, to exclude the noise from the walk parameterisation. Hits near zero charge are dominated by noise so there is no 'correct' value of the walk for these hits. When calibrating data, if the hit is too far below the first walk interpolation point in charge then it is rejected, this will be discussed further in section 4.6. The fraction of channels where noise is included in the walk parameterisation will therefore affect the fraction of noise hits which get rejected by the calibrator. Another cost of including noise in the walk parameterisation is using an interpolation point to sample irrelevant data. If a more reliable way to decide where the threshold lies were devised then it could be added to the walk calibration and/or the calibrator.

Once the lower and upper limits of the charge peak have been identified the charge bins then need to be allotted. Six points sample the main peak of the charge spectrum and four points sample the high charge tail. The range covered by the peak is split into 8 equal sections. The first four sections are given an interpolation point each because the walk is steepest there and so must be well sampled. The second four sections are divided between two interpolation points. The remaining four interpolation points span the range from 0–50, 50–100, 100–200 and 200–300 counts above the upper limit of the charge peak. The range of the chosen charge bins are stored along with the interpolation point level statistics.

The contents of the tq buffer are shifted to a temporary store and the buffer cleared. The buffer is rearranged as:

```
lcn-1:                                          ...TQDATA(99),
lcn  : STATE, N_BIN_1,  TQDATA(1), TQDATA(2), ...TQDATA(99),
              N_BIN_2,  TQDATA(1), TQDATA(2), ...TQDATA(99),
              ...
              N_BIN_10, TQDATA(1), TQDATA(2), ...TQDATA(99),
lcn+1: STATE,...
```

The state number is changed to 1 to indicate the new format. Different channels may be in different states depending on their firing rate. The data in the temporary store is then fed back into the buffer taking care to place the data in the correct charge bin. Data is only added to the buffer if it is within 10 ns of the stored prompt peak time. The walk typically changes by 5 ns from low to high charge and the RMS is typically less than 2 ns, so 10 ns in either direction from the overall centre safely encompasses all data of interest. As each of the charge bin buffers fill, the contents are read out, calibrated and statistics kept on the average charge, time and time squared. The median of the times is added to a running average of the medians. The walk calibration keeps accumulating statistics in this way until the data files are exhausted.

At the start of the termination phase the walk calibration processes the remaining buffered data and then drops the large tq buffer. The accumulated statistics are then processed to generate the walk data. The 10 charge bins are converted into 10 charge/time interpolation points for later use when calibrating data. The charge of the interpolation point is taken to be the average charge of the bin. The time of the interpolation point is the average of the median times for that bin. The median is more reliable than the average because it is less susceptible to being pulled by late hits from scattered light. Recall that hits are only incorporated into the bin statistics if they are within 10 ns of the overall peak. The times of the charge bins vary with charge because of the walk and so the amount of scattered light included in the bin statistics varies systematically with charge. As well as choosing the interpolation points, an attempt is made to record the extremes of the walk function at high and low charge. The time at low charge is the time of the first charge bin that contains at least 10 hits. The time at high charge is taken to be the time of the highest charge interpolation point provided that it contains at least 100 hits. If it does not contain 100 hits then more points are added, up to a maximum of three points.

The walk data is written to a set of titles banks with the name WPCA, where the W stands for walk. There is one bank per crate with the bank number matching the crate number. The bank consists of a 10 word header followed by the channel data. The first word of the header is a duplicate of the crate number. The second word is a version number to simplify the introduction of new versions. The remaining words in the header concern the format of the channel data. The main body of the bank has 12 words per channel. The first word is a summary word, followed by 10 words containing the 10 interpolation points in a packed format, followed by one word containing the value of the walk function at high and low charge, also in a packed format. The summary word is the only one which can be read directly. It is formed as

$$100 \times \text{tube type} + 10 \times \text{tube status} + \text{fit quality}$$

where tube type, tube status and fit quality are all integers. The tube type is 1 for a normal tube and > 1 for all other tube types. The tube status indicates whether the tube produced enough data to attempt a fit and the fit quality indicates whether the interpolation points are considered reliable. The values and meaning of the tube status and fit quality can be summarised as

	tube status		fit quality
0	passed	0	not fitted
1	offline	1	passed
2	zero occupancy	2	poor (bad)
3	low occupancy	3	bad
4	low ECA		

Zero occupancy indicates that the tube was supposedly online but did not fire. Low occupancy indicates that the tube was online but fired less than 100 times during the PCA laserball run. Low ECA indicates that the tube was online and did produce more than 100 hits but less than 100 of these passed the ECA calibration.

The fit quality indicates whether interpolation points were generated for that channel and if so whether they can be trusted. A fit is only attempted for tubes whose status is 'passed', otherwise the fit quality is 'not fitted'. The walk calibration may fail a bin or channel for a variety of reasons. Any bins which do not contain a minimum of 10 ECA calibrated hits are dropped from the interpolation. All other interpolation points are written out to the WPCA banks regardless of whether the channel passed or failed. There are four tests applied to the channel to decide the fit quality. The simplest test is checking that the approximate RMS of the channel is reasonable. The approximate RMS is calculated by combining the RMS of each of the bins, not correcting for the slope of the walk function across the bin. If this RMS is more than 2.8 ns the channel is marked as poor and if it is greater than 5.0 ns then the channels is marked as bad. The level of this RMS cut is returned to in section 4.8.1. The calibrator does not distinguish between poor and bad and both are rejected. The three other tests are designed to catch particular types of channel anomaly and will be explained in section 4.8.

A laserball titles bank is written out at the same time as the WPCA banks. The bank is named LRBP, from LaseRBall Position, and contains the laserball position and wavelength for that run. How the WPCA and LRBP titles banks are used in calibrating data will be reviewed in section 4.6

In addition to the titles banks, the collected calibration statistics are written out to a HBOOK ntuple for inspection in PAW. If desired, some fraction of the raw data can be written out to another ntuple. Figure 4.3 shows the collected raw data for crate 0, card 0, channel 3 with the interpolation points superimposed. Large numbers of these plots were scanned to make sure the algorithm was working as intended.

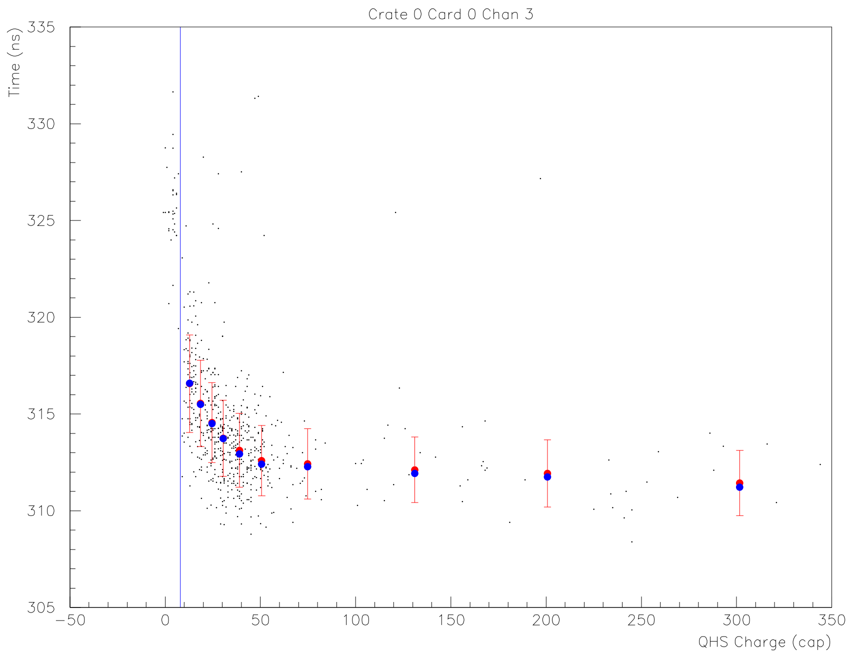

Figure 4.3: The walk parameterisation for crate 0, card 0, channel 3. The blue line shows the approximate level of the DAQ discriminator threshold. The red points show the average time and the spread for that charge bin. The blue points show the result of the averaged medians. It is these blue points which are interpolated between when calibrating.

4.4 The Gain Calibration

The gain calibration scans over the PCA laserball run collecting charge histograms for all channels. These charge histograms are then examined to find the positions of the threshold, peak and high half point. There is a charge histogram collected for each of QHS, QHL and QLX although the QLX histograms do not contain much information at the low light levels of a PCA laserball run. The electronics are structured so that the ratio of QHL:QLX is 12:1 and this information is used to calibrate QLX using the QHL constants.

On initialisation, blocks of memory are set aside to hold the channel level statistics and the charge histograms. The QHS and QHL histograms are allotted

400 words each while the QLX histograms only occupy 200 words as they are both smaller and less interesting. The histogram ranges are not defined at first because it was felt that it was better to estimate the approximate position of the charge zero from the data rather than rely on the ECA pedestal. To estimate the charge zero from the data, the histogram memory is first used to collect the first 400, or in the case of QLX 200, charges from the run, calibrated in ECA counts above pedestal. Once this data has been collected and the histogram memory is full the contents are examined to choose an approximate zero. The charge spectra look like a big peak with a long high charge tail. An 100 count sliding window is passed over the charge spectra, looking for the window with the most entries. In the case of windows with an equal number of entries, the one at the highest charge is preferred. The low edge of the window with the most entries is used as the approximate pedestal. Once the ranges have been chosen, the collected data is transferred to a temporary buffer and the memory restructured to hold the histogram. The high gain charge histograms run from -149.5 to +250.5 counts from the approximate zero and the low gain charge histograms run from -79.5 to +120.5 counts from the approximate zero. This charge is converted back into ECA counts above pedestal when the histograms are fitted in the termination phase. The charges in the temporary buffer are then fed back into the histogram. The histograms continue to get filled with data from the laserball run until the data files are exhausted.

On termination, any channels which are low occupancy and are still accumulating charges to choose the approximate zero, are processed into histogram form. The histograms are then examined to find the threshold, peak and high half points. As a first stage a sliding window 100 counts long is slid over the histogram to find the approximate position of the main peak. It is required that the window with most entries should contain at least 100 hits to continue with fitting, otherwise the channel is failed. The width of the peak can vary sizably so to get a handle on the scale of the charge peak the shortest span containing 65% of the amount of data in the sliding window is found. This span therefore characterises how spread out the bulk of the data in the main peak is. Knowing this range is useful because histograms with a small spread should be examined on a small scale while histograms spanning a large range tend to become more saw-toothed due to the lower statistics in any individual bin and so require examination on a larger scale.

The position of the peak is found by sliding a window of length $0.4\times$shortest span, rounded to the nearest integer, over the data and taking the average of the window with most entries. The height of the peak is taken to be the average contents of the sliding window at this point. The high half point is found by taking a sliding window of the same length and sliding it to the right looking for the point that the number of entries in the window falls to half that at the peak. The centre of the sliding window is then taken as the high half point.

Finding the position of the threshold is a little more involved. The threshold is the point where the charge spectra drops dramatically at the position of the discriminator threshold. Unfortunately, rather than dropping to zero there is often just a valley between the part of charge spectra caused by genuine signal and the zero charge spike caused by pickup noise. Finding the low half point did not work well because in a significant number of cases the spectra fell to half the height of the peak before the hard cut off of the discriminator. To get around this problem a slightly more complicated procedure was adopted.

The size of the window used to search for the threshold was varied depending on the size of the shortest span. The window was 1 count if the span was < 40, 2 counts if the span was between 40 and 80 and 3 counts if the span was > 80. This window was then placed at the peak and slid left until the low half point was found. The average of the charge spectrum between this point and the peak was calculated and the window slid left again until the number of entries in the window fell to half of this new average. The new average gives a better idea of the height of the charge spectrum before the hard cut of the discriminator threshold and the method seems to work well.

With the threshold, peak and high half point found for every channel, the results are then written out to titles banks. The titles banks to which the gain data is written are named GPCA, where the G stands for gain. There is one bank per crate with the bank number matching the crate number. The titles banks consist of a 10 word header, followed by the board IDs, followed by the gain information for each channel. Like with the WPCA banks the first word of the header is a duplicate of the crate number and the second word is a version number. After the header comes the board IDs, these are the identification numbers of the particular electronics at the card level. There are six board IDs for every front end card. Electronics cards are occasionally swapped out for repair and when calibrating it is important to know whether the electronics have changed since the previous PCA. After the board IDs comes the gain information. There are 7 words per channel: a status word followed by the QHS threshold, peak and high half points and then the QHL threshold, peak and high half points. The status word is similar to the walk calibration and is formed as

$$100 \times \text{tube type} + 10 \times \text{tube status} + \text{charge mask}$$

where tube type, tube status and charge mask are all integers. The tube type is the same as in the walk calibration and is 1 for normal tubes and > 1 for all other tube types. The tube status and charge mask are defined as

tube status		charge mask	
0	passed	0	All OK
1	offline	+1	QHS bad
2	zero occupancy	+2	QHL bad
3	low occupancy	+4	QLX bad
4	low QHS or QHL ECA		

The tube status is actually ignored by the calibrator, and is only a summary statistic. The condition for a tube to be marked as low occupancy or low ECA is different than with the walk calibration. The median number of hits received by

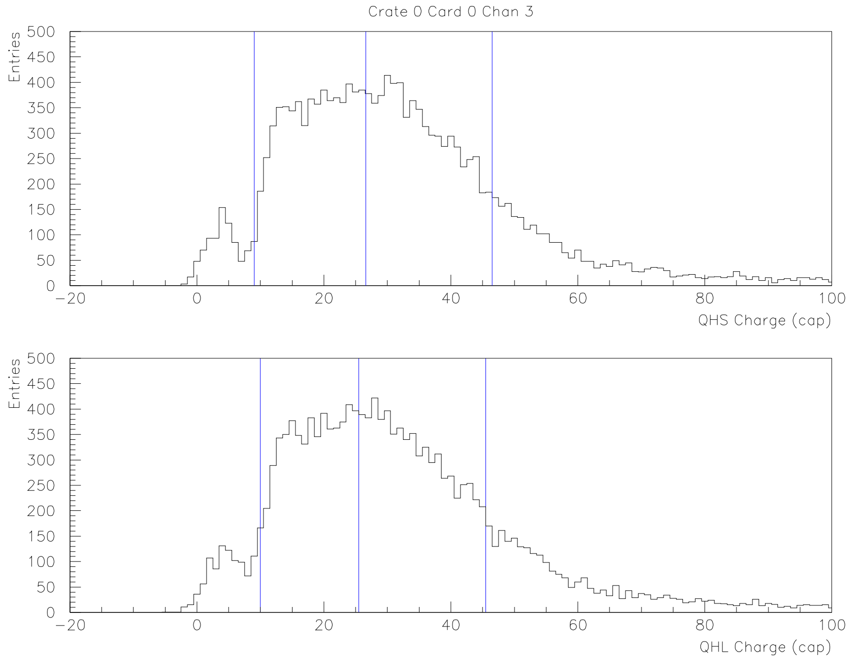

Figure 4.4: This figure shows the QHS and QHL histograms for crate 0, card 0, channel 3. The fitted threshold, peak and high half points are shown.

normal online channels is calculated and if a tube receives less than 20% of that number then it is marked as low occupancy. If a tube is not low occupancy but less than 30% of its hits pass the ECA for QHS or QHL then it is marked as low ECA. To calibrate a charge only requires that the high half point was found and the charge mask is not set. Currently the charge mask is not used.

The collected histograms are written out to an HBOOK file along with an ntuple containing the gain data. Figure 4.4 shows the QHS and QHL histograms for crate 0, card 0, channel 3 with the fitted threshold, peak and high half points superimposed. Again many of these plots were scanned to ensure the gain calibration was working as intended. Figure 4.5 shows a channel where the gain calibration has produced a spurious result due to a large amount of noise.

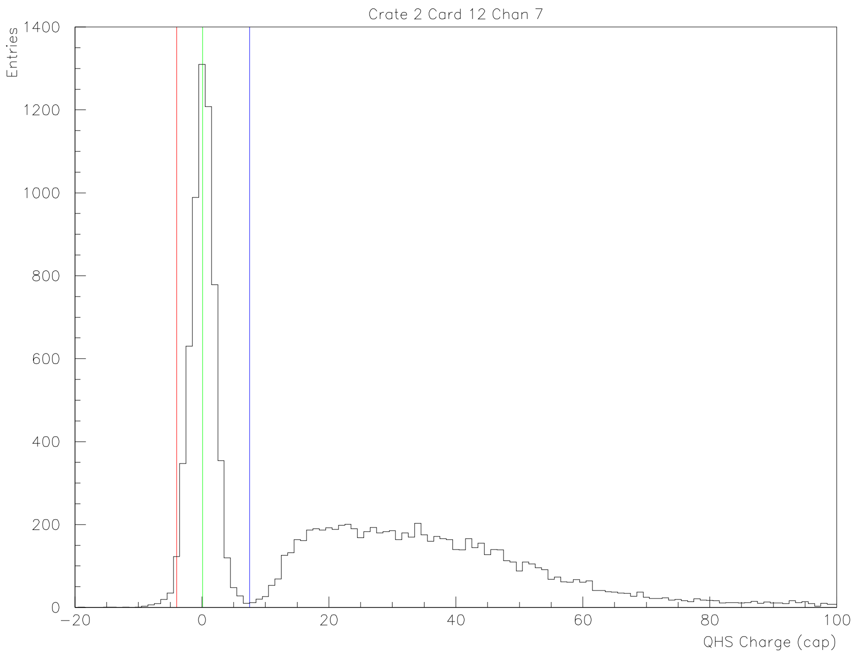

Figure 4.5: The gain calibration can produce a spurious result for channels with a large amount of noise.

4.5 The Extended Walk Calibration

The extended walk algorithm is very similar to the walk algorithm, described in section 4.3, but has some additional features. There are three main differences in the titles banks compared with version 1: the board IDs are recorded in the header, each channel has one additional word indicating which cuts it failed and another word, containing a straight line fit at high charge, replaces the word containing the value of the walk at high and low charge. The bank numbers are given by the crate number plus 100, so that the banks can co-exist in the database with the version 1 constants. The version number in the header is set to 2. In addition, the channel RMS is calculated in a slightly different way.

The extended walk calibration collects statistics in the same way as the standard walk calibration. As in the standard walk calibration the PMT time is tested to see if it is within 10 ns of the stored prompt peak time. If it is then the charge is tested to see which, if any, of the charge bins the hit should be added to. Again, like the standard walk calibration, these hits are stored in the charge bin buffers which get read out once full. In the extended walk calibration, as well as calculating the average value of the charge, time, time squared and median time, the average value of the charge squared and the charge $\times$ time are kept. In other words, for each bin, the average values of q, t, qq, tt and qt are kept along with the average median time, where q and t stand for charge and time respectively. These quantities allow the 'average' gradient of the walk function to be taken into account when calculating the variance of a bin in the termination phase. The variance of the data, about an average gradient m, is given by

$$var = \frac{N}{N-1}\left[(A_{tt} - A_t^2) - 2m(A_{qt} - A_q A_t) + m^2(A_{qq} - A_q^2)\right] \tag{4.2}$$

where A_{qt} stands for the average value of q times t. This formula reduces to its familiar form in the limit that $m = 0$. The question then arises what gradient should be used? The least squares gradient of the bin is given by

$$m = \frac{A_{qt} - A_q A_t}{A_{qq} - A_q^2} \tag{4.3}$$

The gradients that the calibrator will use for charges greater and less than the average charge of the bin are known because the interpolation points are fixed in termination mode. These gradients are used to put an upper and lower limit on the gradient found using the above least squares formula. The variance of the bin is then given by equation 4.2. The individual bin variances are then combined to give an estimate of the total variance, and hence RMS, of the channel.

The straight line fit to high charge data is more straightforward. The average q, t, qq, tt, and tq for the three highest charge data points are amalgamated and then the above equation 4.3 used to derive the gradient. The intercept is given by the median time (this would normally be the average time) minus the gradient times the average charge. Figure 4.6 shows the extended walk fit along with the

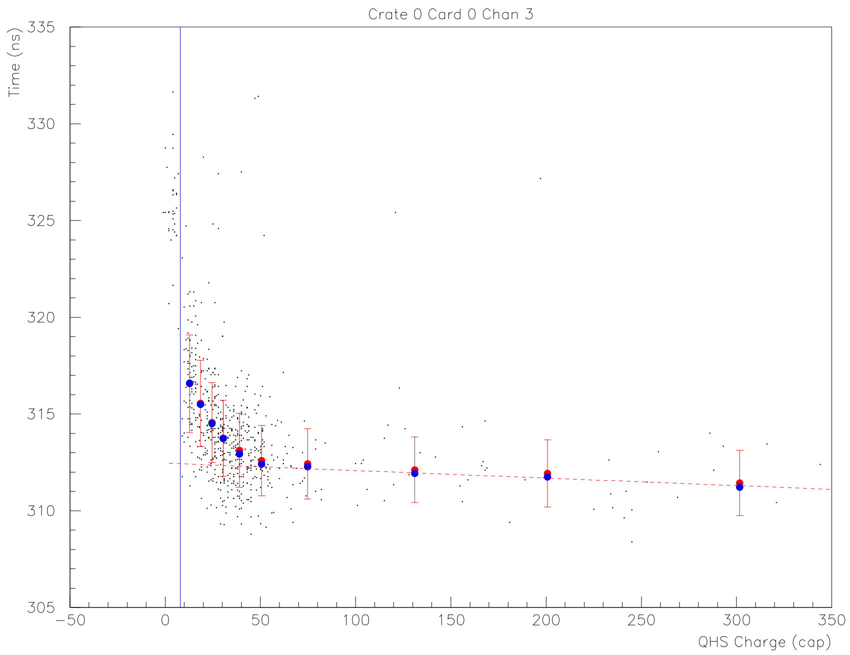

Figure 4.6: The raw data for crate 0, card 0, channel 3. The interpolation points and the high charge straight line fit are shown.

high charge straight line fit.

The summary word is formed in the same way as in version 1, ie:

$$100 \times \text{tube type} + 10 \times \text{tube status} + \text{fit quality}$$

The only difference being in the meaning of the fit quality word

	fit quality
0	not fitted
1	passed
2	no points
3	failed cut
4	failed RMS

'No points' caters for the situation that none of the charge bins has enough data to be accepted as an interpolation point. 'Failed cut' indicates that the channel failed one of the three bad channel cuts, see section 4.8, but the overall channel RMS was good. 'Failed RMS' indicates that the channel RMS was over 2.8 ns, see section 4.8.1.

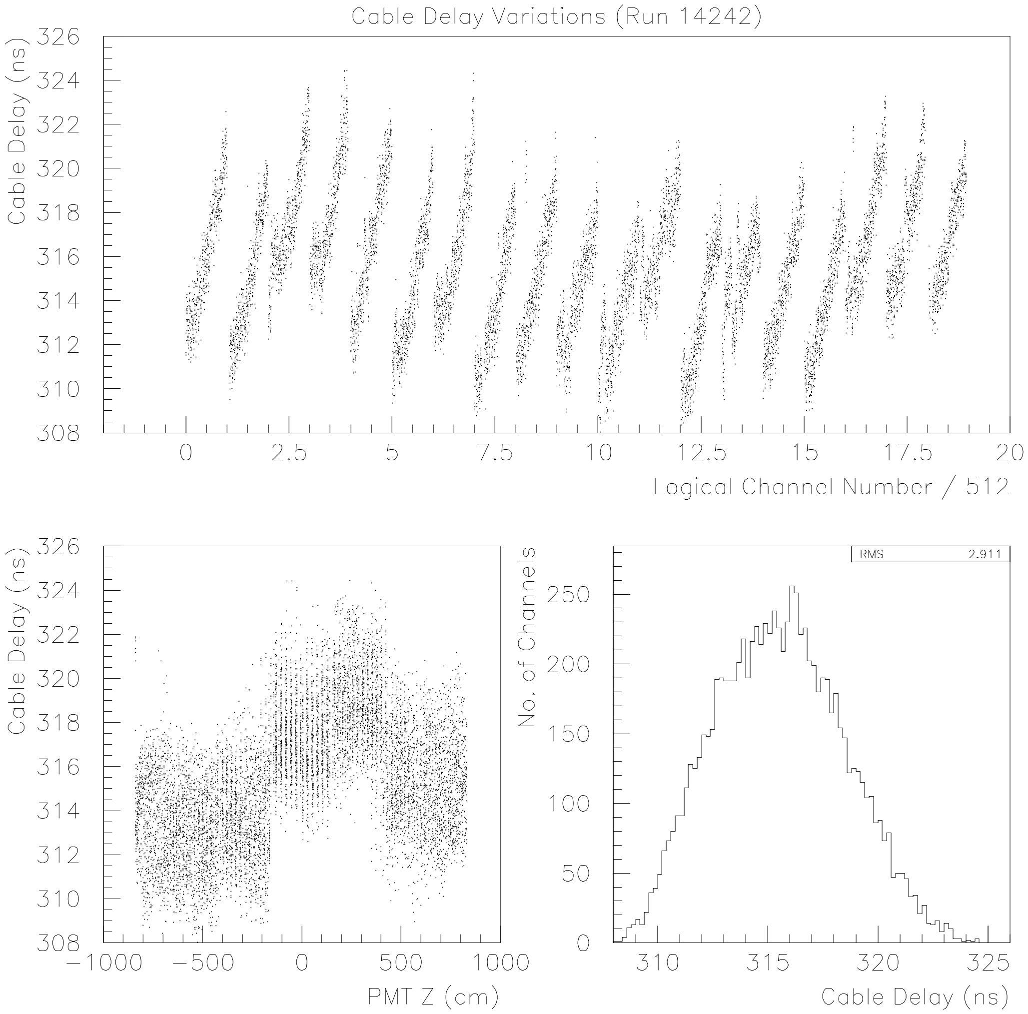

Figure 4.7: Plots showing the amount of variation in the cable delay. The logical channel number is divided by 512 to display the crate numbers. Also shown is the variation of the cable delay with PMT Z and a histogram of the cable delays. Only normal online channels which have a high charge fit are displayed, in addition a small number of outliers (7) have been cut.

Having described how the channel information is collected, a couple of instructive plots can be made from the results. Figure 4.7 shows the channel to channel variation in the cable delay. The cable delay was defined to be the value of the high charge linear fit at 100 counts above pedestal. Figure 4.8 shows the cable delay

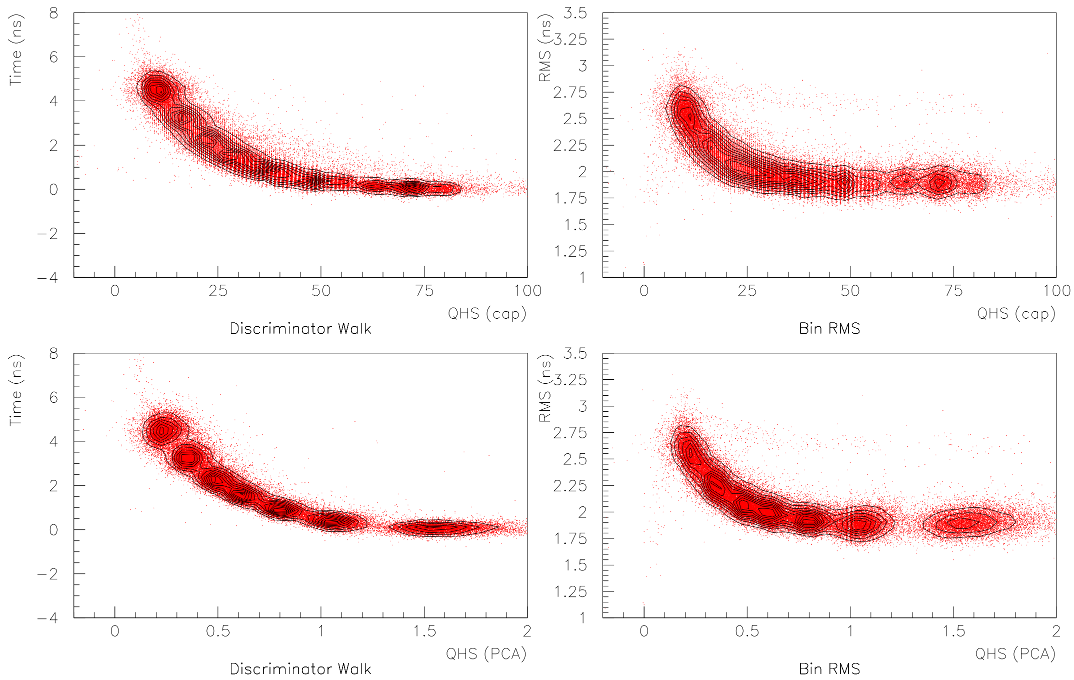

Figure 4.8: The two plots on the left show the cable delay corrected time against charge of the PCA interpolation points for all channels which passed the PCA. The two plots on the right show the RMS of the same interpolation points against charge. The charge in the top two plots is counts above pedestal while the bottom two plots are in PCA corrected time. Contour lines are added in black. The run used was 14242.

corrected time and RMS of the interpolation points against charge in both counts above pedestal and PCA corrected charge. The cable delay corrected time is simply the time of the bin minus the value of the high charge linear fit at 100 counts above pedestal. Comparing the two plots showing the walk against charge it can be seen that the plot against PCA corrected charge has slightly less scatter. The plots of the RMS show that it is strongly correlated with charge. The few points at low charge which have a low RMS are probably due to the first point accidently fitting to noise. The time uncertainty as a function of charge will be investigated further in the next chapter.

4.6 The Calibrator

The calibration processor in SNOMAN is named CAL. The job of CAL is to apply the electronics calibration constants and PMT calibration constants to the raw ADC times and charges which are read in from data files. The user can specify the degree to which they want the time and charge to be calibrated. The basic options are counts above pedestal, full ECA calibration and PCA calibrated. The PCA calibration is only available for normal PMTs, for other types of PMT the calibration is left at full ECA calibration for the time and counts above pedestal for the charge.

When analysing data, it would be very tiresome to have to explicitly load the relevant calibration constants for the date of the data being analysed. SNOMAN solves this problem by using a database [46] which stores the titles banks and can supply SNOMAN with the most recent set of constants for any particular date and time. WPCA and GPCA titles banks are added to the database with a validity range running from the time of the laserball run until a couple of months later. The validity range of the LRBP banks is non-standard. LRBP banks are only valid for the duration of the PCA laserball run. The calibrator gets the date and time of the PCA laserball run used to generate a walk calibration from the WPCA bank creation date. It then looks in the database at that date and time and retrieves the required data from the LRBP bank. The reason for this odd arrangement is that there are separate WPCA banks for each crate. It is therefore conceivable that the walk calibration of different crates could have been generated from different PCA laserball runs. In retrospect this is an unlikely scenario. It is conceivable that a method for fitting the position of the laserball, which is more accurate than the manipulator system, could be devised. In this case laserball banks could be added to the database for all laserball runs, not just PCA runs, without causing any problem.

4.6.1 Applying the Walk Calibration

To generate PCA calibrated time, the calibrator first converts the charge to counts above pedestal and finds the fully calibrated ECA time. The calibrator then finds the data recorded in the WPCA banks for that channel and checks to see if it passed the calibration. If no calibration is available for that channel then the time is returned as undefined (-9999.). For channels with a calibration the calibrator finds the two walk interpolation points which straddle the charge of the hit to be calibrated. A straight line interpolation between these points give the value of the walk function at that charge. If the charge is below the first interpolation point then the first two interpolation points are used to extrapolate back, but only if the difference in charge between the data and the first interpolation point is less than the difference in charge between the first and second interpolation points. If the charge is less than that then the time is returned as undefined. This cut-off is typically a few counts below the DAQ threshold, see figure 4.9. If the charge is above the last interpolation point then the last two interpolation points are used to extrapolate provided that the difference in charge between the data and the last interpolation point is less than half the difference in charge between the last two interpolation points. If the data's charge is even higher than this then, for walk version 1, the estimate of the time at high charge recorded in the WPCA bank is used. When calibrating with walk version 2 the straight line fit at high charge is used. The LRBP bank gives the laserball position and the PCA calibrated time is then calculated using equation 3.1.

It is worth commenting further on the effect of the 'walk threshold' shown in figure 4.9. In any type of run, including neutrino runs, there are channels which fire due to pickup noise. The charge of these hits are mostly below the discriminator threshold and also the walk threshold. This results in a surprising number of channels not being calibrated by the PCA. Figure 4.10 shows the charge spectrum of all hits from a neutrino run and the charge spectrum of hits which fail the time calibration. The upper plot was made using the standard SNOMAN while the lower

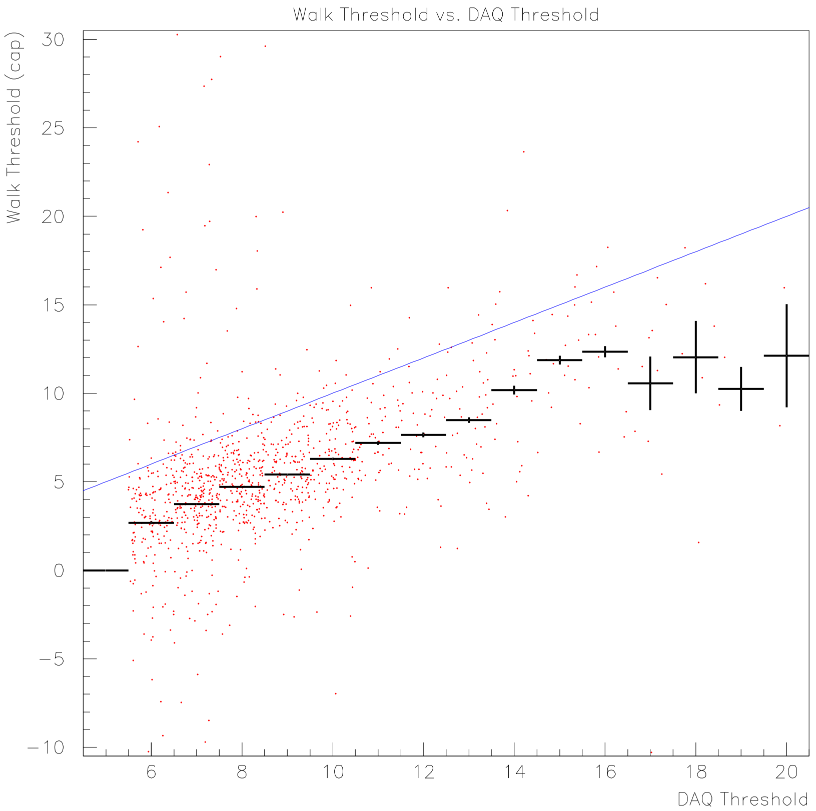

Figure 4.9: The calibrator returns undefined times for hits whose charge lies too far below the first interpolation point. The cut-off point is currently defined as the charge of the first interpolation point minus the difference in charge between the first and second interpolation points. This 'walk threshold' is typically a bit less than the DAQ discriminator threshold. The comparison relies on the accuracy of the pedestal and most of the outliers are due to bad pedestals. Only channels which passed the walk calibration can be plotted. The run used was 11326. The DAQ discriminator threshold is an integer but the data was entered into a 2-D histogram with bins of size 1 count by 1 count to better show the density of entries. The black crosses are a profile histogram of the data (accepting all data with a walk threshold between -10. and 25.) with the vertical size of the cross indicating the error on the mean. The blue line shows $x = y$.

plot was made using a modified SNOMAN which does not have a walk threshold. The modified SNOMAN kept extrapolating back using the first two interpolation points without any cut off. It can be seen that the effect of the walk threshold is that the noise spike near zero charge is mostly uncalibrated. The walk threshold

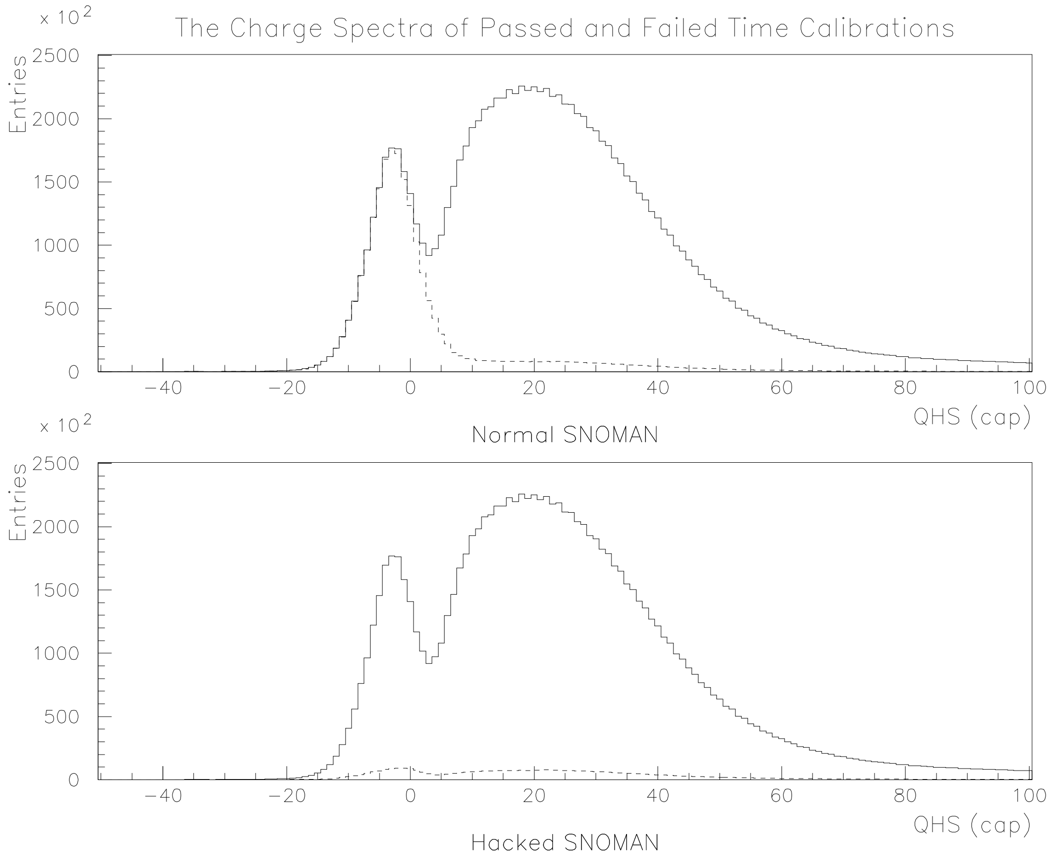

Figure 4.10: The solid line shows the charge histogram of all hits from a neutrino run. The dotted line shows the charge spectrum of hits which fail the time calibration ($\leq$-9990). The upper plot is for normal SNOMAN while the lower plot is for a modified SNOMAN which does not have a walk threshold. See the text for more details. The neutrino run used was 15083, which was normal except that the D_2O was being circulated.

will be returned to briefly at the end of chapter 6 when discussing the effect of the detector rate on the measured charge (note that the noise spike in figure 4.10 is not exactly at zero but is actually around minus 2 or 3).

4.6.2 Applying the Gain Calibration

The generation of the PCA calibrated charge is simpler. Like the walk calibration, if a channel failed the gain calibration then the charge is returned as undefined. For the other channels, the charge in counts above pedestal is simply divided by the high half point for that charge. The high half point for QLX is taken to be the high half point for QHL divided by 12. The gain calibration routine can also return the position of the threshold and the position of the peak if requested.

4.7 Bad Pedestals

The channel anomaly which has proved most problematic to the PCA calibration is the presence of low charge noise. The way which this low noise is excluded from the walk calibration, described in section 4.3, avoids the use of the pedestal. This is obviously a peculiar decision because the most obvious feature of the noise is that it appears at zero charge, below the discriminator threshold. There are two ways in which genuine hits can appear below the ECA pedestal plus the DAQ discriminator threshold: bad pedestals and bad ADCs. This section discusses the effect of bad pedestals; bad ADCs will be discussed further in the next section.

The gain calibration estimates the position of the threshold, in units of ADC counts above pedestal, using the charge histograms it collects. The discriminator threshold is set in units of DAQ counts. DAQ counts and ADC counts are very similar in size and so the gain threshold should roughly equal the DAQ discriminator threshold. Figure 4.11 compares the QHS gain threshold and the QHL gain threshold with the DAQ threshold. The caption explains the details of what is plotted. It can be seen that the gain threshold and the DAQ threshold broadly agree but there are some outliers.

Figure 4.12 shows the difference between the gain threshold and the DAQ threshold versus channel number for QHS and QHL. Clearly there is something peculiar about channel 17. Figure 4.13 shows the charge histograms for the two largest outliers in QHS. The blue line shows the DAQ threshold assuming an accurate pedestal while the red line shows the gain threshold. The charge histograms appear normal and it seems that the ECA pedestal is just wrong.

Although channel 17 is particularly prone to the problem of bad pedestals the

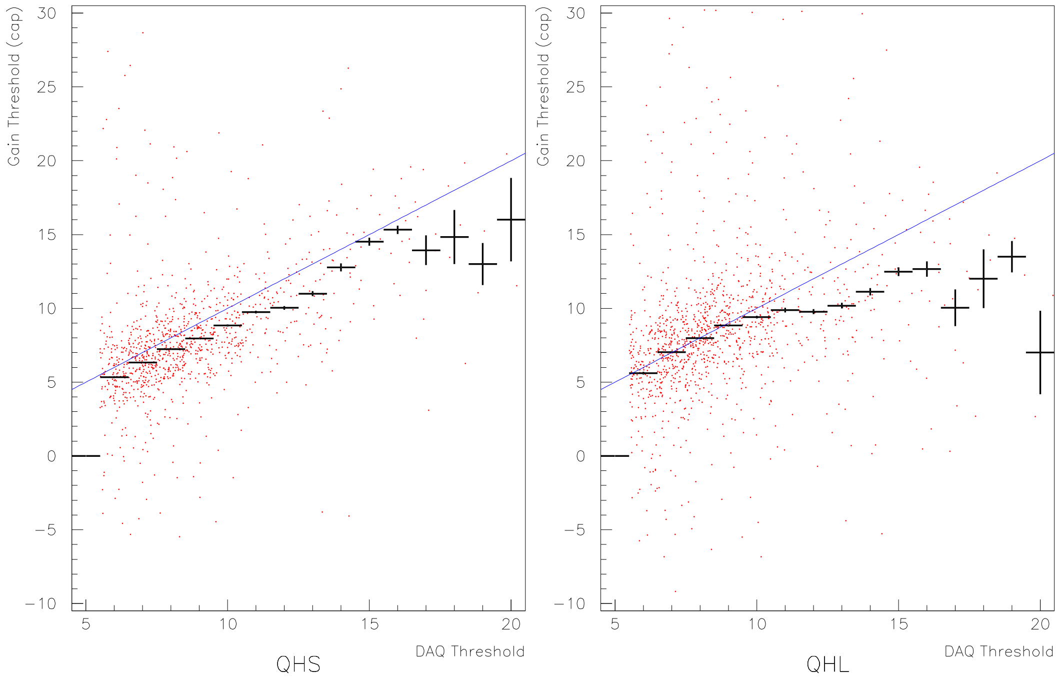

Figure 4.11: This plot compares the value of the threshold as found by the gain calibration for QHS and QHL to the value of the DAQ threshold for PCA run 11326. Only normal, online channels with at least 100 calibrated hits are plotted. The DAQ discriminator threshold is integer and the gain threshold half integer. The data was entered into a 2-D histogram with bin size 1 count by 1 count to better show the density of entries. The black crosses are profile histograms of the data (accepting all data with a gain threshold between -10. and 25.) with the vertical size indicating the error on the mean. The blue lines show $x = y$.

problem is not confined to those channels. The upper plot in figure 4.14 shows the difference between the QHS gain threshold and the DAQ threshold for all channels except the channel 17s. There are still some outliers and 14 of them have a difference of greater than 15 ADC counts. These 14 outliers were hand scanned to see whether the difference was due to a bad pedestal or a failure of the gain calibration. Two of the outliers were due to the gain calibration failing and incorrectly incorporating noise in the charge spectrum (an example of this type of failure was given in figure 4.5). 12 of these 14 outliers were due to bad pedestals. The lower plot in figure 4.14 shows one of these channels.

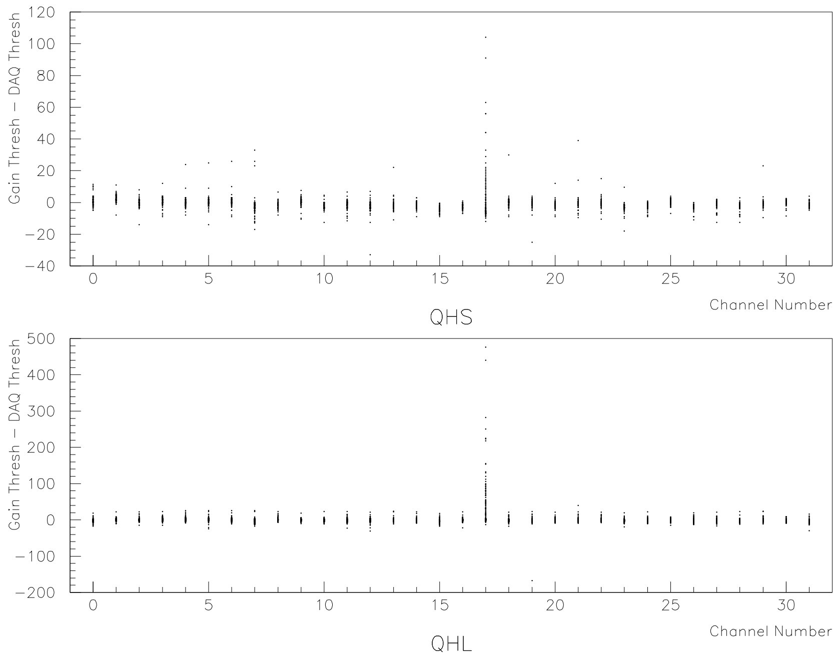

Figure 4.12: The difference between the gain calibration threshold and the DAQ discriminator threshold for QHS and QHL. Only normal, online channels with at least 100 calibrated hits are plotted. The difference is particularly large for channel 17.

What effect will these bad pedestals have on the walk calibration? An identical and unchanging offset in the pedestals of all the cells on a channel does not affect the walk calibration. The interpolation points appear normal apart from the charge offset. When applying the calibration, the pedestal corrected data will also have the offset and so the walk will be applied correctly. If the ECA pedestal changes from ECA to ECA, or from run to run, then this would cause a problem. The ECA has been updated to find and reject channels whose pedestals change from one ECA to the next [35].

Another problem arises when trying the estimate the energy of an event. Clearly it would be desirable to cut any hits which are due to noise. To do this we need to know where the discriminator threshold lies. The threshold used by the walk calibration (the charge of the first interpolation point minus the difference in

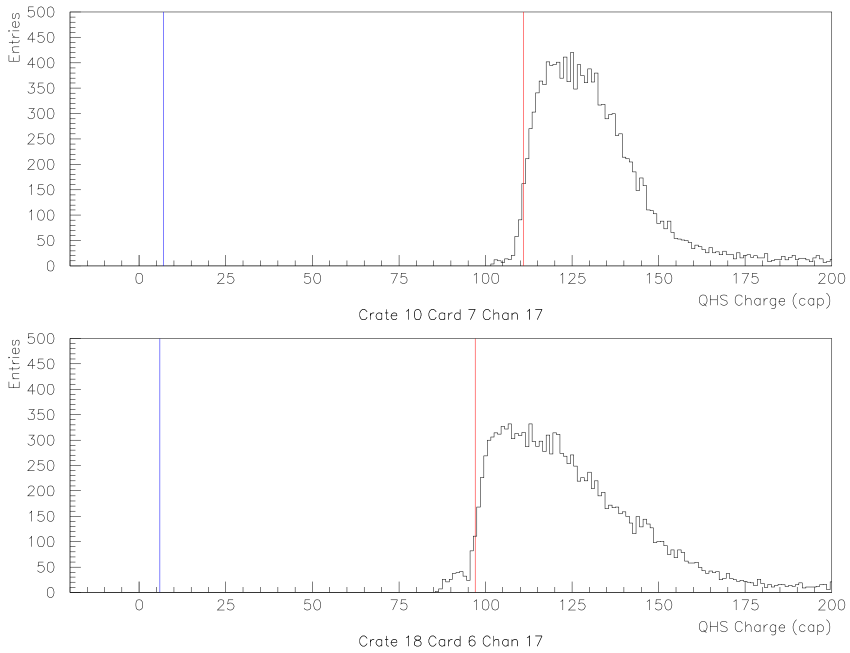

Figure 4.13: The charge histograms for the two channel 17s with the largest difference between the gain threshold and DAQ discriminator threshold. The red line shows the location of the gain threshold and the blue line the position of the DAQ threshold assuming an accurate charge pedestal.

charge between the first two interpolation points) can be too low if noise was included in the walk calibration. We have seen here that the pedestal plus DAQ threshold is not completely reliable either. There is therefore no completely foolproof way of rejecting noise hits.

Bad pedestals also cause a problem for the gain calibration. The gain calibration uses the position of the hhp as a measure of the gain of the channel. Any offset in the pedestal will result in an offset in the hhp. A possible solution would be to call the real hhp

$$\text{real hhp} = \text{gain hhp} - \text{gain threshold} + \text{DAQ discriminator threshold}$$

The trouble with this method is that it relies on the accuracy of the gain threshold. The low charge noise means that the gain threshold is much less reliable than the

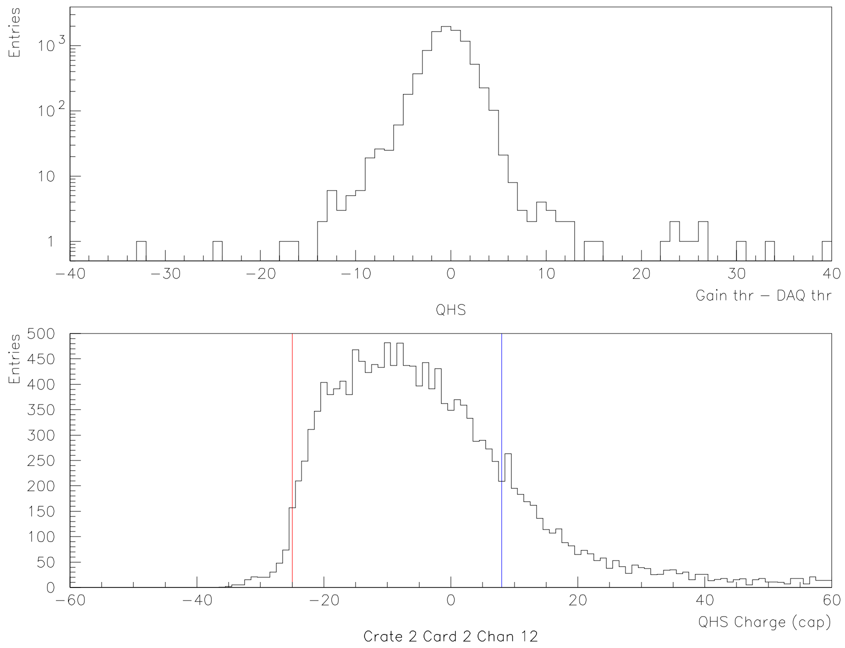

Figure 4.14: The upper plot shows the difference between the QHS gain threshold and the DAQ threshold excluding all channel 17s. Again only normal, online channels with at least 100 calibrated QHS hits are shown. The lower plot is an example of one of the outliers. The red line shows the gain threshold and the blue line shows the DAQ threshold assuming an accurate pedestal.

gain hhp. The choice is between a few bad hhp due to bad pedestals or some bad hhp due to bad gain thresholds. The current choice is to assume the pedestal is OK.

4.8 The Bad Channel Cuts

This section describes the cuts used by the walk calibration to identify bad channels. The form of bad channel behaviour which was most common around the start of production data taking was bad ADCs. A bad ADC has discontinuities in its conversion of voltage to ADC counts. Figure 4.15 shows the effect of a bad charge ADC and a bad time ADC. The ADCs deal with all the data from a card, so whole cards are affected at a time. A temporary fix to the bad ADC problem, involving heating up the chip, was found around the beginning of production data taking in

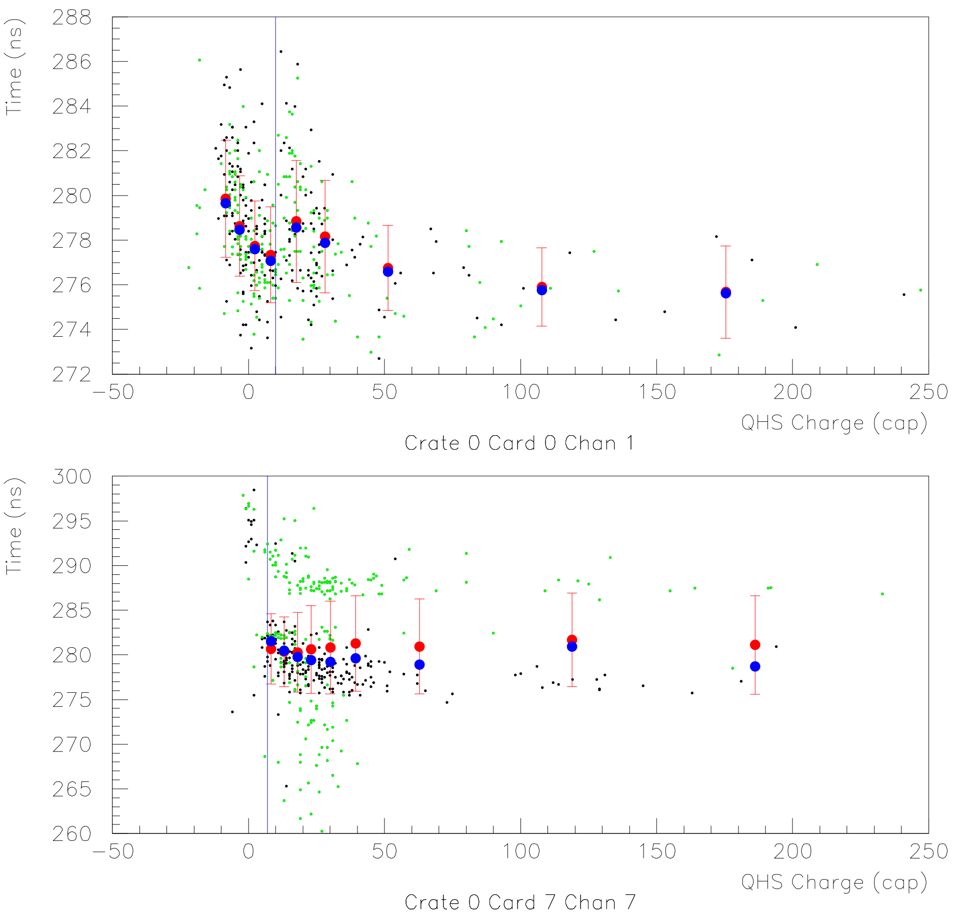

Figure 4.15: The result of the walk calibration for two channels suffering from a bad ADC. The red points show the average time and the spread for each of the charge bins. The blue points show the average of the medians. A small subset of the raw data is plotted in black and green. The black points are from odd cells and the green points from even cells. The upper channel has a bad charge ADC and the bottom channel has a bad time ADC. These plots are taken from run 4603 which was before the high voltage raise.

late 1999. Over time, the ADCs would go bad again and require re-heating. A permanent fix, reducing the voltage on the ADC, was found in August 2000. The walk calibration uses four parameters to identify bad channels. The first one is the approximate RMS of the channel. The calculation of the RMS has already been discussed in sections 4.3 and 4.5. The three other parameter used are named *tstep*, *flate* and *fout* and will be described here. Section 4.8.1 explains how the positions of the cuts on these parameters were chosen.

Tstep is a simple parameter designed to catch bad charge ADCs. Normal walk functions monotonically decrease with increasing charge. Channels which have a double charge spectrum due to a bad ADC often contain a step up instead of a continuous fall, see figure 4.15. The interpolation points are examined in turn, calculating the difference in time from the lowest time point yet found. A step up in time of a bin could conceivably occur due to statistical fluctuations, so once a step up has been found, the significance of the step must be estimated. It is not obvious how to calculate the uncertainty on the average of medians. It was felt that it might be dangerous to use the uncertainty from the spread of the data because every so often, by chance, it might be unusually low. Instead, a simpler, more robust method was used. The typical time RMS for a normal channel is around 2 ns so the uncertainty on the time was taken to be 2 ns divided by the square root of the number of entries in that bin. The error in the time difference between bin x and bin y is therefore approximately given by:

$$\sigma_{\mathrm{tdiff}} \simeq 2\sqrt{\frac{1}{n_x}+\frac{1}{n_y}}$$

Steps up in time are only considered significant if they are more than 3 sigma, as calculated using this formula. The size of any significant steps up are then tested to see if they exceed some predefined size. Figure 4.16 shows the size of the largest step up for all normal channels from all PCA runs since the high voltage raise until mid-February 2001, along with the distribution of the channel RMS. The current level of the cut *tstep* cut is 1.5 ns (this will be discussed further in the following section).

Flate and *fout* are simple cuts which check the fraction of hits which are contained in the fit. Figure 4.17 schematically shows the regions where the walk calibration records the number of entries. Regions TLESS and TMORE contain hits which are more than 10 ns from the overall time peak. Regions QLESS and

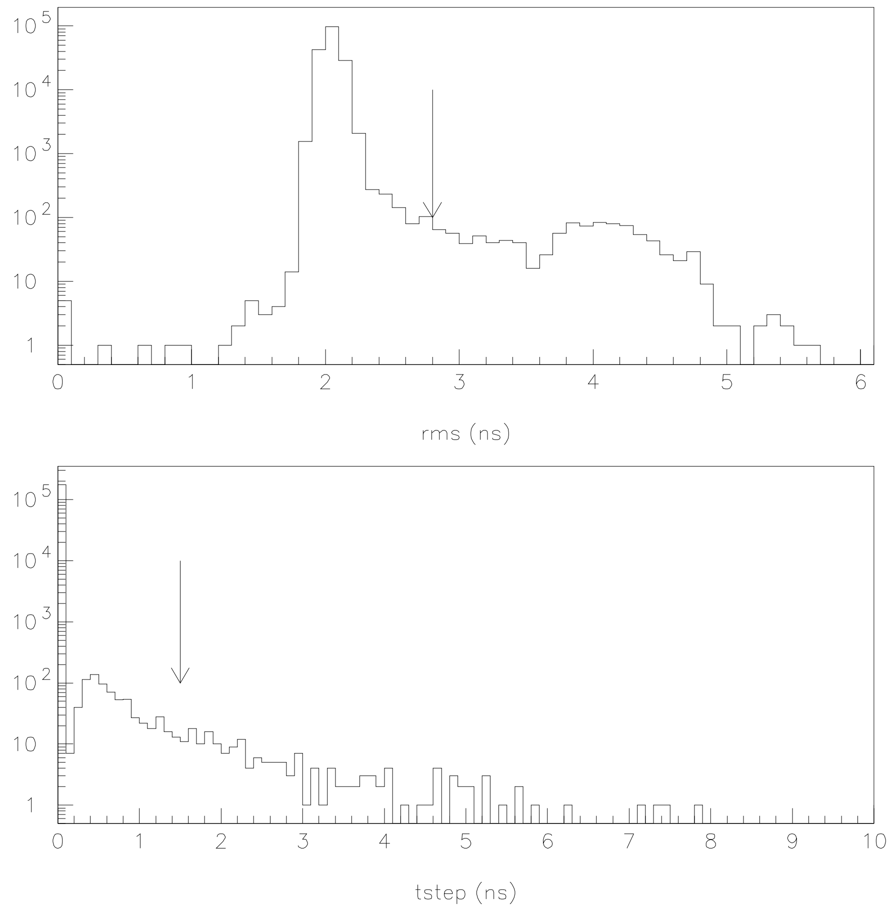

Figure 4.16: Histograms of the channel RMS and largest *tstep* for every fitted channel in all PCA runs after the HV raise until mid-February 2001. Channels are only fitted if they are normal, online channels with at least 100 calibrated charge time pairs. *Tstep* has a spike at zero because most channels do not have a significant time step up. The arrows show the position of the bad channel cuts.

QMORE contain hits which fall below the low edge of the lowest charge bin and above the high edge of the highest charge bin. The cut *flate* examines the fraction of the total number of hits which are in the region TMORE. This cut cannot be made too stringent because genuine hits from scattered and reflected light can fall in this region. The cut *fout* examines the fraction of hits which are in regions TLESS and QLESS. This cut can be made somewhat harder because in the absence of noise

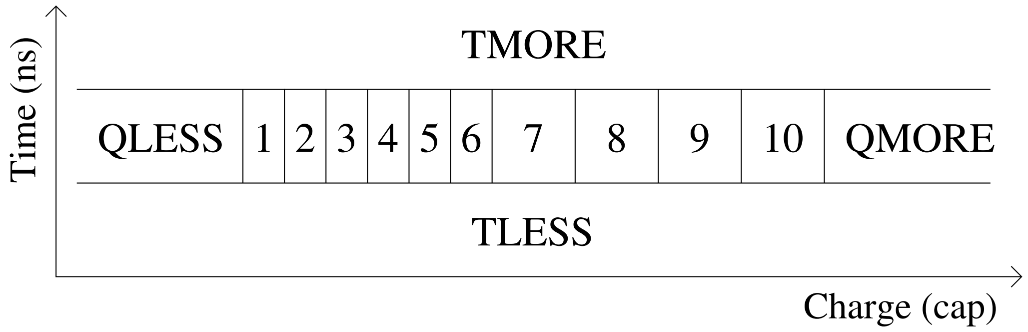

Figure 4.17: This is a schematic showing the regions where the walk calibration records the number of entries. The horizontal bar is centred on the estimate of the ECA time peak and includes all points within ± 10 ns. The charge bins are labelled 1–10. The regions TMORE, TLESS, QMORE and QLESS are discussed in the text.

no hits should fall in these regions. Figure 4.18 shows histograms of *flate* and *fout* for all normal channels from all PCA runs since the high voltage raise until mid-February 2001. The current level of *flate* is 0.4 and the current level of *fout* is 0.1. These levels are discussed in the following section.

4.8.1 Setting the Level of the Bad Channel Cuts

Armed with four quantities which are sensitive to bad channels, how is the level of the cut defined? The tails of the distributions of RMS, *tstep*, *flate* and *fout* are composed of two types of tube: genuinely bad channels and channels where the fit is perfectly OK but the presence of noise has resulted in a suspicious looking value of the parameter. Examples of one bad and one good channel which fail each of the cuts are given in figures 4.19, 4.20, 4.21 and 4.22. When setting the level of the cuts one would really like to know the underlying distribution of good and bad channels. In the absence of any clear separation between good and bad channels, there is no easy way to do this. The only way to decide whether a channel is good or bad is to hand scan the raw data and fits and make a judgement. If there were a way to automate this then the automation routine would become a new and improved cut. It is possible that another data set, not so susceptible to noise, could be used to decide which channels are good and bad. If such a method were pursued then it

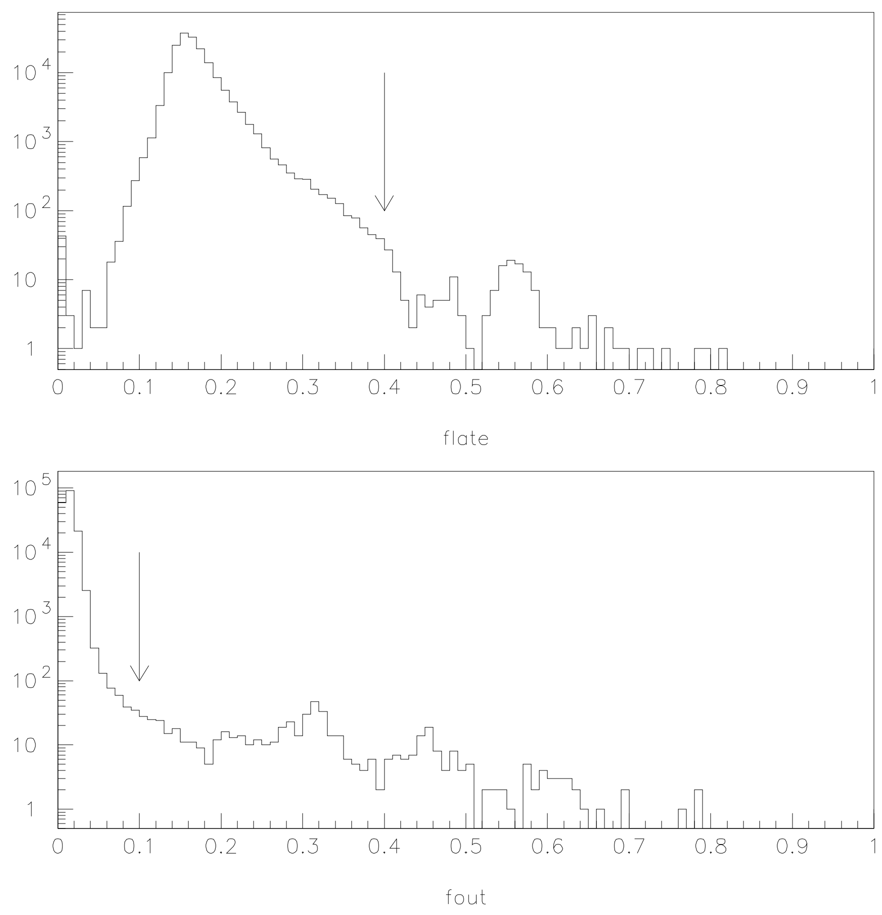

Figure 4.18: Histograms of *flate* and *fout* for every fitted channel in all PCA runs after the HV raise until mid-February 2001. Channels are only fitted if they are normal, online channels with at least 100 calibrated charge time pairs. The arrows show the position of the bad channel cuts.

would make sense to have it as a separate calibration with its own titles banks.

The RMS cut was placed as close to the peak as seemed prudent given the channels which are slightly noisy but otherwise OK. For *tstep*, the peak at around 0.5 ns seemed to contain a large number of reasonable looking channels and so the cut was placed at a higher level, using the force. The *flate* cut was placed at the end of the shoulder that is falling away from the main peak. The noise is mostly late and so affects *flate* more than the other cuts, so the shoulder contained a large

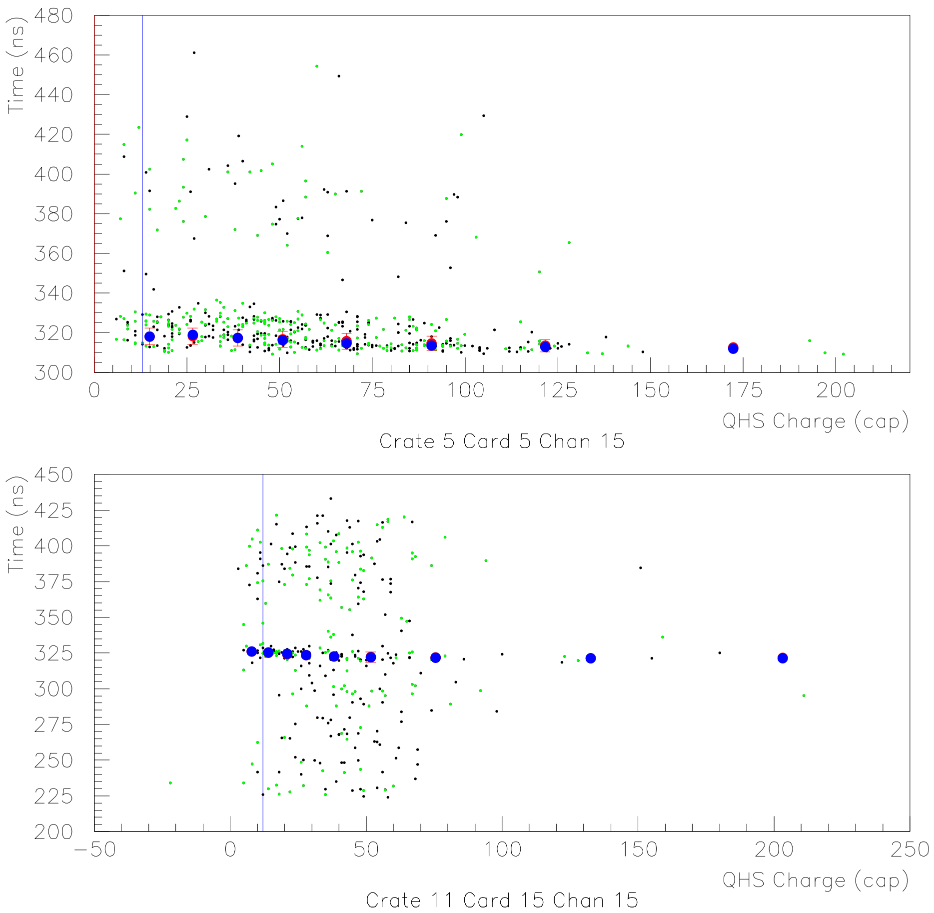

Figure 4.19: The result of the walk calibration for two channels which fail the RMS cut. There are half a dozen channels similar to the upper plot. The lower channel is normal except that it is noisy and so has an artificially high RMS. The RMS cut is the most reliable cut for identifying weird channels.

fraction of good channels. The level of the *fout* cut was placed close to the peak.

Figure 4.23 is an important plot. It shows the pass and fail rate for all PCA calibrations since the HV raise until mid-February 2001. The lower plot shows the number of channels which fail each of the cuts. *Tstep*, *flate* and *fout* are typically responsible for an addition 24 channels failing each PCA above the number failing due to a high RMS. The number of channels which fail is around the 1% level.

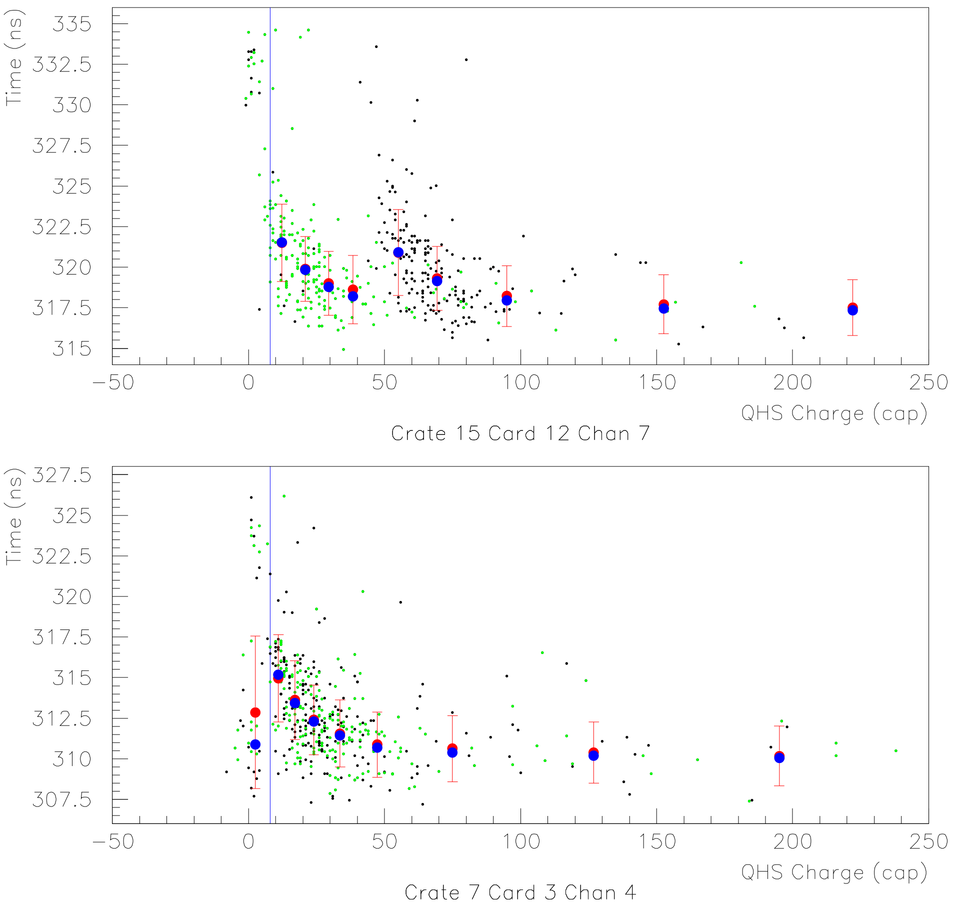

Figure 4.20: The result of the walk calibration for two channels which fail the *tstep* cut. The upper channel has been correctly identified as bad while in the lower plot the cut has been fooled by what appears to be low charge noise. These plots have been taken from PCA run 11326.

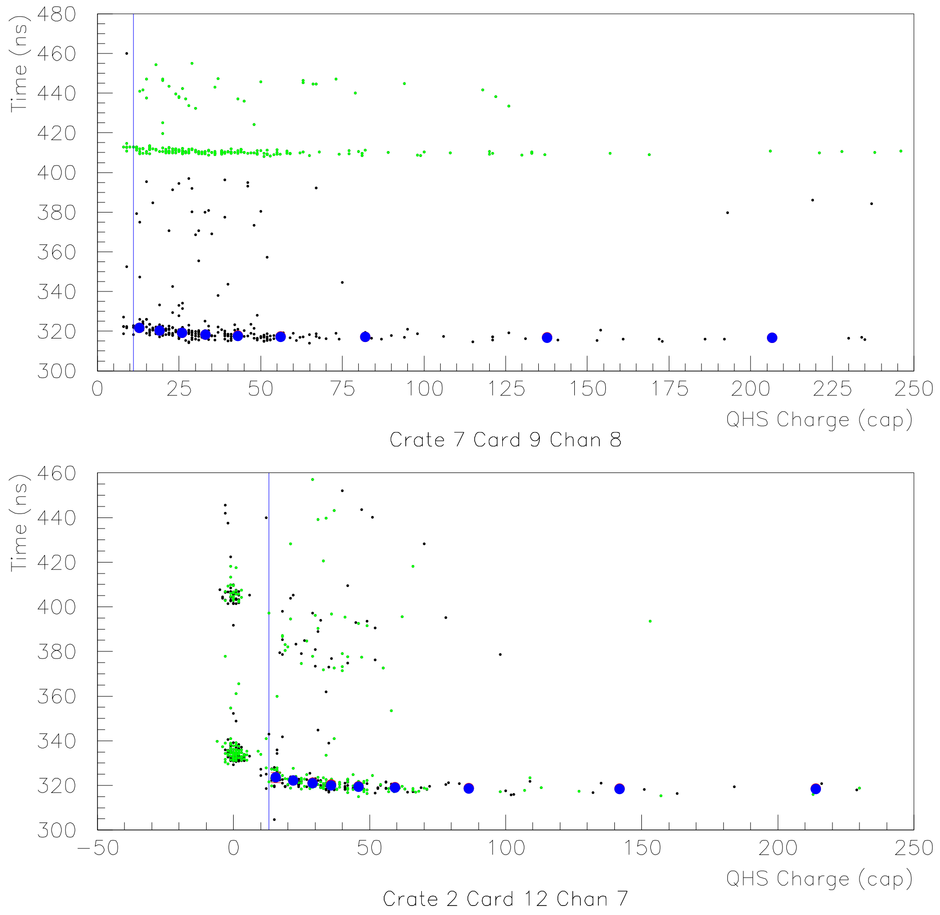

Figure 4.21: The result of the walk calibration for two channels which fail the *flate* cut. The upper plot has been correctly identified as bad while in the lower plot the cut has been fooled by low charge noise. These plots have been taken from PCA run 14444.

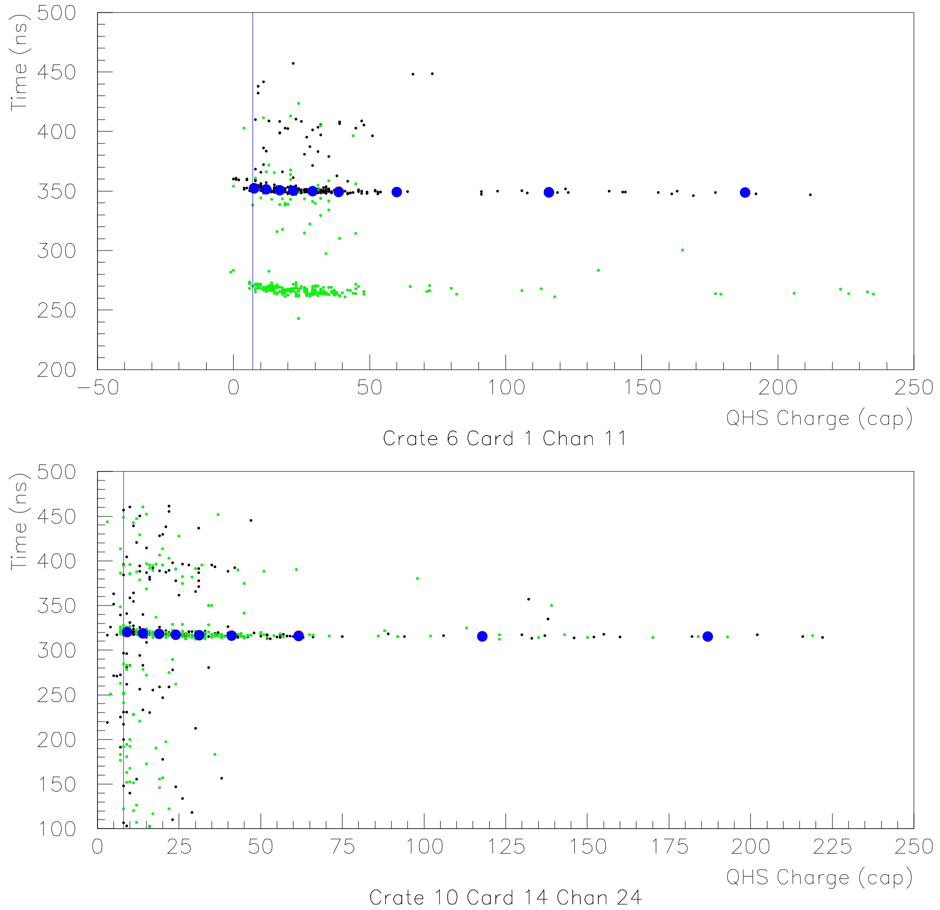

Figure 4.22: The result of the walk calibration for two channels which fail the *fout* cut. The upper plot has been correctly identified as bad while in the lower plot the cut has been fooled by a noisy channel. The upper plot is from run 11326 and the lower plot from run 10623.

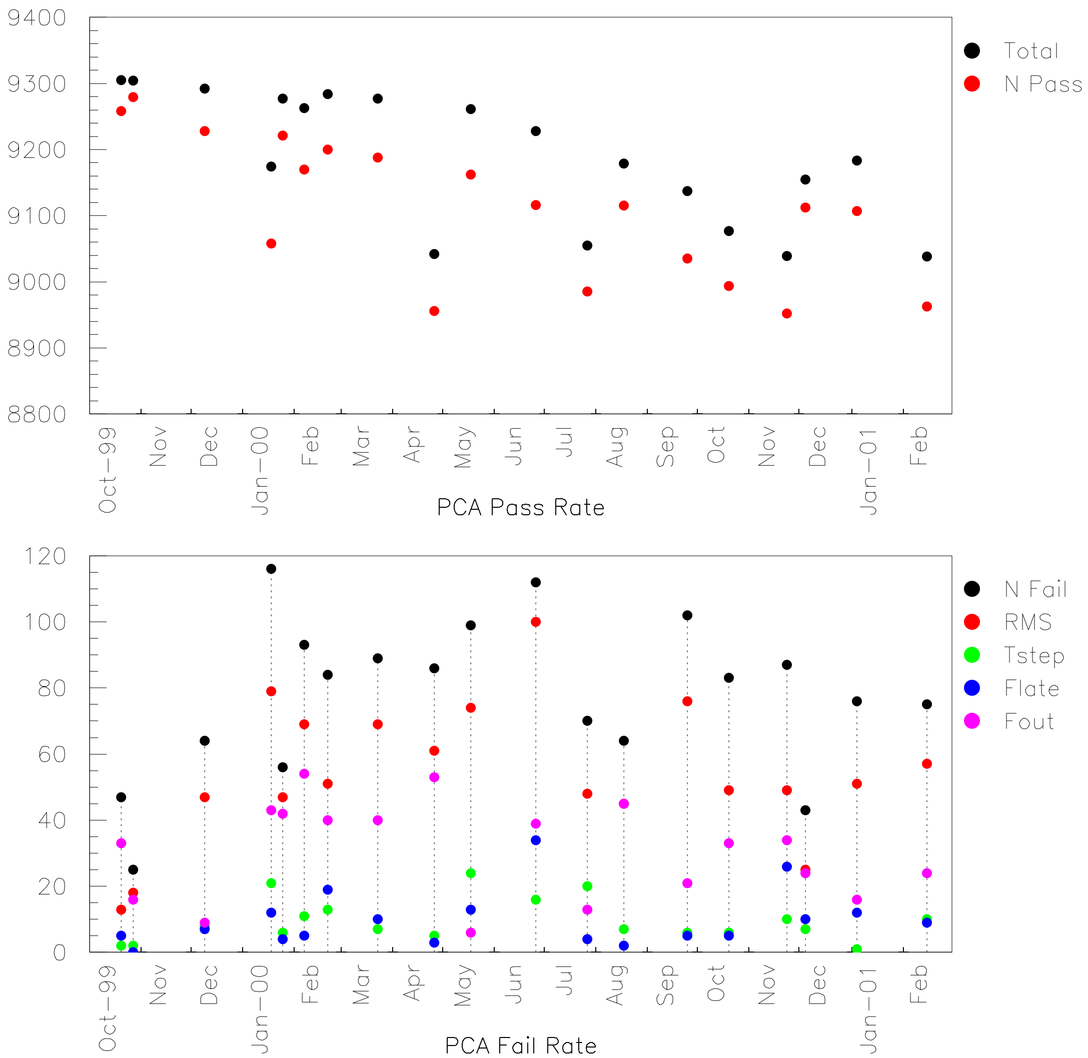

Figure 4.23: These plots summarise the pass and fail rate for all PCAs after the high voltage raise until mid-February 2001. In the upper plot the black points show the number of normal, online channels which produce at least 100 ECA calibrated hits. The red points show the number of channels which pass the PCA calibration. The lower plot gives a break down of the number of channels which fail each of the bad channel cuts. The black points indicate the total number of channels which fail, the dotted lines merely connect points pertaining to the same PCA run.

Chapter 5

The Performance of the PCA

This performance of the PMT calibration is investigated in this chapter. The chapter is divided into four main sections. Section 5.1 examines the adequacy of the PCA walk correction and looks at the timing of the laserball and other calibration sources. Section 5.2 is concerned with verifying the path correction used by the PCA. In particular the position of the laserball is fitted and compared to the position reported by the manipulator. The effect of the intensity of light on both the time and charge distributions is investigated in section 5.3. Finally the stability of the PCA constants is discussed in section 5.4.

There is a minor complication affecting plots of optical path corrected PCA calibrated time. When applying the PCA calibration, the calibrator adds an arbitrary 100 ns to the calibrated time for all channels. This 100 ns has been subtracted from the calibrated time in some plots in this chapter but not in others. If the prompt peak appears at 0 ns then the offset has been subtracted, if it appears at 100 ns then it has not.

5.1 Testing the Walk Calibration

This section investigates with how well the PCA calibration removes the discriminator walk and how the resulting calibrated times vary with charge.

The simplest test of the calibration is simply to apply it to the laserball run

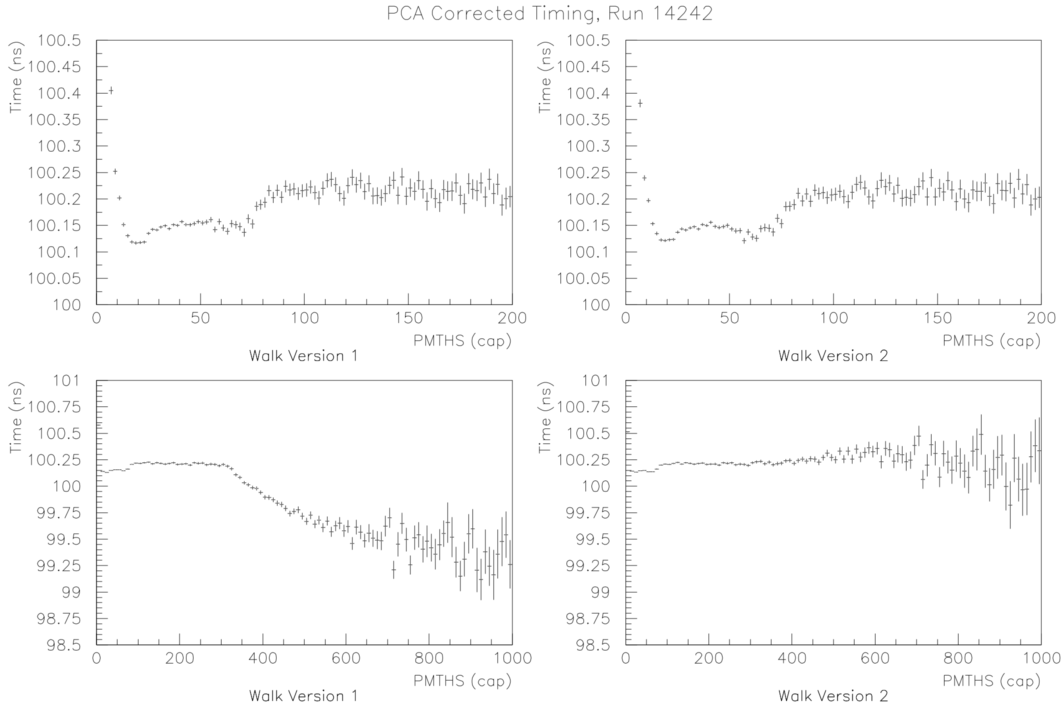

Figure 5.1: Profile histograms of the path corrected time of the prompt peak of laserball run 14242 for PCA versions 1 and 2. All times between 90 and 110 ns were admitted to the histogram.

which was used to generate the PCA constants and see whether the resulting calibrated time is independent of charge. In addition to PCA calibrating the PMT times, the optical path to the PMTs needs to be subtracted because the laserball might be slightly off centre and the PMTs are at slightly different radii. The subtracted optical path is identical to the optical path added during calibration so the resulting PMT time is simply the ECA time minus the interpolated walk time. Figure 5.1 shows profile histograms of time against charge for PCA run 14242 for both versions 1 and 2 of the walk calibration. Most hits have a charge of less than 100 counts above pedestal but there is a long tail to high charge so the correction is shown for two charge ranges. The main difference between walk versions 1 and 2 is the linear fit at high charge, the effect of which is very obvious above 300 counts above pedestal. At low charge there is a plough shape present in both versions of the

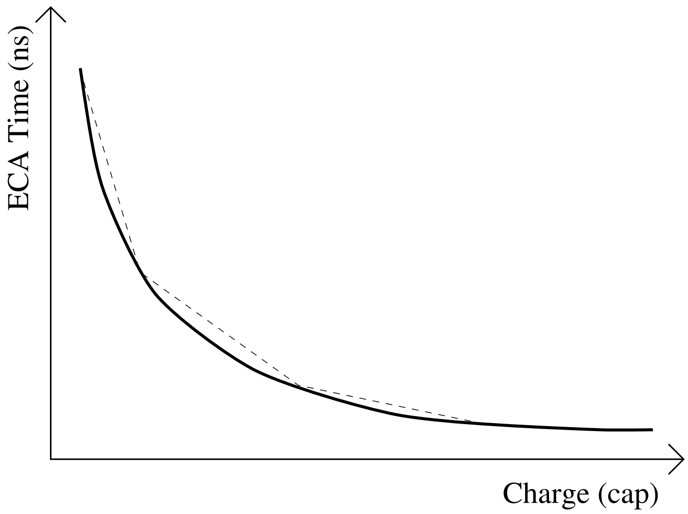

Figure 5.2: This diagram shows how a linear interpolation will slightly overestimate the value of the walk at low charge.

walk calibration with an amplitude of around 0.1 ns. This plough shape is probably due to the linear interpolation overestimating the height of the walk function at low charge, see figure 5.2. The time jitter of the PMTs is around 1.5 ns (this figure will be justified later) so correcting to better than 0.1 ns will not benefit analysis.

We have seen that the standard PCA calibration removes the discriminator walk to within 0.1 ns but how far could it be improved? The standard calibration uses a running median to find the times of the interpolation points, but would a simple average perform better? Perhaps QHL provides a less noisy measurement of the charge which fired the discriminator and would provide a better calibration. Perhaps using more interpolation points would provide a significantly better calibration. The PCA only produces an 'average' walk for the cells on a channels, so if there were slight errors in the measurement of the charge pedestals for different cells then this would result in the wrong walk being applied at low charge but not at high charge where the walk gradient has flattened off. Figure 4.8 demonstrated that the RMS of the time is larger at low charge than at high charge—could this be a consequence of cell pedestal variations? To address these questions five non-standard PCAs were done: a PCA which uses QHL instead of QHS, a PCA which uses the average time

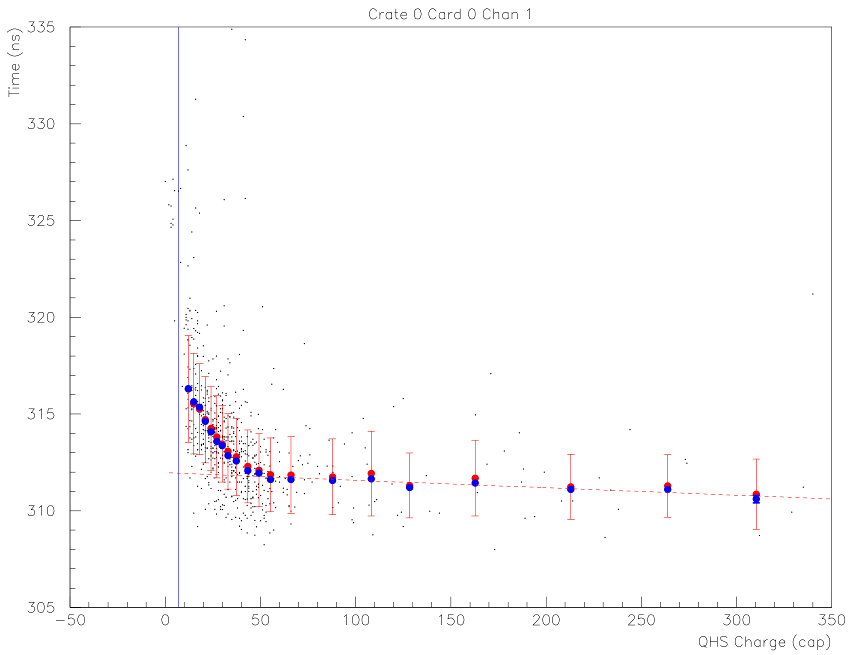

Figure 5.3: The interpolation points for the PCA with increased sampling. 20 interpolation points were used.

instead of the running median time, a PCA with 20 interpolation points instead of 10, a PCA which only uses data from cell zero and a PCA which randomly rejects 15/16ths of the data to compare with the cell zero PCA. Figure 5.3 shows a typical channel from the 20 interpolation point PCA.

The non-standard PCA calibrations were tested in two ways: a simple Gaussian fit was made to the overall prompt peak of the calibrated data and a set of Gaussian fits were made to a set of charge bins. The fit to the overall peak will be discussed first. For each of the calibrations, a histogram of the PCA calibrated and path corrected prompt peak was collected using the first sub-run of run 14242. A Gaussian was fitted, using PAW, to the peak over a range from -2.5 to +1.5 sigma, using a starting sigma of 1.5 ns and iterating 5 times. The range was curtailed at late times to avoid the non-Gaussian tail from scattered light. Figure 5.4 shows the Gaussian fit to the prompt peak calibrated using the normal PCA. The fit results

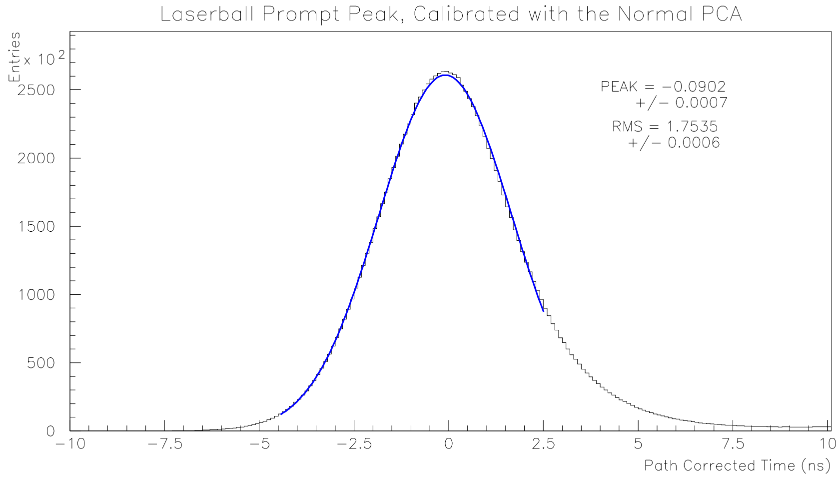

Figure 5.4: Example of a Gaussian fit to the prompt peak of the laserball. The fit range is from -2.5 to +1.5 sigma, using a starting sigma of 1.5 ns and iterating 5 times

Cal Type	Peak Height	Peak Position	RMS
Normal (QHS)	260765. ± 104.	-0.0902 ± 0.0007	1.7535 ± 0.0006
QHL	258027. ± 102.	-0.0954 ± 0.0007	1.7634 ± 0.0006
Average	261142. ± 103.	-0.2467 ± 0.0007	1.7446 ± 0.0006
20 Point	260656. ± 104.	-0.0675 ± 0.0007	1.7522 ± 0.0006
1/16th	16353. ± 25.	-0.1005 ± 0.0027	1.7420 ± 0.0025
Cell Zero	16395. ± 26.	-0.0956 ± 0.0027	1.7411 ± 0.0025

Table 5.1: Results of a Gaussian fit to the prompt peak of run 14242 subrun 0 for various PCA calibrations.

for each of the calibrations are summarised in table 5.1. The height of the fit is not as instructive as the RMS but is included for completeness. The central value does not indicate anything about the quality of the calibration because a constant offset will not affect reconstruction. It does show that the average walk calibration must be more influenced by the late tail because the corrected time is lower (the calibrated time goes as the ECA time minus the interpolated calibration time). The quantity of interest is the RMS and it is very similar for all the calibrations. The cell zero calibration and the 1/16th calibration have very similar widths indicating that variations in the pedestal calibration from cell to cell is a negligible effect. The

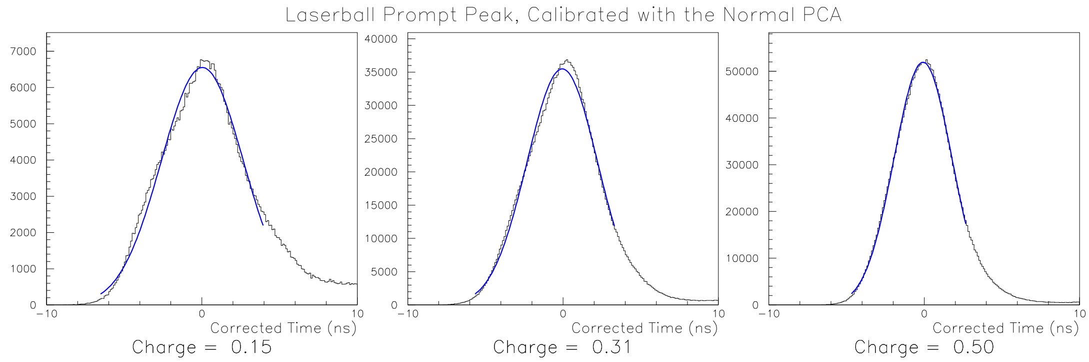

Figure 5.5: Histograms of the prompt peak of the laserball, calibrated using the normal PCA, for the charge bins 0–0.2, 0.2–0.4 and 0.4–0.6 in units of PCA calibrated QHS. The fit range of the Gaussian is from -2.5 to +1.5 sigma, using a starting value of 1.5 ns and iterating 5 times.

average calibration slightly outperforms the normal, running median, calibration. The reason for sticking with the running median will be discussed when examining the residual variation of the calibrated time against charge.

The second method used to examine the performance of the various types of PCA calibration was to look at the residual variation of the calibrated time against charge. This was done by collecting and fitting a set of prompt peak histograms for a range of charges. The charge bins, in units of PCA calibrated charge, were:

`0-0.2, 0.2-0.4, 0.4-0.6, 0.6-0.8, 0.8-1.0, 1.0-1.3, 1.3-1.6,`

`1.6-2.0, 2.0-3.0, 3.0-4.0, 4.0-5.0, 5.0-6.0, 6.0-8.0, 8.0-10.0`

The charge of the bin was taken to be the average charge of the entries in the bin. The fits to the first three bins using the normal PCA calibration are shown in figure 5.5. The first couple of histograms are not perfectly Gaussian. In addition to collecting the central value and sigma of the Gaussian fit, an average and RMS were computed for the histograms using all data between -3.0 and +3.0 sigma, where the sigma was taken from the fit. The two ways of characterising the position and width of the peak are to give an idea of the sensitivity of these values to the method used. The average and RMS include more of the late light tail than the Gaussian fit and so the average tends to be later and the RMS tends to be greater than the values from the fit. Figure 5.6 shows the fitted peaks and the peak average against charge for the various calibrations. The zero cell and 1/16th calibrations are similar to the normal calibration and are therefore not shown. The equivalent plots of the width against charge are indistinguishable, so only the one for the normal PCA is shown in figure 5.7. Comparing the normal PCA to the average PCA it can be seen that the average PCA suffers from a greater negative slope at low charge. This is probably because when the points are being collected in the PCA the same time range is used for all charge bins so bins with an earlier bin time include more of the late tail resulting in a negative slope in the calibrated times. Comparing the normal PCA to the 20 point PCA it can be seen that the increased sampling removes the 'plough' shape. Any new version of the PCA might benefit from increased sampling but care would have to be taken that all bins receive adequate statistics even for the shortest PCAs. It should be borne in mind that 0.1 ns is a tiny systematic and the improvement would not have any visible effect on analysis.

The RMS of the laserball in figure 5.7 does not show the true uncertainty in time of a hit from Čerenkov radiation because the laserball has an intrinsic width. To extract the timing of a Čerenkov source three runs were examined:

Run Number	Source
14273	High Rate ^{16}N
14276	Low Rate ^{16}N
14348	^{8}Li

To derive the residual calibrated time and prompt peak width against charge a method almost identical to the one for the laserball was used. The only difference is that the position and time at which the light is emitted is not known for the ^{16}N source and the time of the event is not known for the ^{8}Li. The ^{16}N is a gamma emitter and the gammas typically convert $\sim$60 cm from the source, see section 2.6. The source vertex and the event time were fitted using the path fitter [47][48]. The ^{8}Li is a beta emitter, see section 2.7. Although the source position is known for these events the events were still fitted varying all 4 parameters. As all these runs

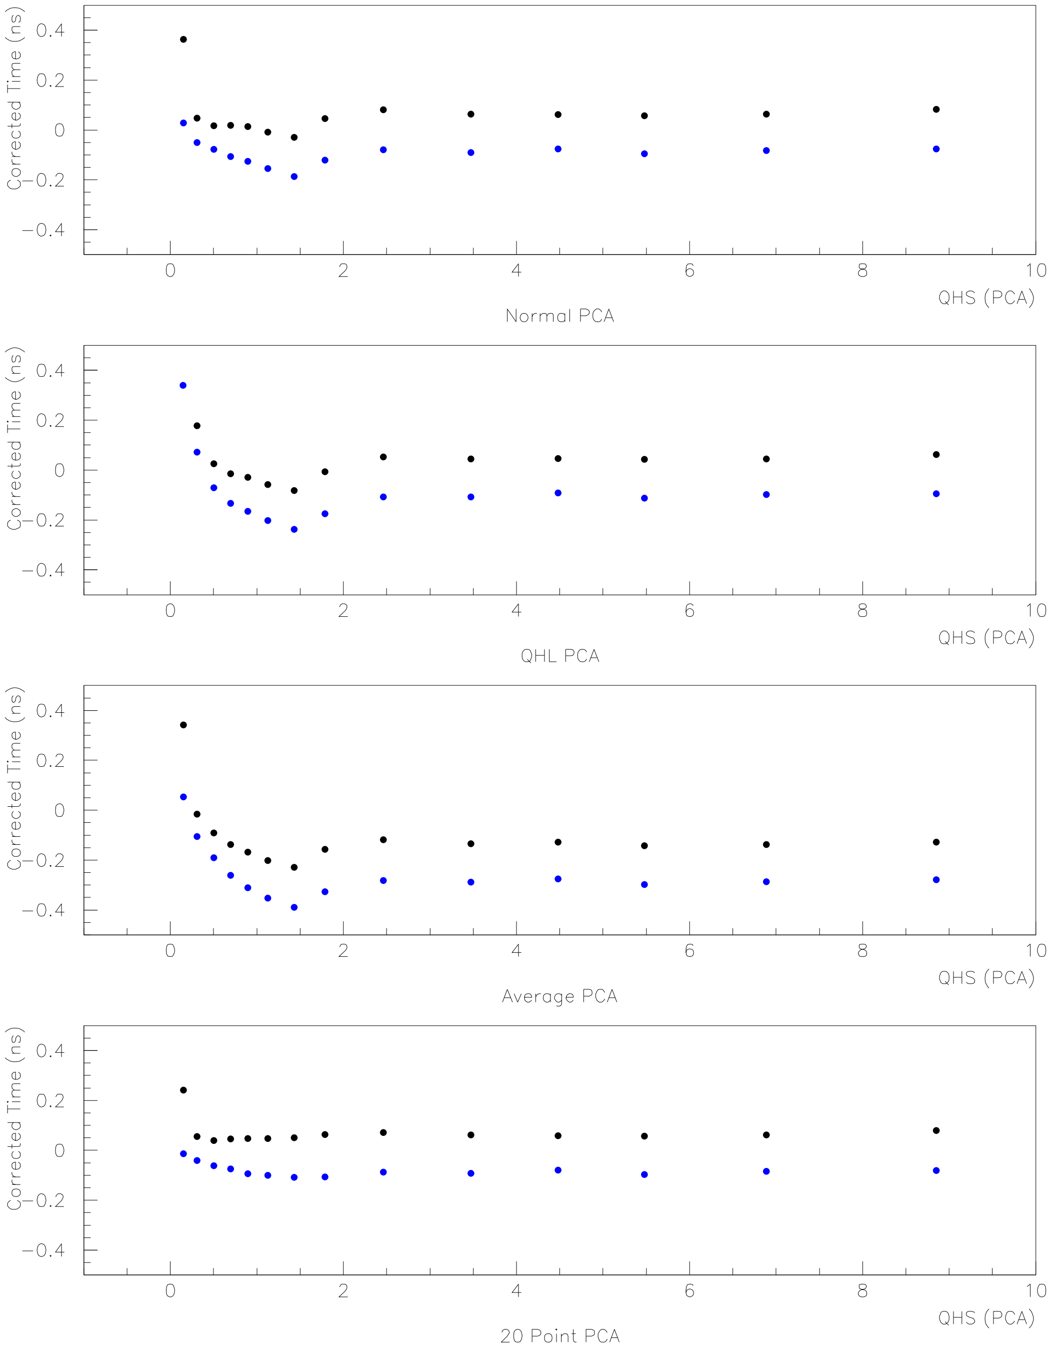

Figure 5.6: The residual calibrated time against charge for various PCA calibrations. The blue points are the result of a Gaussian fit while the black points are the result of an average. See the text for more details.

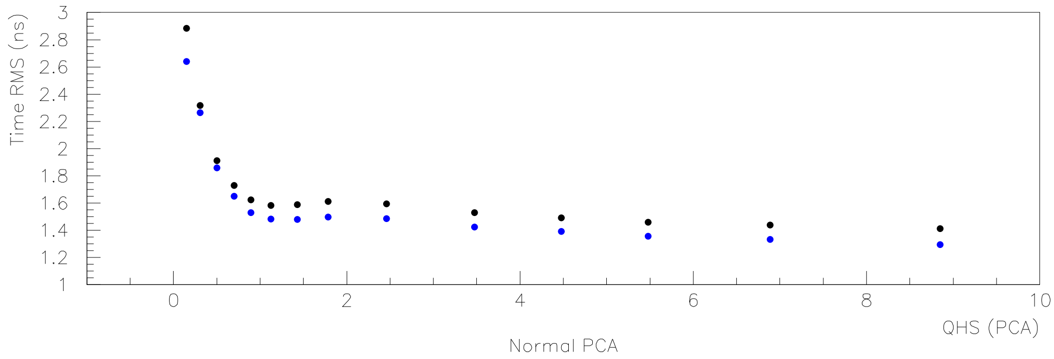

Figure 5.7: The width of the prompt peak of the laserball as a function of charge. The blue points show the sigma of the fitted Gaussian fit while the black points show the RMS of the data between -3 and +3 sigma. See the text for more details.

were taken very close to the centre of the detector, any event which reconstructed more than 200 cm from the centre was rejected. In addition to the normal PCA calibration the low rate ^{16}N and ^{8}Li times were also history calibrated—the history calibrator will be discussed in the next chapter. Figure 5.8 shows the residual PCA calibrated and path corrected time against charge. The high rate ^{16}N source is clearly reasonably flat, however the low rate ^{16}N and the ^{8}Li, also a low rate source, have a curve which spans almost a nano-second. The reason for this curvature is that the charge pedestal is lower at low rate than it is at high rate. The rate dependence of the pedestals will be discussed further in the next chapter. Figure 5.9 shows that the curve has the expected sign. The magnitude is also plausible; the gradient of the walk at low charge is around 0.2 ns/ADC count and the charge spectrum slides down by around 4 counts from high to low rate giving an expected shift of around 0.8 ns at low charge.

A Gaussian was fitted to the overall prompt peak of the ^{16}N and ^{8}Li sources. The method was the same as that used for the laserball, adjusting the fit range to between -2.5 and +1.5 sigma using a starting value of 1.5 ns and iterating 5 times. Table 5.2 gives the fitted peak positions and sigma. It can be seen that the

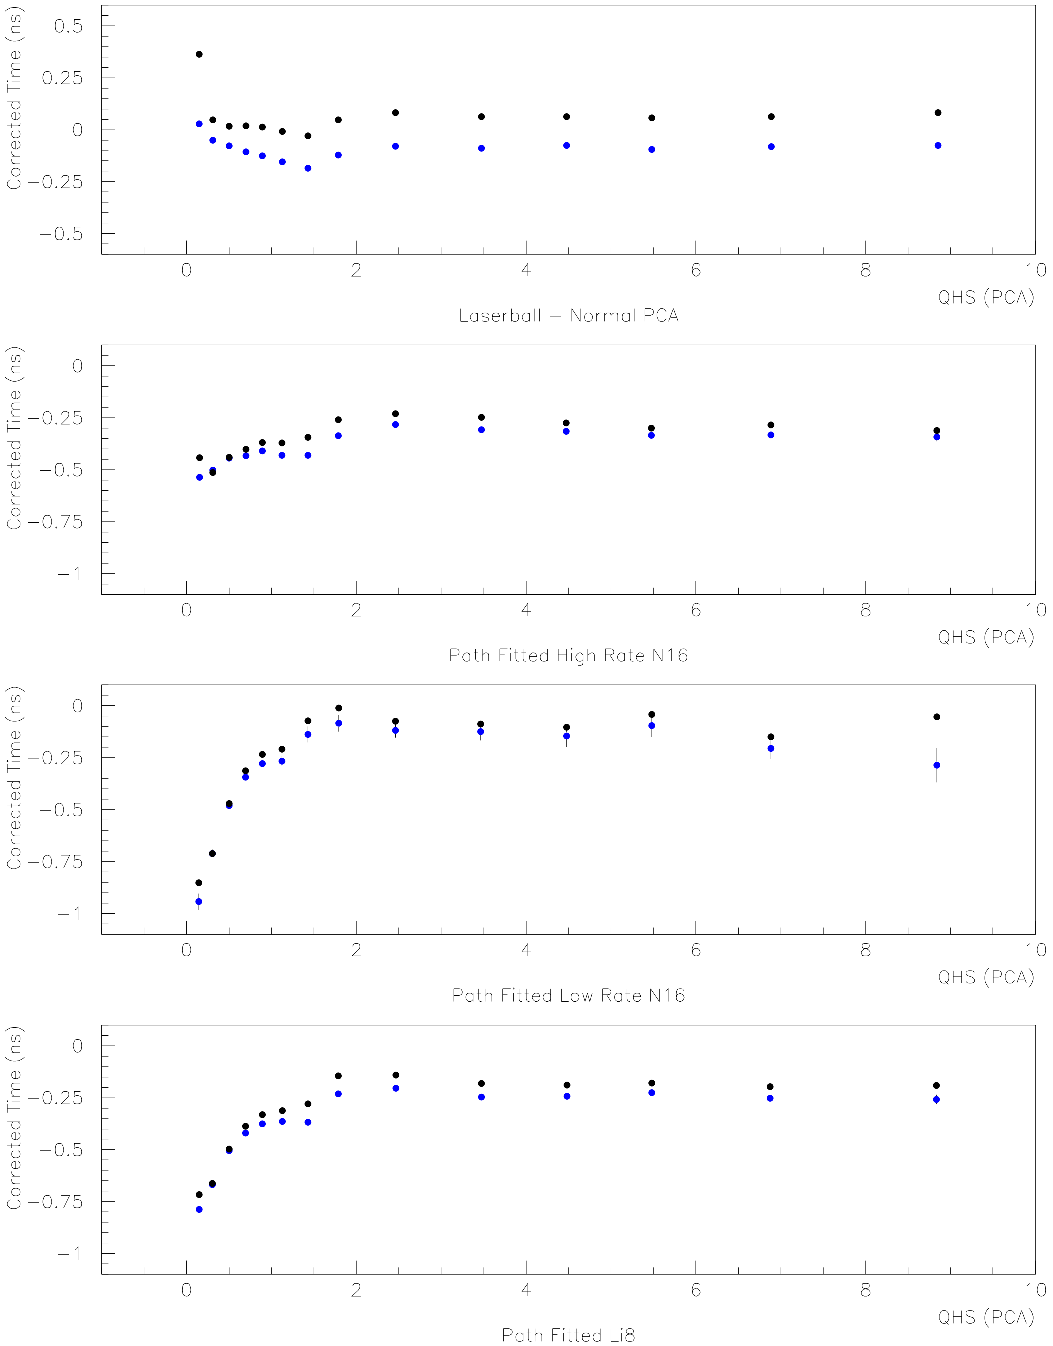

Figure 5.8: The residual PCA calibrated time against charge for various sources. The blue points are the result of a Gaussian fit while the black points are the result of an average. See the text for more details.

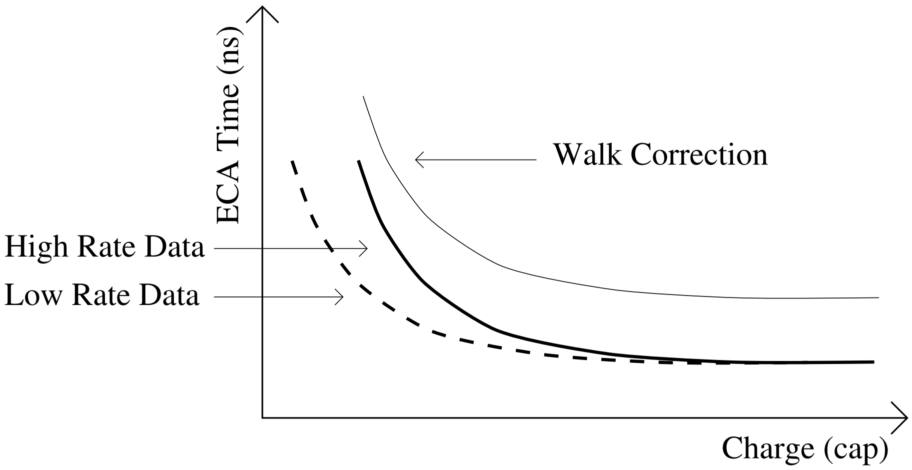

Figure 5.9: Schematic showing that the curve seen in calibrated low rate data has the expected sign. The pedestal at low rate is lower than the pedestal at high rate so the charge spectrum for low rate data shifts down. The walk correction, which gets subtracted from the data during calibration, is collected at high rate and therefore too much time gets subtracted off at low charge for low rate data.

Source	Peak Position	RMS
Laserball	-0.0902 ± 0.0007	1.7535 ± 0.0006
High Rate ^{16}N	-0.4271 ± 0.0019	1.5213 ± 0.0018
Low Rate ^{16}N	-0.4008 ± 0.0069	1.5679 ± 0.0066
^{8}Li	-0.4439 ± 0.0021	1.5349 ± 0.0020

Table 5.2: Results of a Gaussian fit to the prompt peak of various sources, see the text for more details.

timing performance does degrade slightly at low rate because of the charge pedestals changing. It is clear that the Čerenkov sources are somewhat sharper in time than the laserball, this is due to the intrinsic time width of the laserball. Figure 5.10 shows the width of the prompt peak against charge for the Čerenkov sources and the laserball.

This section has focused on the prompt peak of the timing distribution of the sources. Figure 2.3 on page 32 showed the overall timing distribution from the laserball. The Čerenkov sources produce very similar distributions, see figure 5.11. Figure 5.12 shows the laserball time spectra at various wavelengths. The most obvious feature of these spectra is that the late tail is much longer at short wavelengths.

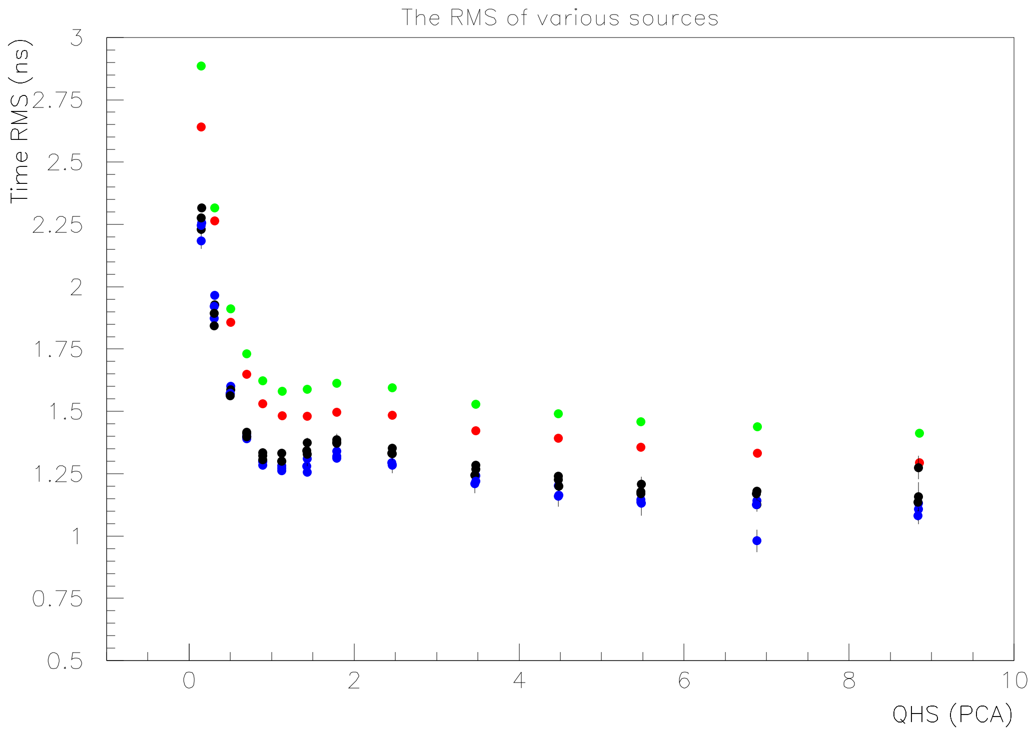

Figure 5.10: The width of the prompt peak versus charge for Čerenkov sources and the laserball. For the Čerenkov sources (high rate ^{16}N, the low rate ^{16}N and ^{8}Li) the blue points indicate the result of Gaussian fits while the black points are an average. For the laserball the result of the Gaussian fits are plotted in red while the average is plotted in green.

A possible explanation for this tail is fluorescence in the optical fibres which carry light from the laser to the diffuser ball. PCA calibrations are done at 500 nm because it was the longest wavelength, and therefore the sharpest timing, available when the first calibrations were taken. There is one unusual feature of these time spectra, an anomalous extra peak around 22 ns after the prompt peak. This peak is peculiar to these runs and does not normally appear in time spectra.

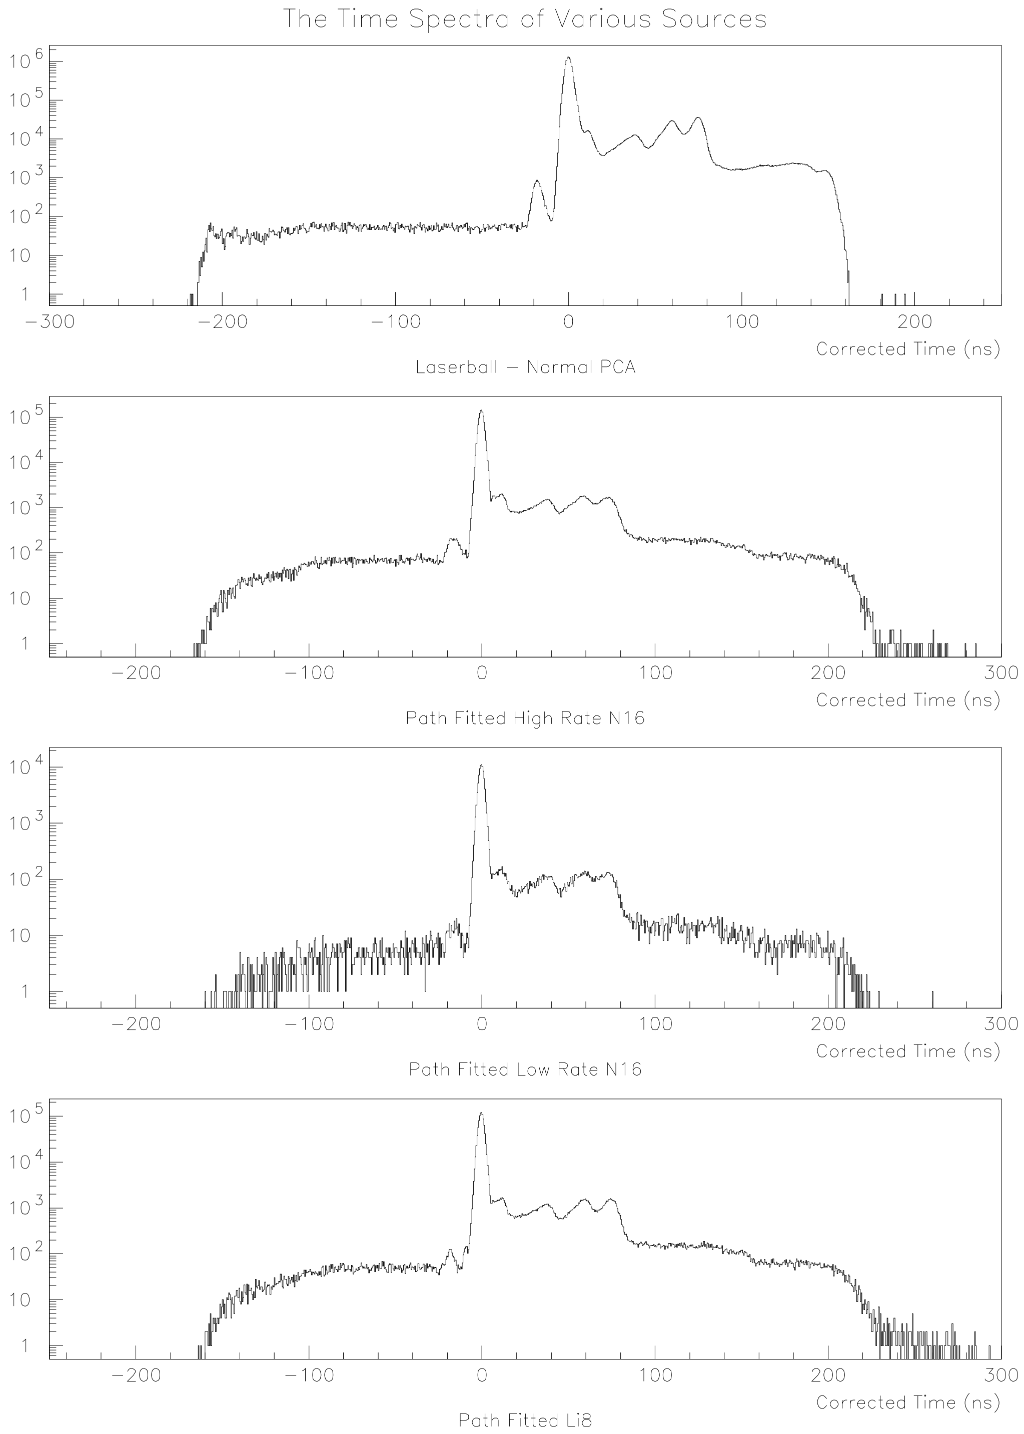

Figure 5.11: The time spectra of various sources.

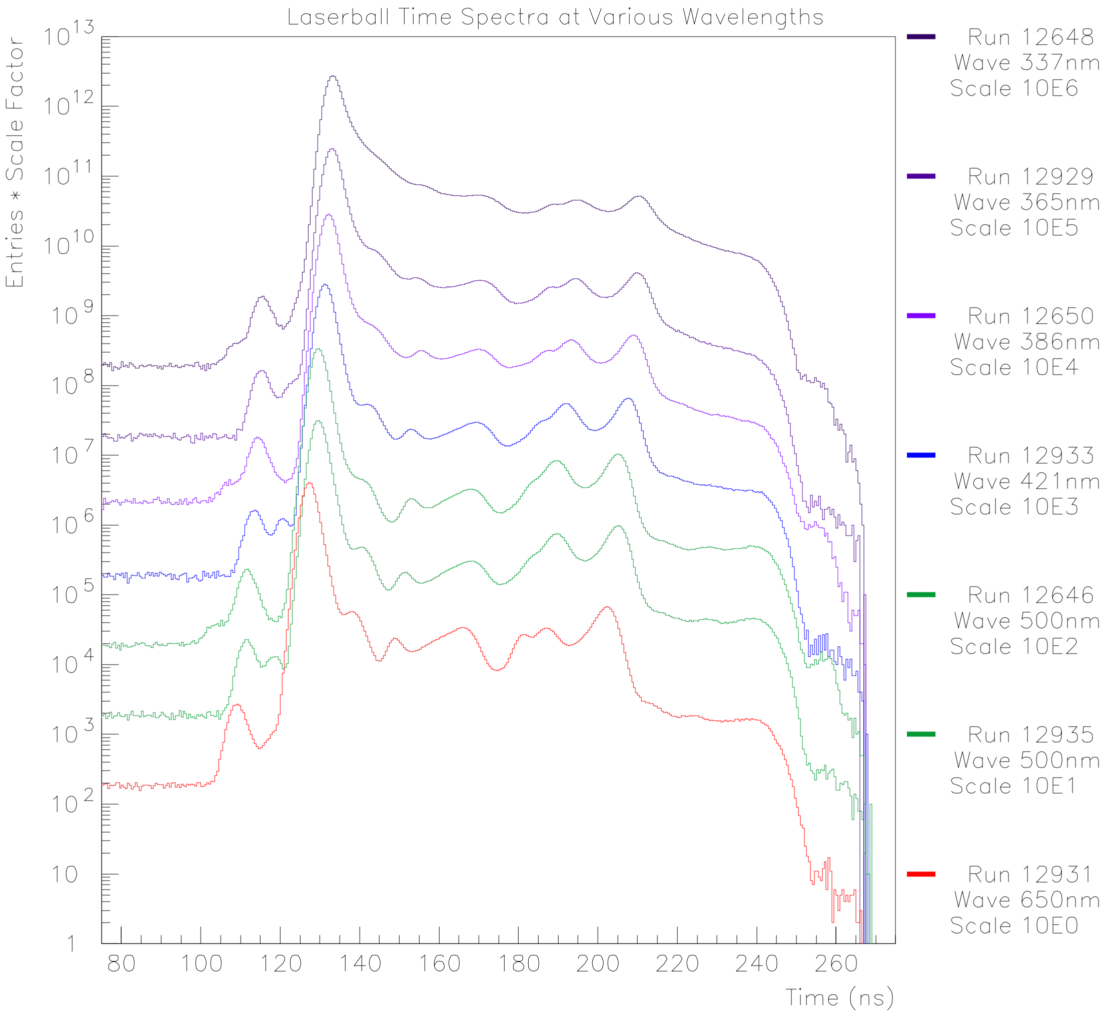

Figure 5.12: The laserball time spectrum at various wavelengths.

5.2 The Laserball Path Correction

This section is concerned with verifying the optical path correction used by the PCA. To calculate the optical path from the laserball to the PMTs the PCA assumes a simple straight line path. How accurate is this calculation? Figure 5.13 shows the number of times PMTs were hit during a typical PCA run as a function of Z and with position in the detector. The flat map highlights some interesting features. The 10 rings of high and low occupancy tubes are due to distortions introduced by the belly plates. (The belly plates are plates of acrylic positioned around the equator of the acrylic vessel which have circular grooves in them. The 10 loops of rope which support the acrylic vessel are housed in these grooves.) Also visible is a line running from top to bottom in the middle of the plot, this is caused by some water pipes which run down the outside of the acrylic vessel. How much do these structures affect the PCA?

The adequacy of the optical path correction is investigated in section 5.2.1 using the Monte Carlo and in section 5.2.2 using real data. Section 5.2.3 addresses the important question of whether the laserball is where the manipulator system reports it to be.

5.2.1 Testing the Path Correction in the Monte Carlo

This section examines the accuracy of the straight line optical path correction using the Monte Carlo. An isotropic light source was simulated, using the Monte Carlo, at X=0,Y=0,Z=-100. To make the prompt peak as sharp and distinctive as possible Rayleigh scattering was turned off, the PMT noise rate was set to zero and the analysis used the unjittered times. Histograms of the straight line path corrected times were collected for every PMT in the detector using bin widths of 0.2 ns. The path corrected prompt peak time was found by finding the bin with the most entries and calculating the average of all times within 3 ns. Figure 5.14 shows the path corrected prompt peak time and the number of entries in the prompt peak as a

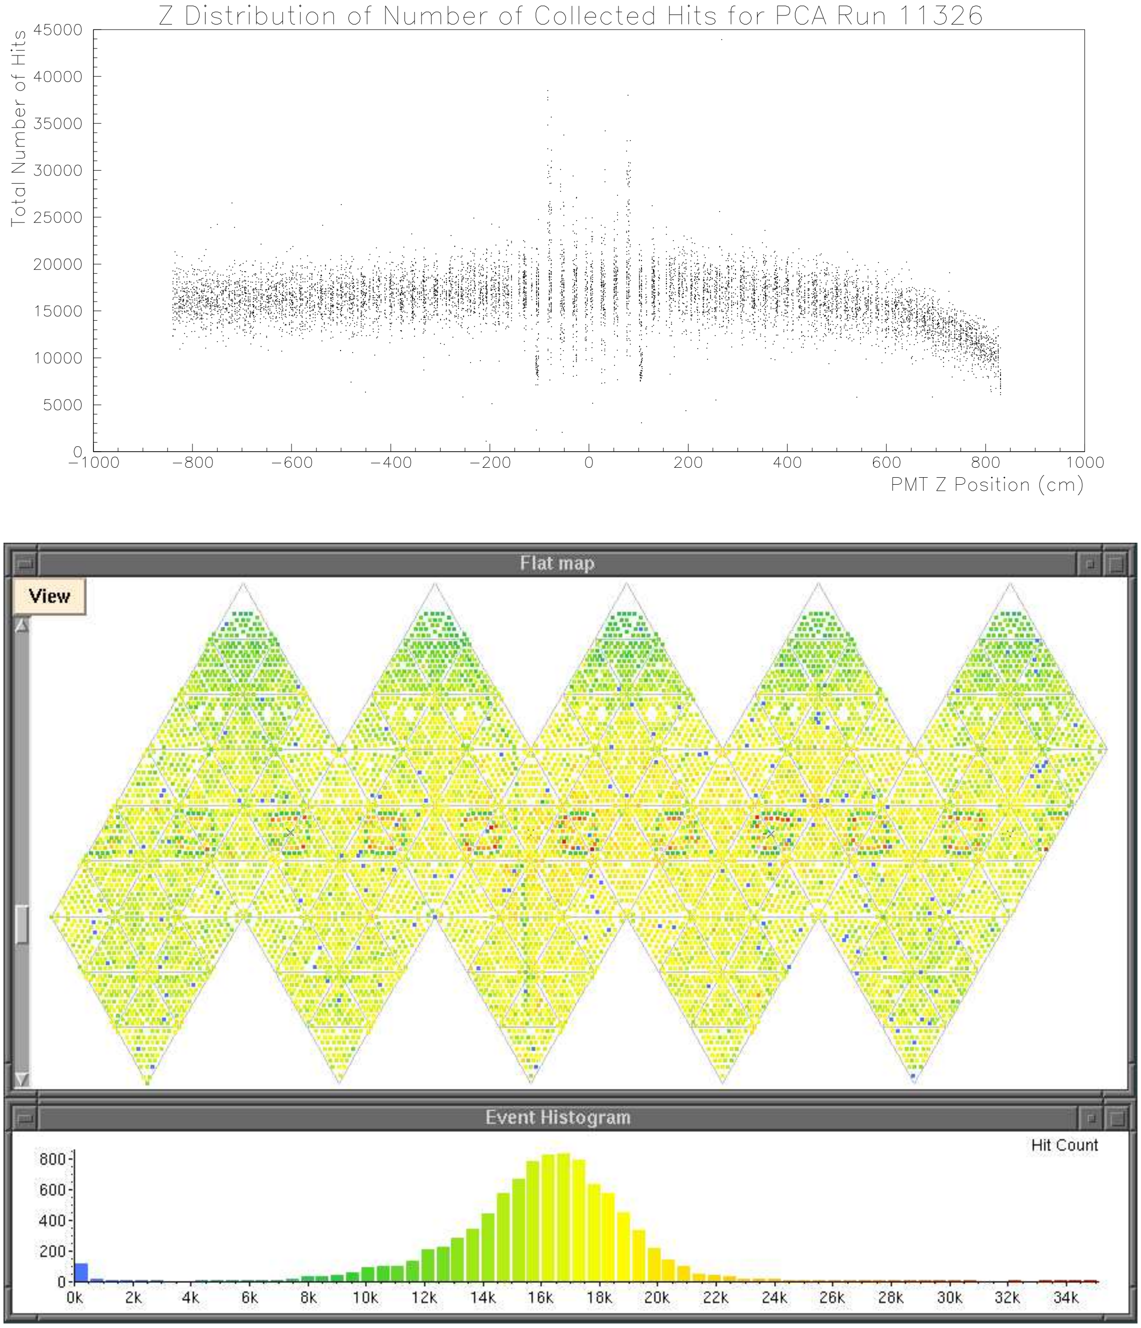

Figure 5.13: The upper plot shows the total number of hits collected by all normal online channels versus PMT Z for PCA run 11326. It can be seen that the intensity of the laserball drops off near the top of the detector. The lower plot shows an opened out display of the detector for the same run. The colour of the points indicates the total number of times the PMT fired. All normal PMTs are shown regardless if they are online or offline.

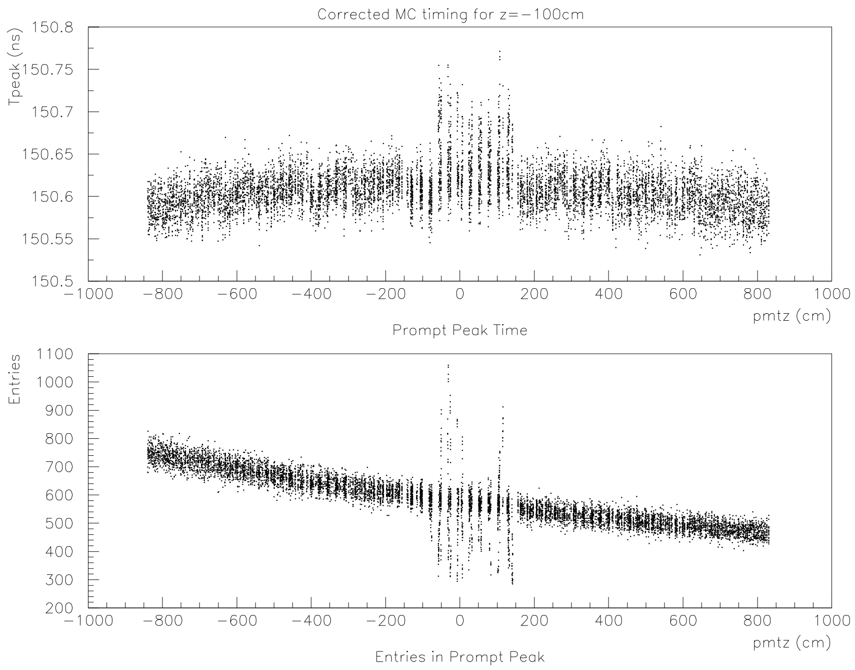

Figure 5.14: The upper plot shows the straight line path corrected prompt time against Z position for of an isotropic source as generated in the Monte Carlo with full geometry. The lower plot shows the number of hits recorded by each PMT as a function of Z. The effect of the belly plates around the equator is small but visible.

function of Z. The effect of the belly plates is visible but very small in the Monte Carlo.

5.2.2 Testing the Path Correction with Data

Real laserball data provides a good way to test the straight line optical path correction used the PCA. Laserball runs were selected from a Z scan and a later X/Z scan, see table 5.3. For each of the laserball runs, histograms of the PCA calibrated time were collected for all calibrated channels in the detector. The histograms were then fitted with a Gaussian from -2.5 to +1.5 sigma, using a starting sigma of 2 ns and iterating 5 times. The fitted peak position was then path corrected. This path correction was the same straight line path correction used by the calibrator (using the appropriate group velocity for that wavelength). If the straight line path correction

Run Number	λ	X	Y	Z	Occupancy	PMT Rate (Hz)
10988	500.	0.	0.	0.	0.056	2.57
11004	500.	0.	0.	-400.	0.038	1.35
11006	500.	0.	0.	-200.	0.044	1.56
11012	500.	0.	0.	200.	0.045	1.60
11014	500.	0.	0.	400.	0.040	1.43
11030	500.	0.	0.	0.	0.051	2.32
10990	365.	0.	0.	0.	0.053	2.42
11002	365.	0.	0.	-400.	0.035	1.27
11008	365.	0.	0.	-200.	0.042	1.51
11010	365.	0.	0.	200.	0.040	1.42
11016	365.	0.	0.	400.	0.033	1.17
11027	365.	0.	0.	0.	0.037	1.34
12848	500.	0.	0.	0.	0.050	2.06
12862	500.	-70.	0.	-70.	0.052	2.12
12937	500.	70.	0.	70.	0.050	2.04

Table 5.3: The laserball runs used to investigate the path correction. The quoted positions are approximate. The occupancy is the typical fraction of PMTs which fire in an event. The PMT rate is the read out rate and not the firing rate. Run 10988 is a PCA calibration run and all runs starting 11 were taken within 24 hours of it. The runs which start 12 were taken about three weeks after a PCA.

is inadequate for any channels then these will show up as outliers.

Figure 5.15 shows the path corrected prompt peak time against logical channel number and against PMT X, Y and Z for runs 10988 and 11030. Both runs were taken at the centre of the detector with a wavelength of 500 nm. The plots for the PCA run 10988 are gratifyingly flat, as might be expected, while run 11030 contains a number of outliers. It can be seen that the outliers are clustered in logical channel number rather than with position in the detector. The bad cards and paddle cards were identified by hand, there are 193 bad channels out of the 9169 channels fitted. The bad channels are plotted in red with the good channels plotted on top in black. These bad channels will be discussed further later. In addition to the outliers there is a slightly increased spread in from run 10988 to run 11030 for all channels.

Having identified the bad channels, the Z scan can be examined for other effects. Figure 5.16 shows the path corrected time against PMT Z for the remaining

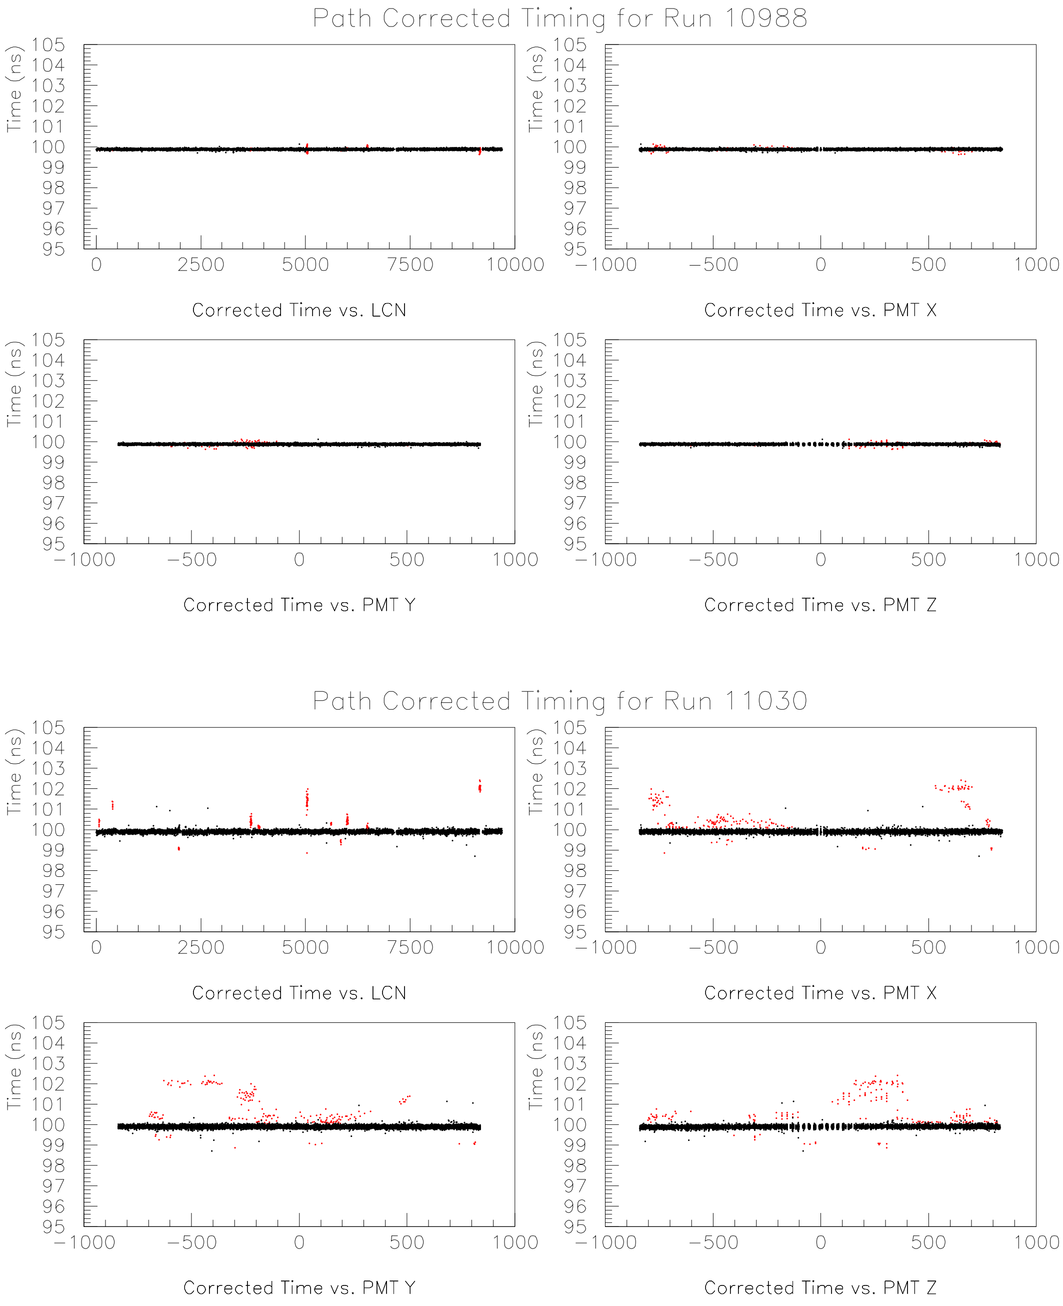

Figure 5.15: The path corrected time against logical channel number and against PMT X, Y and Z for runs 10988 and 11030. The laserball is in the centre of the detector and at a wavelength of 500 nm in both cases. Run 10988 is the PCA run used to calibrate the time. Bad channels are plotted in red (193 out of 9169).

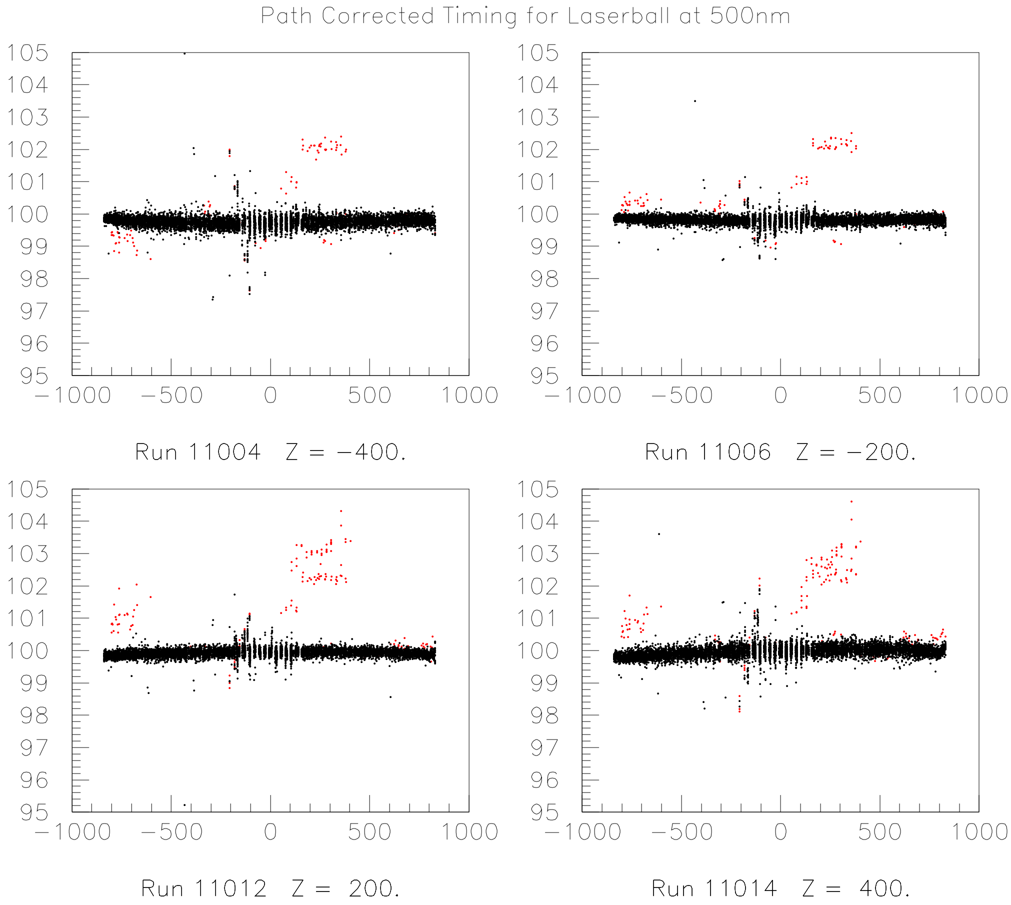

Figure 5.16: The path corrected prompt peak time versus PMT Z for a set of runs at 500 nm from a laserball Z scan. Bad channels are plotted in red. The effect of the belly plates is visible around PMT Z = 0 for the off centre runs.

500 nm runs from the Z scan. Figure 5.17 shows the 365 nm Z scan runs. The channels which were identified as bad are plotted in red. It can be seen that for the off centre runs there are some PMT outliers near the equator. Similar plots can be made in ϕ (not shown) and these show the outliers are in 10 groups equally spaced in ϕ. Clearly these are due to distortions introduced by the belly plates. The maximum size of this effect is around a nano-second. There is noticeable broadening in the width and a slight curvature in the path corrected time of runs which are a long way off centre. This is a negligible effect for the path correction of the PCA

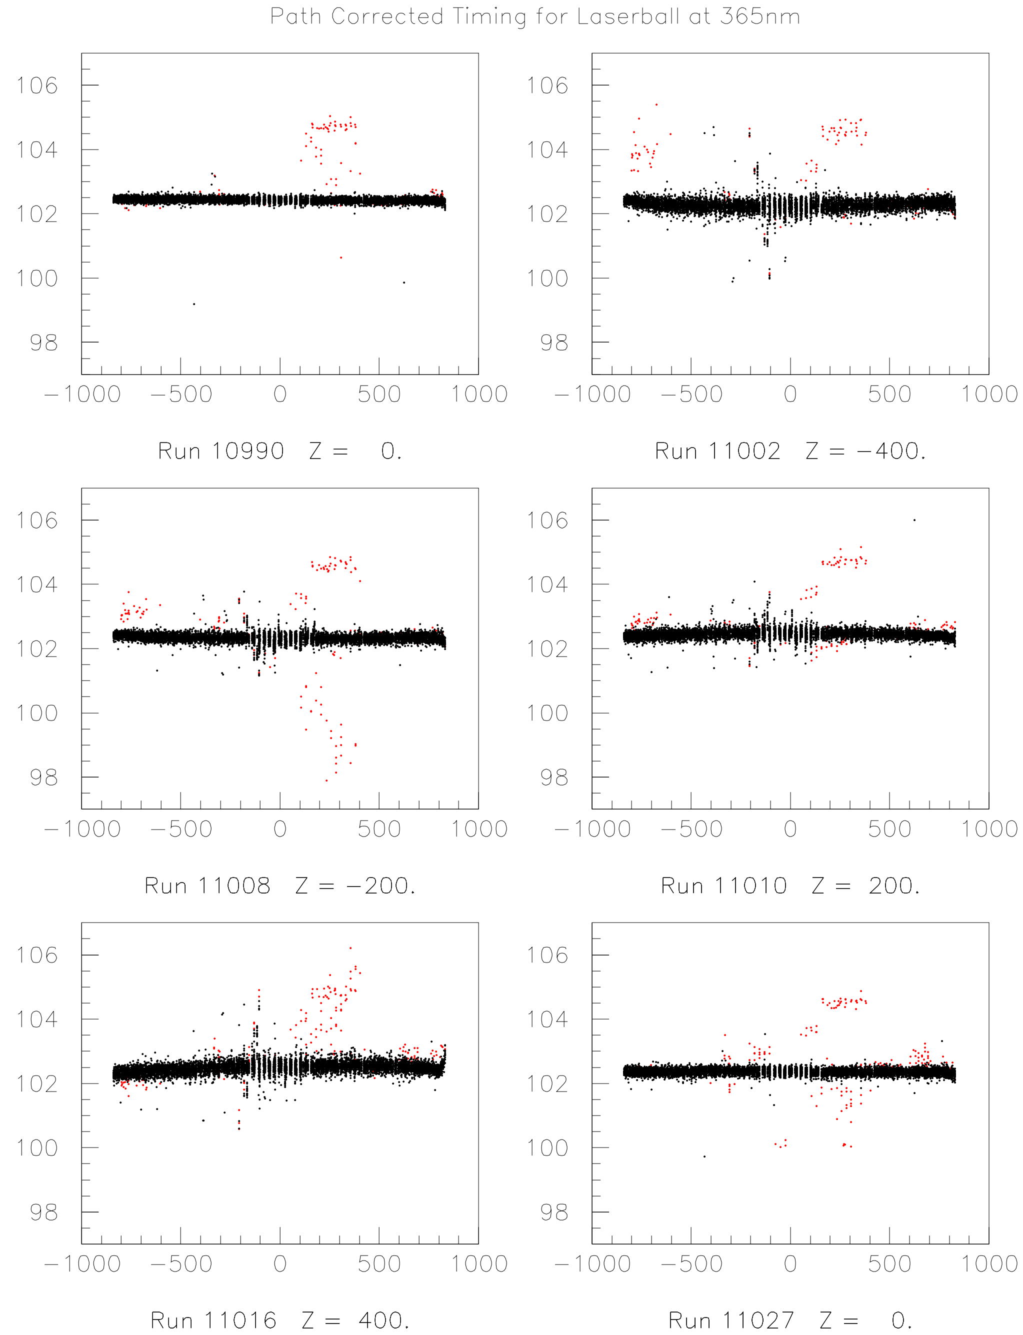

Figure 5.17: The path corrected prompt peak time versus PMT Z for a set of runs at 365 nm from a laserball Z scan. Bad channels are plotted in red. The effect of the belly plates is visible around PMT Z = 0 for the off centre runs.

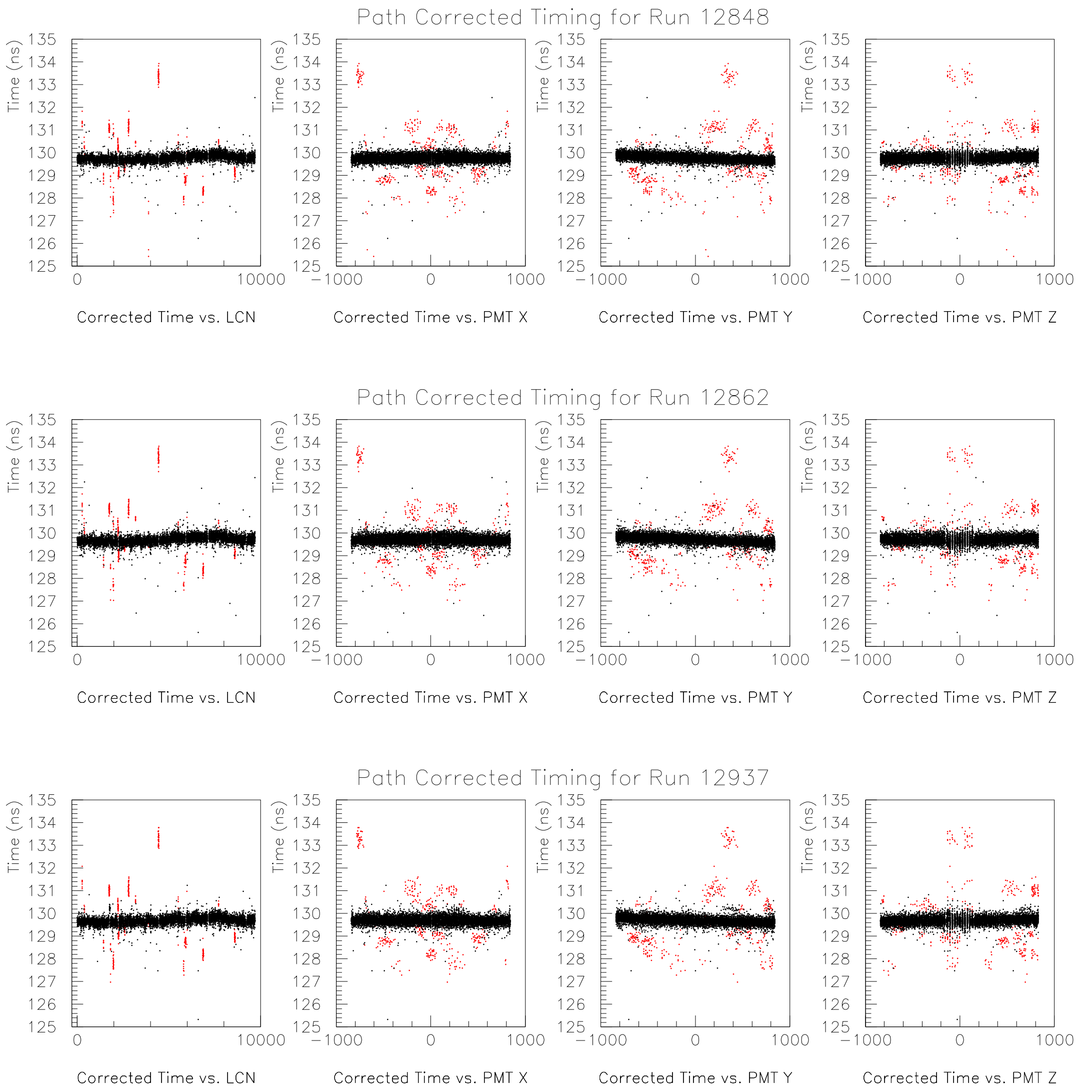

Figure 5.18: The path corrected prompt peak time against logical channel number and PMT X, Y and Z for runs 12848, 12862 and 12937. In run 12848, the uppermost plot, the laserball was in the centre. In the other two runs the laserball was approximately 1 metre off centre. The bad channels are plotted in red, there are 447 of of the 9113 plotted.

because all PCAs taken so far are within 30 cm of the centre of the detector.

Figure 5.18 shows the path corrected prompt peak time against logical channel number and against PMT X, Y and Z for runs 12848, 12862 and 12937. In run 12848 the laserball was at the centre of the detector and this run was used to identify the bad channels. Again the outliers are clustered in logical channel number rather than with position in the detector. These bad cards were again identified by hand for removal from the other runs and are plotted in red. This time there are 447 bad channels out of a total of 9113 fitted. The spread in the corrected time for the run in the centre, 12848, is greater than the width for run 11030, 10990 and 11027. This could be due to drifting in the pedestals over the three weeks from the last PCA calibration. The effect of the belly plates is hard to distinguish and is less than a nano-second.

Figures 5.15, 5.16, 5.17 and 5.18 have a number of bad channels which were plotted in red. The question arises as to whether these bad channels are cut by the ECA validation cuts. Runs 11030 and 12848 were put through the same analysis but with the ECA validation cuts at full (setting words 15 and 16 of the TCAL bank to 4 and 2 respectively, see [43][35]). For run 11030, with the ECA validation cuts applied, the number of channels which had been identified as bad and were still present was 100 out of 8585 (without ECA cuts the number was 193 out of 9169). For run 12854, with the cuts, the number of bad channels was 285 out of 8794 channels. (without ECA cuts the number was 447 out of 9113). See figure 5.19. The cuts have removed some of the bad channels, but are not a silver bullet.

5.2.3 The Reflection Fitter

In order to get the path correction to the laserball right, it is obviously essential to know where the laserball is. The time offsets of all the PMTs are adjusted assuming the position reported by the manipulator system is accurate. If the position is incorrectly reported to be, say, 10 cm below where it actually is then an ordinary fitter will reconstruct everything 10 cm low, including the laserball. Clearly it is important to cross check the position given by the manipulator system.

There are two methods which can reconstruct the laserball position without a knowledge of the relative timing of the PMTs. One method examines the change

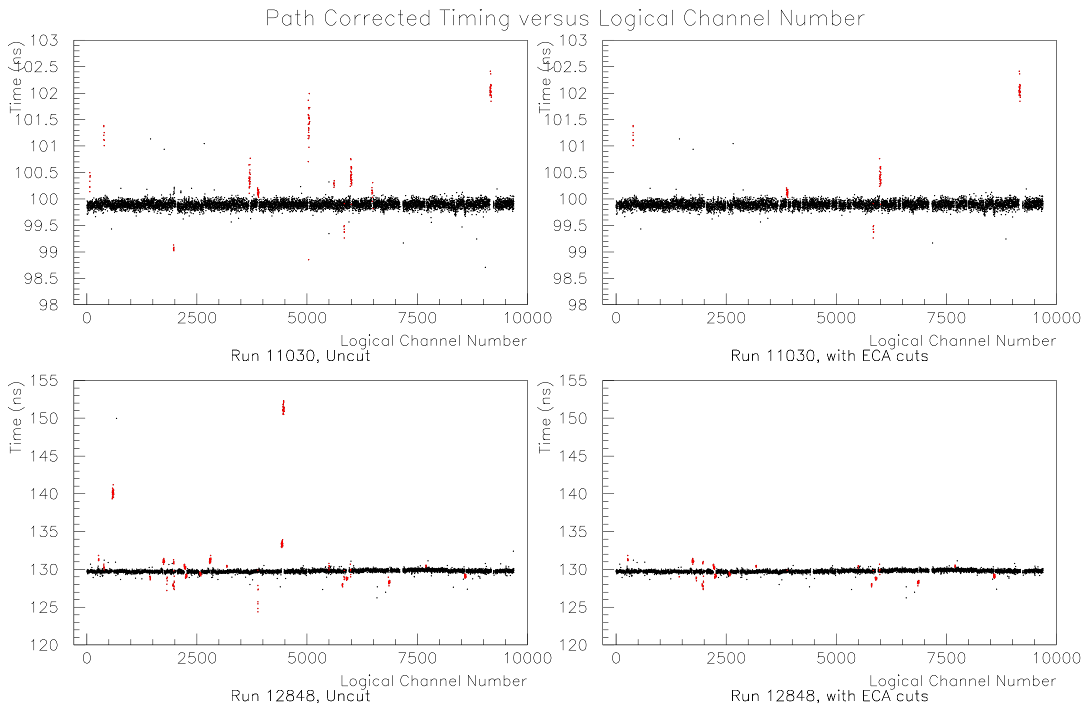

Figure 5.19: The path corrected prompt peak time against logical channel number for runs 11030 and 12848. The channels in red are bad channels identified by hand. The plots on the left are without any ECA validation cuts. The plots on the right were generated with the ECA validation cuts at full.

in the prompt time for the same PMT but with the laserball in different positions. The other method looks at the time difference between the prompt peak and the 180° reflection peak for all PMTs. The first method, done by Xin Chen, has the advantage that it uses the high statistics prompt peak and should work for a wide variety of positions in the detector. The second method will only work near the centre of the detector because the reflection peak must be identifiable. The second method only requires a single laserball run near the centre of the detector and so is suitable for checking the PCA.

Figure 5.20 shows the basic principle of the reflection fit. For a PMT with position $\underline{P}$ the time difference between the reflection peak and the prompt peak minus twice the radial crossing time of the detector, for a laserball at position $\underline{\delta}$, is

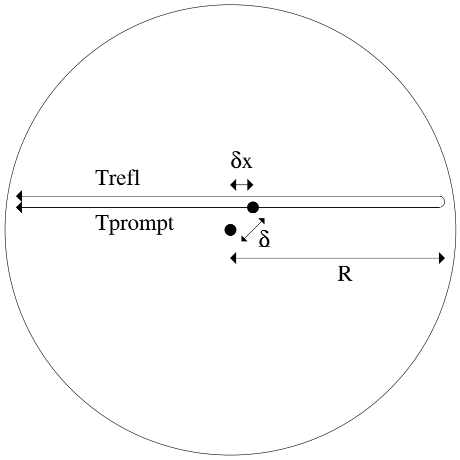

Figure 5.20: If the laserball is a small distance $\underline{\delta}$ away from the centre then the time difference between the 180° reflection peak and the prompt peak for a PMT which coincides with the negative x axis is given by $\mathrm{Trefl} - \mathrm{Tprompt} = 2R/c - 2\delta_x/c$ where c is the group velocity of light in the D_2O.

given by:

$$T_r - T_p - \frac{2R}{c} = \frac{2P_x}{|\underline{P}|c}\delta_x + \frac{2P_y}{|\underline{P}|c}\delta_y + \frac{2P_z}{|\underline{P}|c}\delta_z + \delta_t \tag{5.1}$$

where c is the group velocity of light in the D_2O. δ_t should in principle be zero but is there to allow a slight bias in the calculation of the left hand side of the equation. All vectors are from the geometrical centre of the PSUP. If the left hand side of equation 5.1 were measured for all PMTs, it would constitute a data set with which to perform a linear least squares fit to δ_x, δ_y, δ_z and δ_t.

The immediate problem with fitting the position of the reflection peak is that the statistics are low for a single tube. It is possible to combine the time spectra of PMTs in the same region of the detector by lining up their prompt peaks. The position of the prompt peak as a function of charge is recorded by the PCA. The calibrator gives the time of the PMT relative to the prompt peak with the addition of a path correction. Subtracting the path correction again gives the times of the PMTs relative to the prompt peak. A convenient way to divide the detector into regions is to use the geodesic PMT support structure, the PSUP. The PSUP has 92 vertices evenly spaced around the detector. Assigning every PMT to the nearest

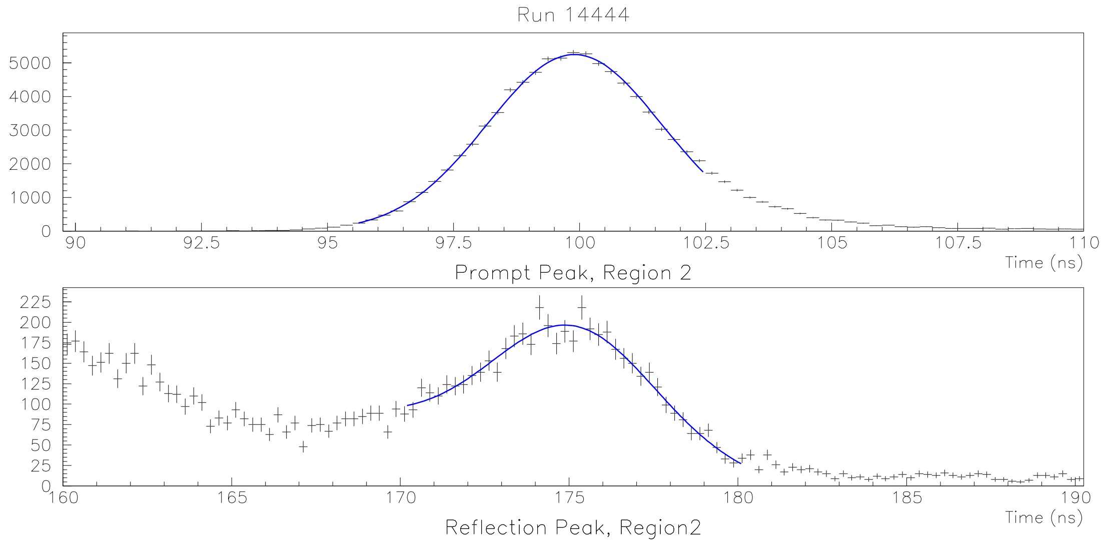

Figure 5.21: These two plots show two portions of the region 2 time histogram, collected from the first subrun of run 14444. The upper plot shows the prompt peak and a Gaussian fit. The lower plot shows the 180° reflection peak with a Gaussian plus a linear background fit.

vertex results in 92 regions containing around 100 PMTs each. The position of the PMT in equation 5.1 is replaced by the resultant of all the vectors to the PMTs in a region.

Figure 5.21 shows a histogram of the prompt peak and a histogram of the reflection peak for region 2 (region 1 corresponds to the position of the neck of the acrylic vessel and only includes 34 PMTs). The data was collected from the first sub-run of PCA run 14444. The first sub-run of every PCA is kept on disk so being able to fit using only one sub-run considerably reduces the effort required to fit all PCAs. Figure 5.21 also shows the result of fits to the peaks. To fit the prompt peak, the starting values of the height and position were simply the height and position of the bin with the most entries. The prompt peak was then fitted with a Gaussian running from -2.5 to +1.5 sigma, using a starting value of 2 ns for sigma and iterating 5 times. Finding the starting values of the position and height of the reflection peak is more involved. With the laserball in the centre of the detector the reflection peak comes about 76 ns after the prompt peak but if the laserball is 150 cm off centre then the peak could be anywhere between 62.5 ns and 89.5 ns later. With the laserball in the centre, the 180° reflection peak is the largest peak after the prompt peak so the approximate height is estimated by finding the bin with the most entries between 50 and 100 ns after the prompt peak. The high edge of the reflection peak is then found by stepping back from the end of the histogram looking for the point where the number of entries in the bin exceeds one third of the approximate height of the reflection peak. The starting value of the position and height of the reflection peak is then taken to be the same as that of the largest bin between the high edge of the reflection peak and 12 ns before. The reflection peak is then fitted with a Gaussian from -0.75 to +2.0 sigma, using a starting value of 2.5 ns for sigma (a typical value is actually around 2.3 ns) and iterating 5 times. To improve the fit, the reflection peak is then fitted with a Gaussian plus a linear background over the range ± 5 ns from the peak, again iterating 5 times. These fits were done using a non-linear least squares algorithm from Numerical Recipes [49]. The χ^2 for the prompt peak fit is a bit high because the statistics are good enough to show that the data is non-Gaussian. The χ^2 is very reasonable for the reflection fit because the statistics are too poor to show the model is inaccurate. The difference in the time of the peaks is easily calculated and the errors combined. The poor χ^2 of the prompt peak should not be a problem because the error in the difference is totally dominated by the error in the position of the reflection peak. The reason for including δ_t in the linear fit is to allow for any bias in the fitted position of the reflection peak, adding say 5 ns to all the time differences will be absorbed by the δ_t and should not affect the fitted δ_x, δ_y and δ_z. (A bias in the time of the reflection peak could easily arise because there is no reason to suppose that it should be a sharp peak as opposed to a cut off. The reflection peak could be a combination of reflections over the range 160–180° shifting the overall peak earlier.)

With the left hand side of equation 5.1 evaluated for each of the 92 regions the

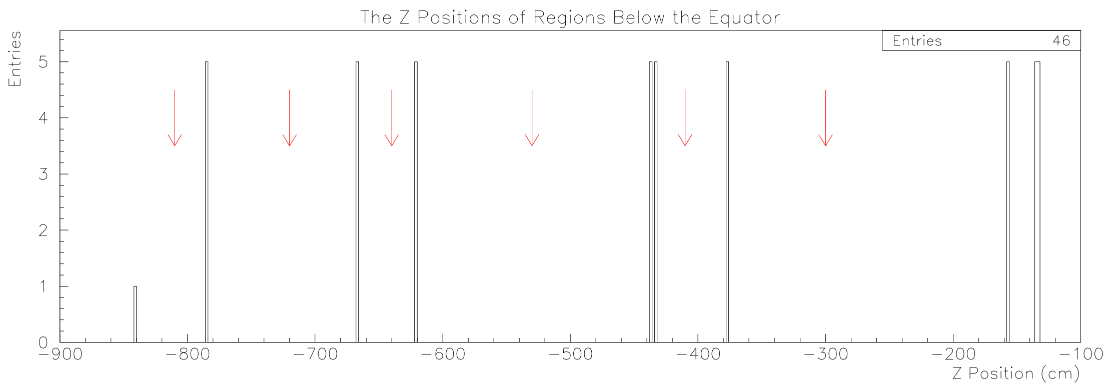

Figure 5.22: Histogram of the z position of PMT regions below the equator. The red arrows indicate the position of cuts made to the data.

linear fit can be performed. There is one further complication. The laserball does not emit light isotropically and in particular does not illuminate the PMTs near the neck of the acrylic vessel as well. Furthermore, there is a gap in the reflective sphere of the PSUP to make room for the neck of the acrylic vessel. The position of reflection peak for PMTs near the bottom of the detector might be systematically different from its position in other areas of the detector. To test how much of an effect this is it was decided to place cuts on the Z position of regions included in the fits. Figure 5.22 is a histogram of the Z position of regions below the equator. The red arrows indicate the positions of cuts made to investigate any bias.

The linear fit was carried out on PCA run 14444 using a Numerical Recipes routine for each of the cuts on region Z and with no cut. Figure 5.23 shows the results of the fits with the uncut fit being plotted with a Z cut of -1000. It can be seen than the fitted position hardly varies with the position of the cut but the χ^2 per degree of freedom improves as the cut on Z is raised. Most of the improvement in the χ^2 has been made by -650 cm so this is the value of the cut used for the rest of this section. The χ^2 with a cut of -650 cm is 114 with 81 degrees of freedom, the probability that the data fits the model is therefore 0.38%. Clearly the χ^2 is not outstanding. There are two possible reasons for this: the error in the position of the

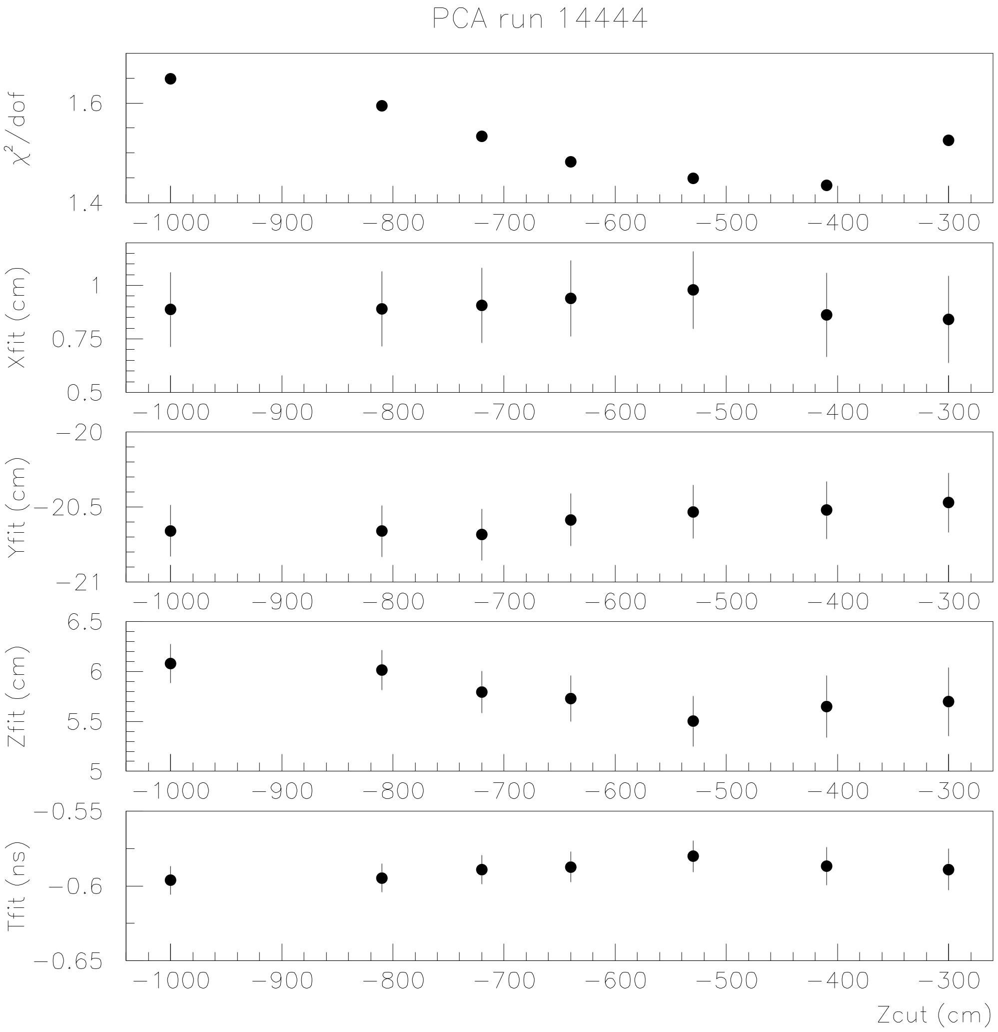

Figure 5.23: Plot showing the fitted δ_x, δ_y, δ_z and δ_t. The χ^2 per degree of freedom is shown at the top.

reflection peak could be overly optimistic or the PSUP is sufficiently non-spherical to introduce systematics differences in the reflected time. Both are plausible. The error in the reflection fits would only need to be increased by $\sqrt{1.48}$ from around 0.1 ns to make the χ^2 reasonable. The variation in the average radius of the PMTs in each region span around 1 cm corresponding to an extra 0.1 ns variation in the reflected time. The result of the fit with a Z cut of -650 cm and the position reported by the manipulator are:

Variable	Linear Fit	Manipulator
X	0.94 ± 0.18	0.0
Y	-20.59 ± 0.18	-21.6
Z	5.73 ± 0.23	-0.5
T	-0.587 ± 0.010	—

The most serious discrepancy is in the fitted Z position. Note that the error is around 0.2 cm which is well inside the diffuser ball radius of 5 cm.

It can be demonstrated that the above fit is reasonable by examining the raw data. Figure 5.24 shows the raw data used for the linear fit plotted against the region X, Y, Z and a dot product of the region position with unit vector in the direction of the fitted position. Figure 5.25 explains why the fourth plot is expected to be a straight line. The results of the linear fit can be read off from these diagrams. Referring back to figure 5.20 it can be seen that the maximum time difference minus the minimum time difference will be equal to $4\delta/c$. From the data it can be seen that the maximum minus the minimum is almost 4 ns. Therefore the laserball is approximately $(max - min)c/4 = 22$ cm off centre. Again referring back to figure 5.20 it can be seen that the time difference of the region that intersects with the positive side of an axis minus the time difference of the region that intersects with the negative side of the axis is equal to 4/c times the offset of the laserball along that axis. Therefore figure 5.24 shows that the laserball is very close to X = 0, most of the offset is in the negative Y direction and there is a small positive offset in Z. The final diagram is expected to be a straight line and shows that δ_t is around -0.6 ns. The final diagram also shows the scatter of the data about the expected form and gives some indication that the χ^2 of the fit will be a little high. These diagrams give a good feeling for how well the data distinguishes the laserball position. There are still some lingering doubts, the method is an approximation for small δ, so how far out can it be applied?

To test the fit method the laserball was simulated in 11 different positions using the Monte Carlo. The X position of the laserball was held fixed and the Y

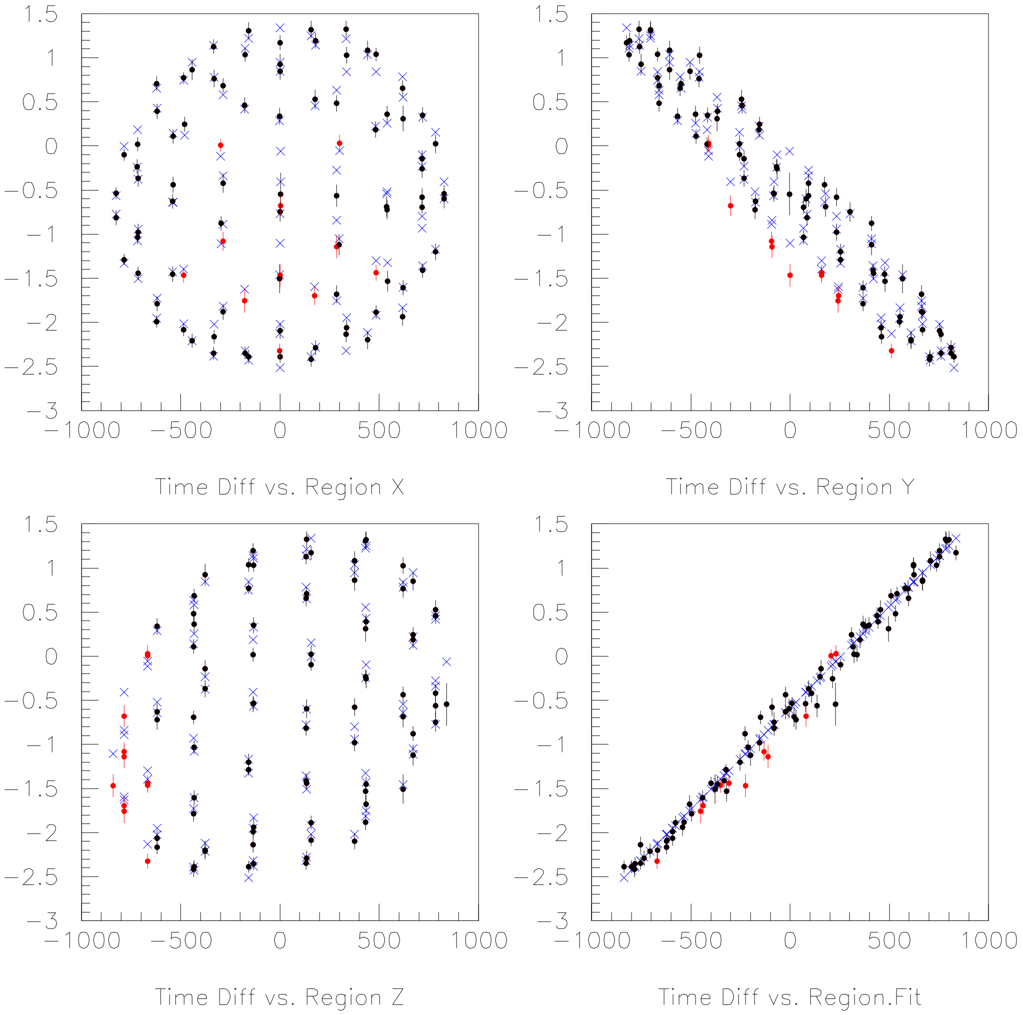

Figure 5.24: Diagram showing the time difference between the reflection peak and the prompt peak minus $2R/c$ against the X, Y and Z of the region. Also plotted is the time difference against the dot product of the region vector with a unit vector in the direction of the linear fit. The red points correspond to regions below the z cut of -650 cm. The blue crosses show the fitted time difference for every point. The last diagram is designed to be a straight line (which would go through the origin if there were no δ_t).

and Z incremented in 10 cm steps yielding the positions: 0,0,0 , 0,10,10 , 0,20,20 … 0,100,100. The laserball was simulated as a series of bursts of photons using the angular distribution of the laserball as implemented in SNOMAN. The time distribution was simulated as a Gaussian where the width was adjusted to match

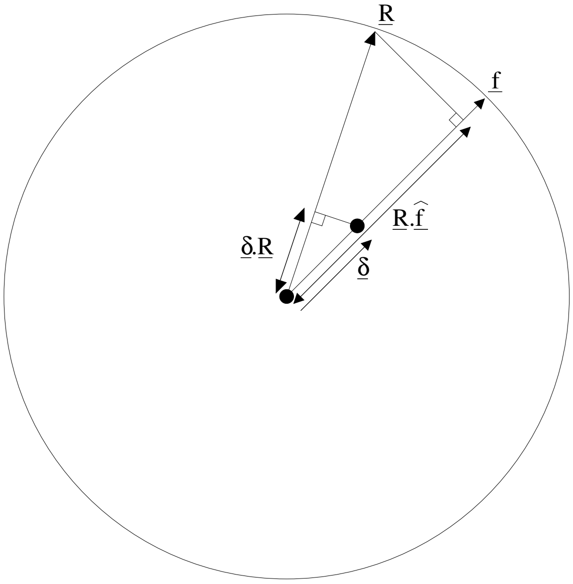

Figure 5.25: The diagram shows two similar triangles with hypotenuses of length R and δ. It follows that $\underline{\delta}.\underline{R}/\delta = \underline{R}.\underline{\hat{f}}/R$. The time difference between the reflection peak and the prompt peak is equal to $2\underline{\delta}.\underline{R}/c$ and therefore the time difference is proportional to $\underline{R}.\underline{\hat{f}}$.

that of the real laserball. Figure 5.26 shows the results of the linear fit to the position. The fit performs well for all radii with a systematic pull in of around 2 cm at a radius of 141 cm. The χ^2 per degree of freedom starts to rocket at around 60 cm. All PCAs taken so far claim to be inside 30 cm. The fit looks to be reliable to at least a centimetre within that range.

The results of the linear fit on all PCAs since the high voltage rise until mid-February 2001 are shown in figure 5.27. There is good agreement between the fit and the manipulator position in X and Y, except for a brief period in January 2000 where the discrepancy in Y grows to over 5 cm. There is clearly a systematic difference between the fit and the manipulator of around 5 cm in Z. A possible reason for this discrepancy is that the fit is done relative to the centre of the PSUP while the manipulator position is given relative to the acrylic vessel. If the acrylic vessel were sitting 5 cm high then the discrepancy would be explained. Overall it can be stated that the manipulator position is accurate to around 5 cm. Analyses of neutrino data

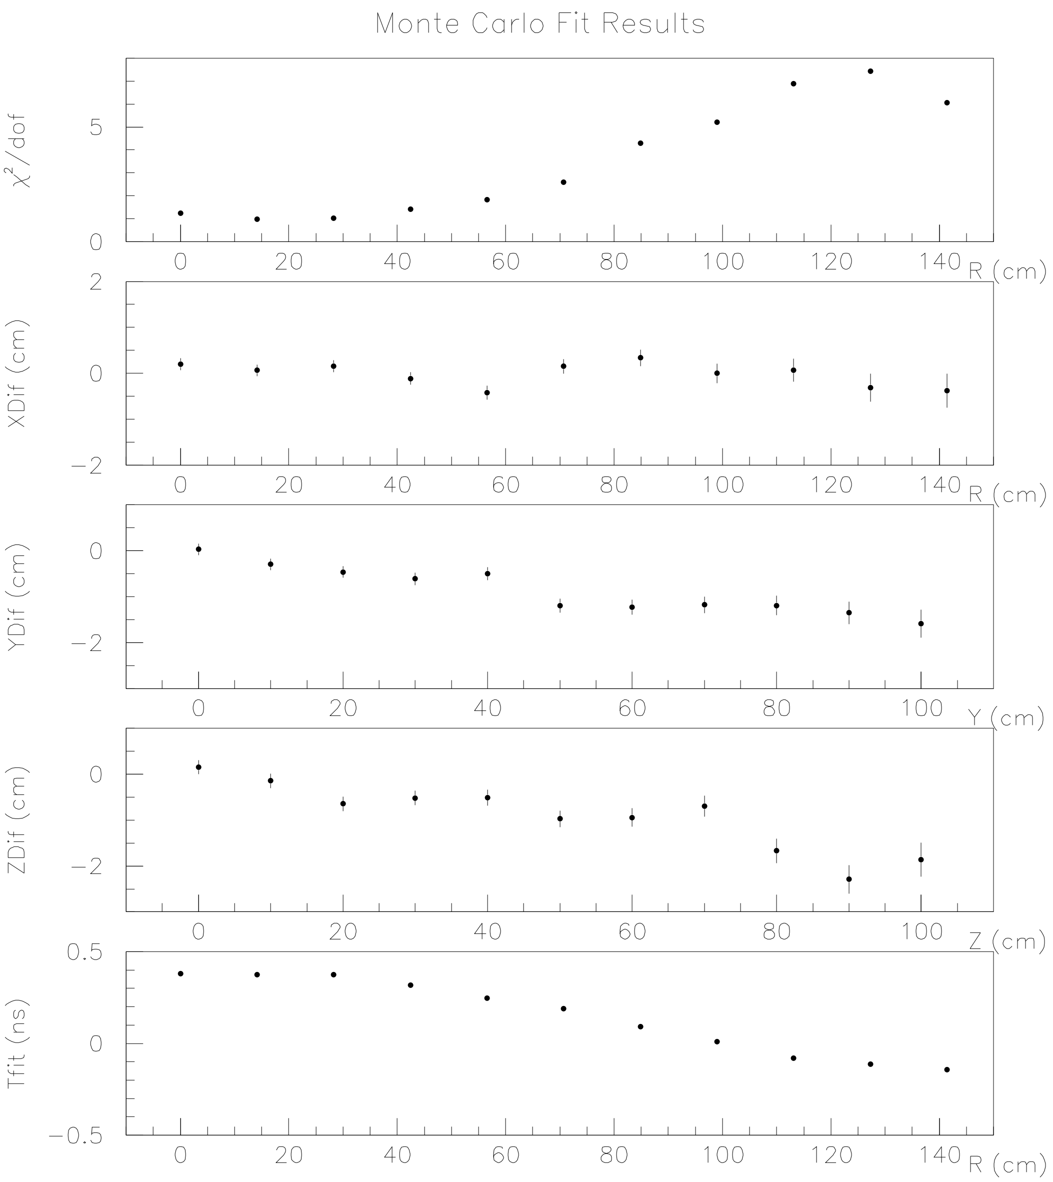

Figure 5.26: Fit results for the Monte Carlo Laserball. The laserball was simulated at 0,0,0 , 0,10,10 , 0,20,20 ... 0,100,100. The χ^2 per degree of freedom, the difference between the fitted X and simulated X and the fitted δ_t are plotted against the simulated radius. The difference between the fitted Y and Z and the simulated Y and Z are plotted against the simulated Y and Z.

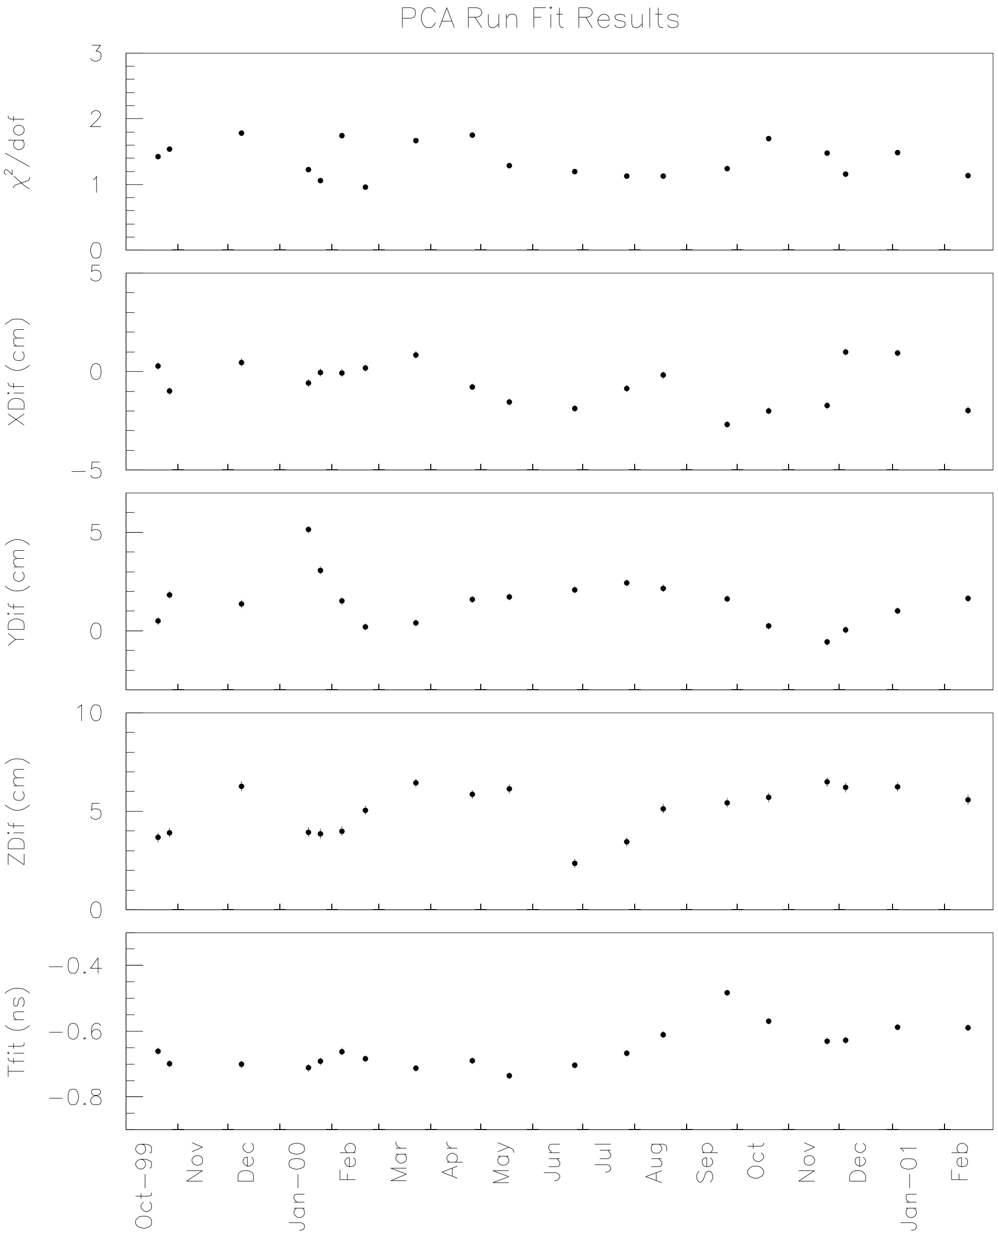

Figure 5.27: The fit results for all PCAs since the high voltage raise until mid-February 2001. The difference between the fitted X, Y and Z and the reported manipulator position is plotted against date. The χ^2 per degree of freedom and the fitted δ_t are also shown. The error bars are plotted but are very small, typically 0.18 cm for the fitted X and Y and 0.23 cm for the fitted Z. It can be seen that the laserball typically fits 5 cm higher than the reported manipulator position.

which use a radial cut of 550 cm will not be affected by this uncertainty because all it will result in is a slightly different region of D_2O being studied. That said, the cause of the discrepancy should still be tracked down and fed back into the analysis.

5.3 The Effect of Occupancy

This section investigates the effect of the intensity of the laserball on the walk and gain calibrations. PCA laserball runs are normally taken with the intensity of the laserball adjusted so that 5 to 10% of online PMTs fire in response to each flash. The fraction of PMTs which fire in an event is commonly referred to as the occupancy. As the intensity of the laserball is increased, so the fraction of events where the PMT receives more than one photon will increase. The discriminator fires in response to the first pulse which crosses its threshold so it might be expected that the measured time would shift earlier with increasing occupancy. The magnitude of this effect is quantified in section 5.3.1.

Clearly multiple photons will increase the amount of charge received by a PMT. Therefore, it might be expected that the charge spectrum will become broader with increasing occupancy, affecting the high half point. The effect of occupancy on the gain calibration is studied in section 5.3.2.

The effect of occupancy on the timing and the gain calibration was studied using a set of laserball runs. The run numbers and their occupancies are:

Run No.	Occupancy
11328	0.005
11330	0.013
11332	0.051
11326	0.103
11334	0.204
11344	0.304
11336	0.342
11338	0.736
11340	0.932
11342	0.987

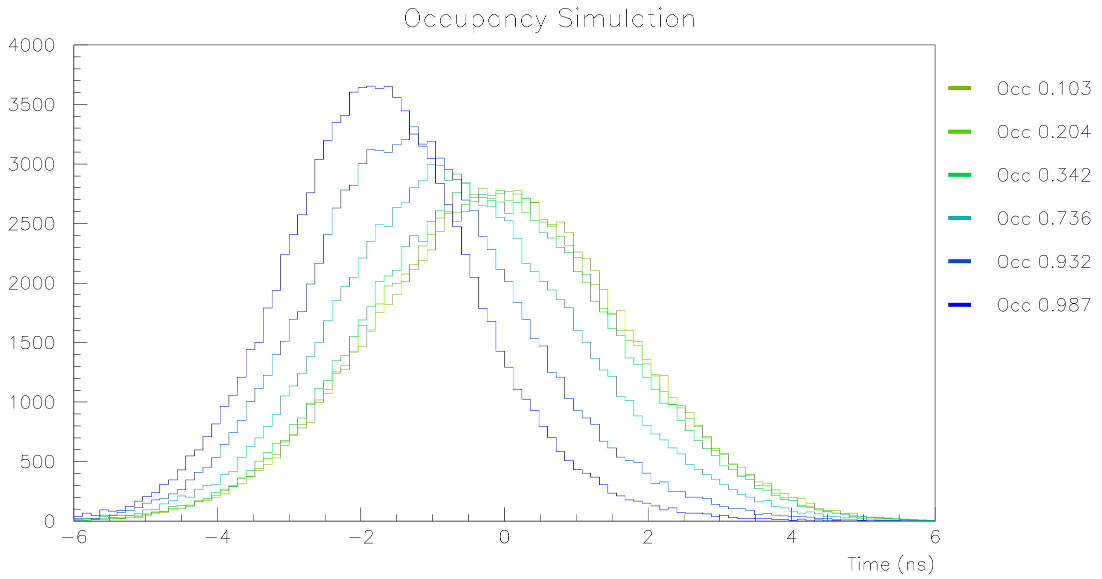

Figure 5.28: The results of a simple Monte Carlo drawing points from a Gaussian of width 1.75 ns and keeping the earliest one for six different occupancies.

5.3.1 The Effect of Occupancy on the Timing

To get a rough idea of how much earlier the time might move at a given occupancy, a very simple Monte Carlo was done. 100,000 discriminator firings were simulated for a range of occupancies. For each event, a number of photons were chosen from a Poisson distribution with mean $-\log(1-\text{occupancy})$. The chosen number of points were then drawn from a Gaussian of width 1.75 ns and the earliest one kept as an estimate of the time the discriminator would have fired. This method is very crude. Firstly the number of photons will have been overestimated because the occupancy includes all PMTs that fired and not just the fraction that fired in response to prompt light. Secondly in a real PMT the pulses will merge somewhat on their way down the dynode stack. Figure 5.28 shows the result of the Monte Carlo for six different occupancies. It can be seen that as the mean number of photons striking the PMT increases, the time at which the discriminator fires moves earlier and the distribution narrows. Figure 5.29 shows the actual timing distributions at several

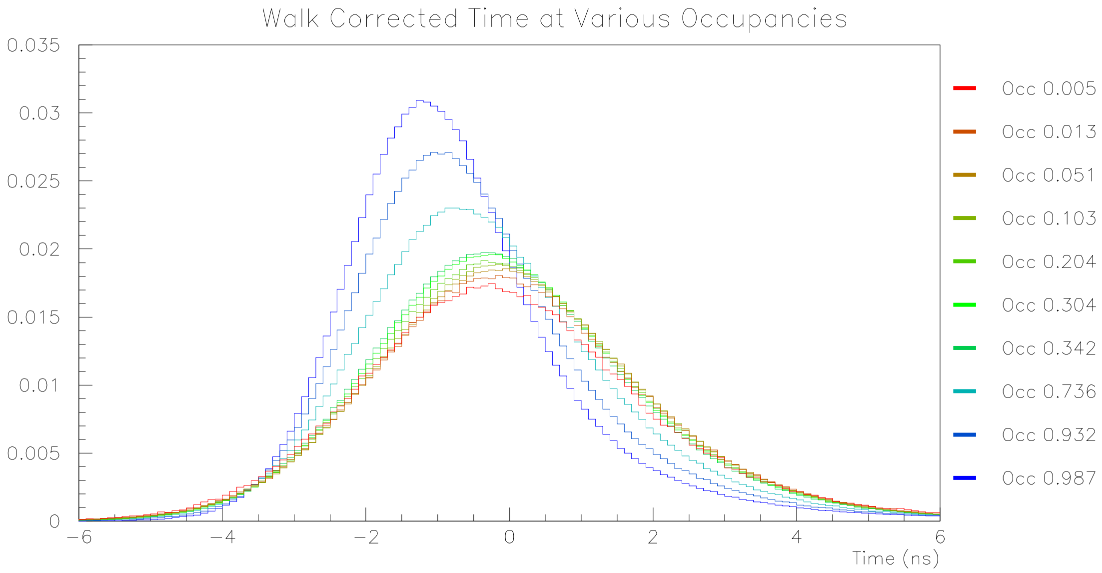

Figure 5.29: The actual timing distributions at several different occupancies.

different occupancies. Superficially the two sets of distributions are quite similar. The actual timing distributions below 34% occupancy have indistinguishable means and widths. The normalised height of the prompt peak below 34% drops slightly with decreasing occupancy, this is an artifact from the amount of late light included in the histogram. The histograms were collected over a range from -10 to +10 ns. As the occupancy increases so the fraction of tubes which can fire in response to late light decreases, resulting in a larger normalised prompt peak. The most obvious difference between the simple Monte Carlo and the data is that at high occupancy the data has a tail to late times. This tail will be discussed further after examining the timing as a function of charge.

To investigate the timing as a function of charge, profile histograms of the cable delay corrected time and full PCA corrected time were generated. To reduce the chance of including noise in the histograms a data point was only added if the charge was greater than the DAQ threshold and the DAQ threshold agreed with the gain threshold to at least 3 ADC counts. The PCA time was path corrected before entering the time into the histogram. To get the cable delay corrected time, the optical path and the value of the walk at high charge were subtracted. To get the time of the walk at high charge, 100 ns (the time of the prompt peak) was uncalibrated using a QHS of 3 in PCA units (i.e. at high charge). The range of times accepted into the PCA time histogram was -10 to +10 ns while the range for the cable delay corrected times was -10 to +15 ns. Figure 5.30 shows the cable delay corrected time, the PCA corrected time and the RMS of the data about the PCA correction as a function of QHS with close ups at low charge. As the occupancy is raised the average time at which the discriminator fires gets later at low charge and earlier at high charge. The width of the distribution gets broader at low charge and narrower at high charge. It is the late times at low charge that are responsible for the tail noticed when comparing the simple Monte Carlo to data. The fact that, at low charge, the PMTs seem to fire later with increasing occupancy is, at first sight, very peculiar. A possible mechanism is through electronic pickup from signals on adjacent channels, this is studied below. The most important thing from the point of view of the timing of neutrino data is that the overall timing distributions are very similar below 34%.

A feature of pickup is that it tends to come late in time. The charge cut at the DAQ threshold which was used above was designed to eliminate pickup, but it may not have done so completely. Suppose that a channel is hit but does not produce enough charge to fire the discriminator, a small amount of pickup from an adjacent channel could be enough to drive the discriminator over the edge while still integrating a respectable amount of charge. If this model is accurate then it would be expected that the charge of a channel will be slightly less on average if the charge on adjacent channels were large. It would also be expected that the average time at which a channel fires will be later on average if the charge on adjacent channels were large.

Figure 5.31 shows charge and time profile histograms versus the charge of

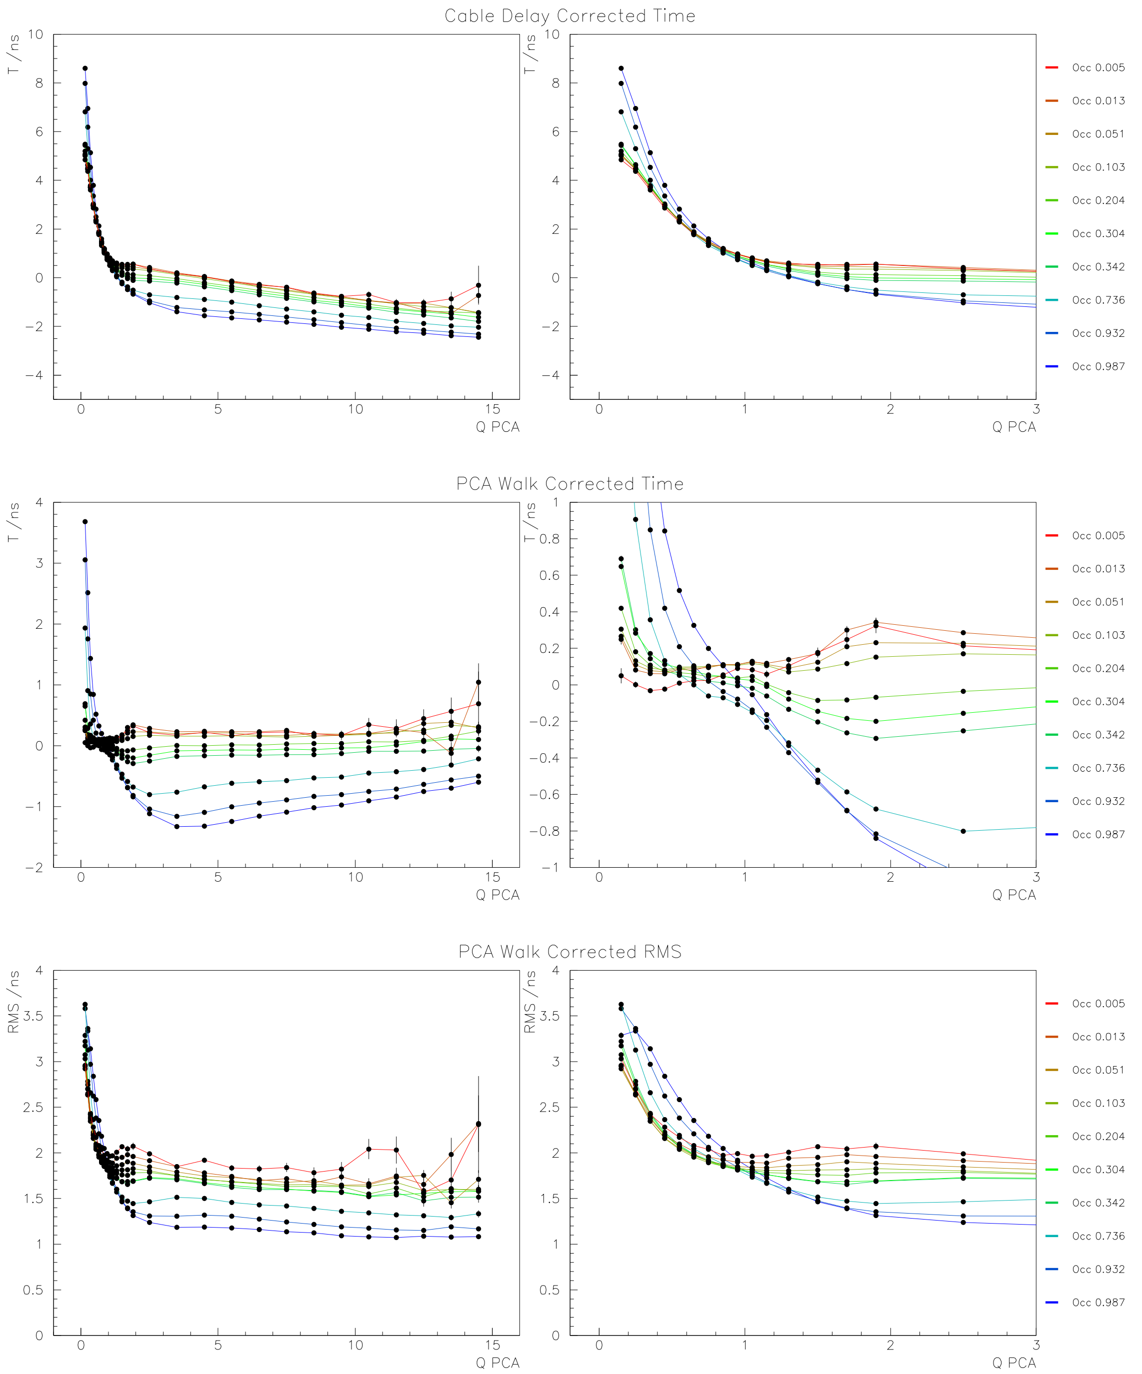

Figure 5.30: The average cable delay corrected time, full PCA corrected time and RMS as a function of charge at several different occupancies. All data between -10 and +15 ns was included in the cable delay and all data between -10 and +10 ns was included in the average and RMS.

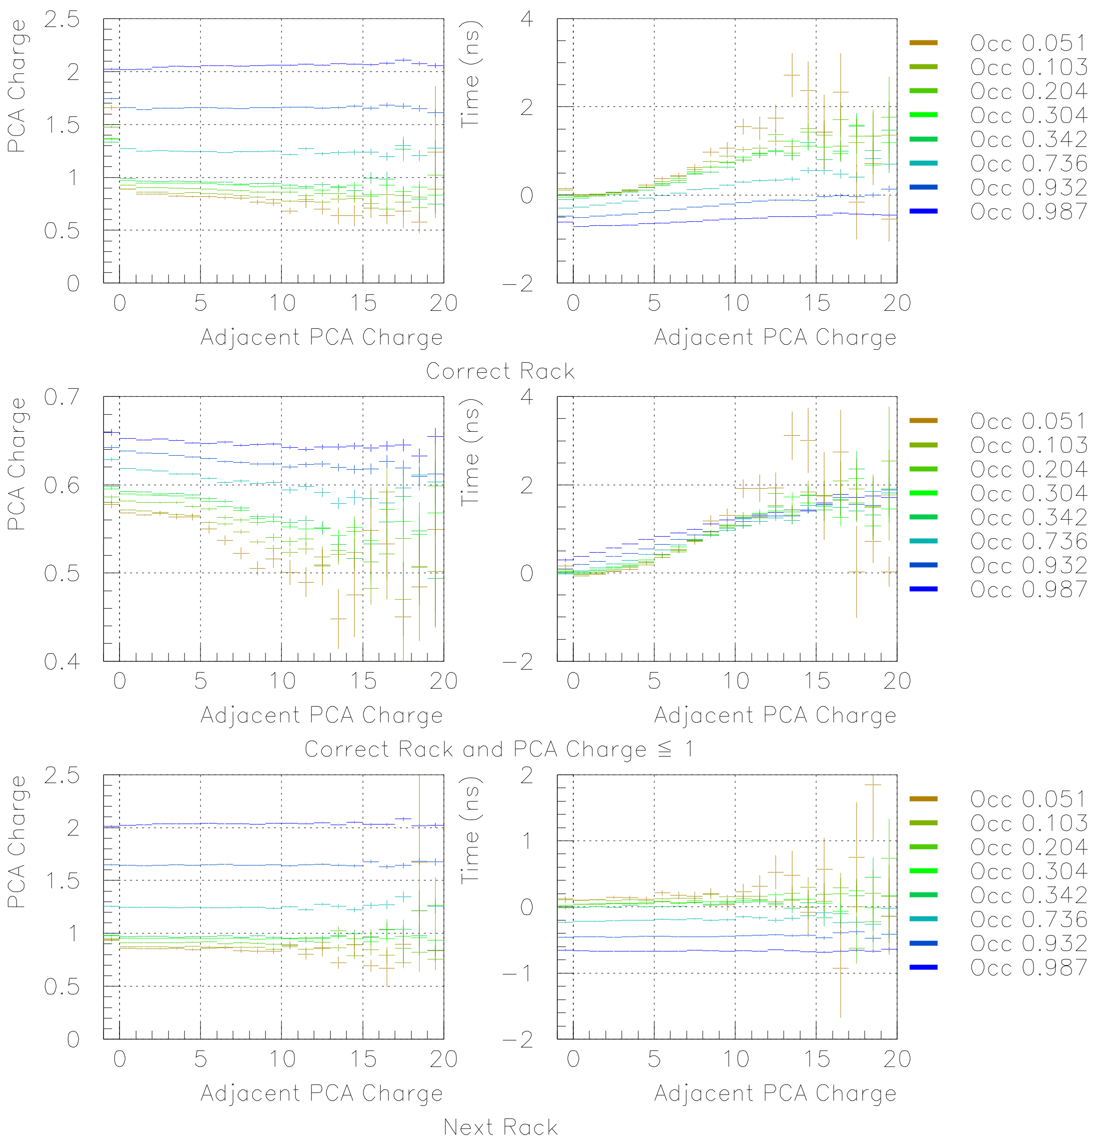

Figure 5.31: The left hand plots show profile histograms of the charge of a channel versus the charge of any neighbouring channels which also fire. The right hand plots show profile histograms of the time at which a channels fires versus the charge of any neighbouring channels which also fire. The top two plots have no cuts applied. Channels were only entered into the middle pair if their PCA charge was less than 1. The bottom pair have no cuts applied but the adjacent channel is redefined to be the adjacent channel in a different electronics rack. See the text for more details.

adjacent channels. There are three pairs of plots. The top pair show the charge and time of a channel against the charge of any adjacent channels on the same card. If a channel has no adjacent channels which were hit then it is not entered into the histogram. If a channel has two adjacent channels (the maximum possible) then there is one entry for each of the neighbours. The y limits of the profile histograms are from -1 to +5 PCA units for the charge and -10 to +10 ns for the time. The middle pair of histograms shows the same thing except that channels were only entered if their PCA charge was less than 1 PCA unit. The motivation for this was that the channels with the lowest charge are more likely to be caused by pickup. The bottom pair of plots are the same as the top pair of plots, but instead of looking at adjacent channels on the same card, looking at 'adjacent' channels on the same card in a different rack (cycling the crate number by 2). In the model where late times are caused by pickup there should be no correlation in this case.

Examining the three pairs of plots it can be seen that there is a very weak correlation between the charge of a channel and the charge of adjacent channels in the direction prescribed by the pickup model. There is a much stronger correlation visible in the time, again in the direction prescribed by the pickup model. The effects are clearest in the lowest occupancy runs. At high light levels, there will be a high incidence of adjacent channels both firing by chance, even without the effect of pickup. This means that the fraction of pairs which have fired due to pickup could drop with increasing occupancy even if the fraction of hit channels which are caused by pickup increases. The two lowest occupancy runs are not plotted because they did not contain many adjacent channels and so the error bars obliterate the rest of the plot. Overall it can be stated that pickup is the most likely cause of the late tail seen in the timing distributions.

5.3.2 The Effect of Occupancy on the Gain

The laserball occupancy scan can also be used to investigate the effect of occupancy on the gain calibration. Histograms of the QHS and QHL distributions were collected in units of counts above pedestal for each of the laserball runs. To avoid including channels with bad pedestals in the histograms it was demanded that the DAQ threshold of a channel agree with its gain threshold to at least 3 ADC counts. No cut was made at the DAQ threshold. Figure 5.32 shows two examples of QHS histograms along with the fitted threshold, peak and high half point. The threshold has been incorrectly fit below the noise peak in the higher occupancy histogram. Figure 5.33 shows the results of fits to all the collected histograms as a function of occupancy. The gradient of the QHS high half point versus occupancy, for occupancies less than 35%, is 14.6±3 ADC counts, as fitted in PAW. PCAs are taken with an occupancy of between 5 and 10%. Using the above gradient, a change of 5% in the occupancy will alter the measured hhp by ∼0.7 ADC counts.

5.3.3 Rates and Loading

It was shown in section 5.1 that the position of the charge pedestal depends on the rate at which the detector is read out, this will be discussed further in the next chapter. The state of the detector during the laserball occupancy scan therefore needs to be considered. The left hand plot of figure 5.34 shows the average rate at which the PMTs were read out as a function of occupancy. This plot is not linear because the laserball rate was reduced as the occupancy was increased. If the PMTs are read out at more than about 4 Hz the circular buffer in the builder starts to fill. The loading on the crate and cards (the number which are read out in one go) is proportional to the occupancy.

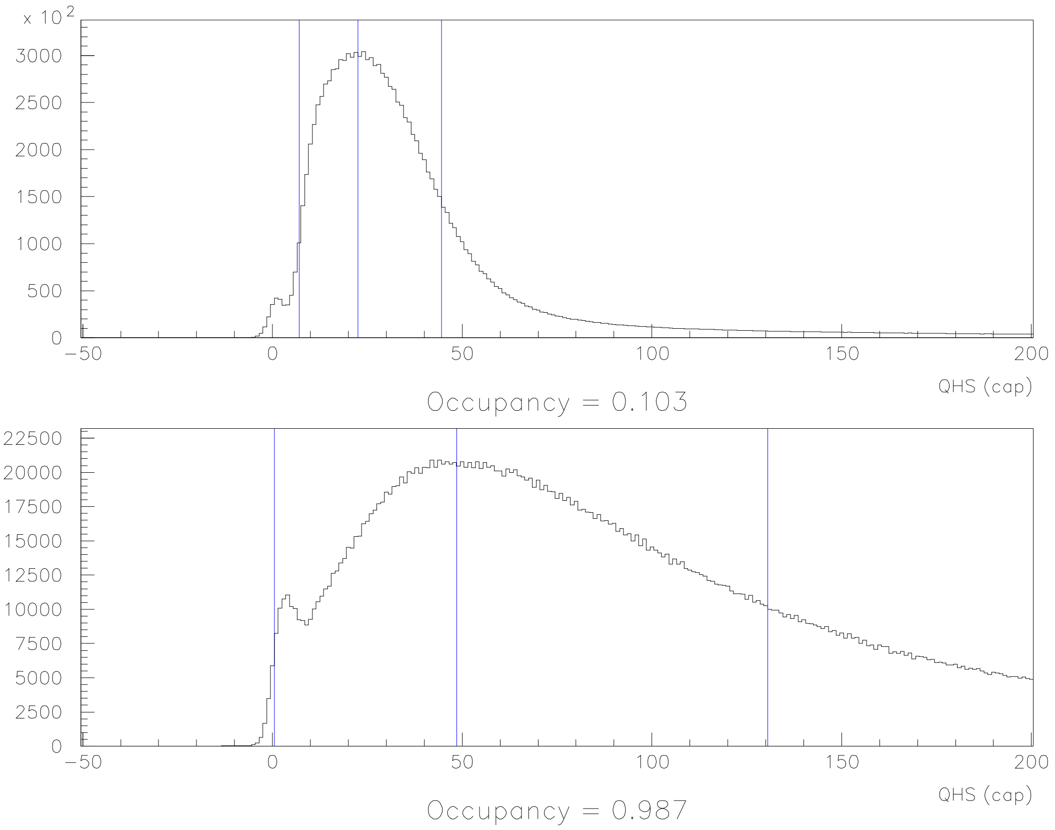

Figure 5.32: QHS histograms for two different occupancies. The fitted threshold, peak and high half point are shown.

5.4 The Stability of the Constants

A PCA calibration is taken roughly once every month. This section takes a brief look at how much the constants change between PCAs. To see how much the high half point (hhp) varies over time all hhps for each PCA were histogrammed and fitted with a Gaussian from -2 to +2 sigma. Figure 5.35 shows example fits for QHS and QHL and the fitted peak position as a function of time. The typical QHS hhp has been stable over time to within an ADC count.

To take a more detailed look at the changes in individual channels the differences between six PCAs up to run 14444 (taken in mid-February 2001) were

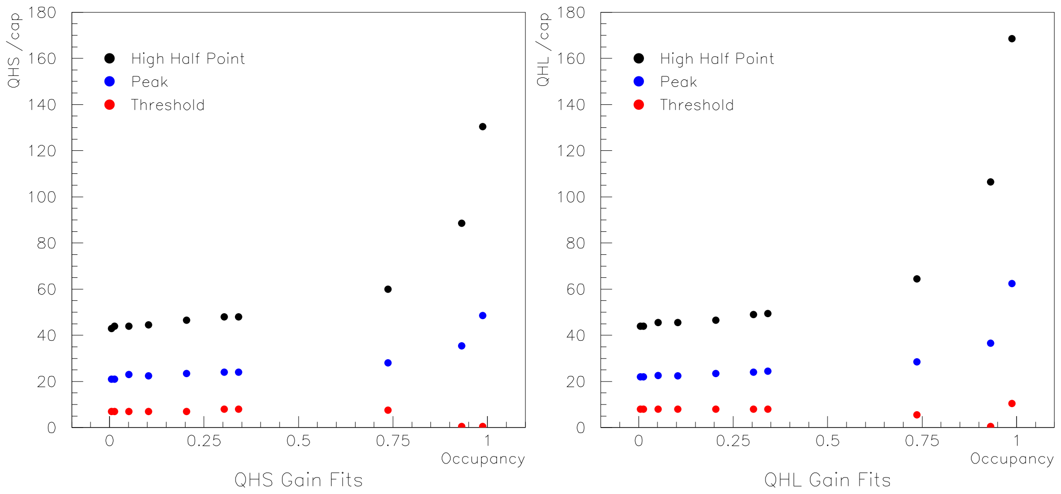

Figure 5.33: The fitted threshold, peak and high half point as a function of occupancy for QHS and QHL.

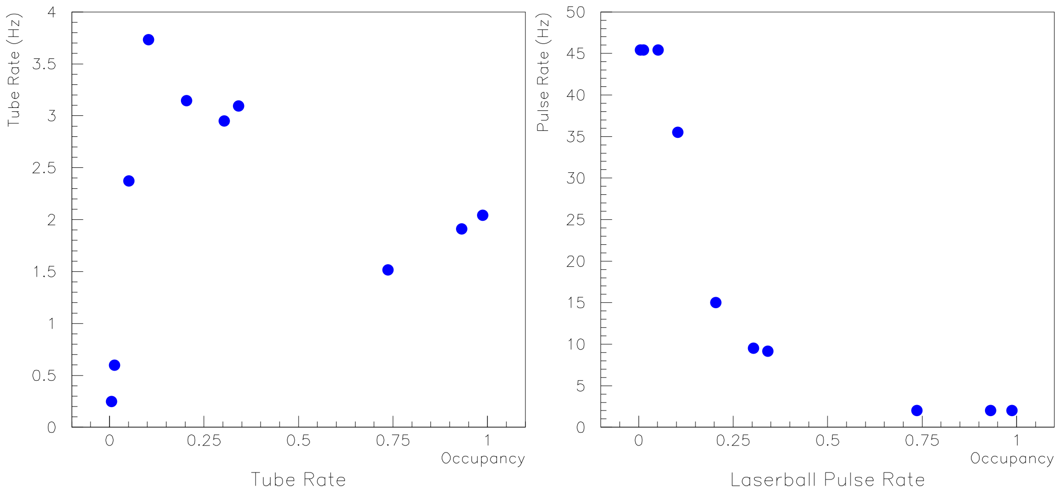

Figure 5.34: The plot on the left is the average PMT read out rate over the run at different occupancies. The reason that it is not linear is that the laserball rate was turned down for the higher occupancies, see the plot on the right.

examined. The differences are only plotted if the electronics board IDs are the same for both runs. Figures 5.36 and 5.37 show the change in the QHS and QHL hhp for the six runs against logical channel number cutting at ± 20 counts. The ECA validation word was checked for all channels in both runs (this was done by calling the ECA calibration routine for every channel with charges of 1000 ADC counts and a TAC of 1800 ADC counts and recording the ECA validation word in an ntuple). If the channel failed the ECA validation in either run then it is plotted in red (the test was an OR of bits 8, 10, 11, 12 and 14, see [43][35]). Figures 5.38 and 5.39 shows histograms of the change, only cutting channels whose board IDs changed. Figure 5.40 shows the change in the cable delay for the six PCAs. The cable delay time was defined as the PCA correction at 100 counts above pedestal. The channels which failed the ECA validation in either run are plotted in red.

It is obvious from the plots that the majority of channels remain stable between PCAs, but how are the outliers to be interpreted? A change in the PCA cable delay could come from two sources: the physical cable delay could change or the ECA time returned for a particular delay relative to the trigger could change. Similarly with the hhp changes could occur because the gain changed or because the difference between the measured and true pedestal changed. Even if the cable delays and gains were static, the PCA would still need to be taken regularly to keep up with the current state of the detector/electronics. Once corrections are available for all the various electronics idiosyncrasies the PCA constants are likely to be far more stable. In spite of this the PCA would benefit by testing how much the constants change between PCAs and making a cut available to the analyser, much in the same way as now happens with the ECA. This is unlikely to be the most sensitive discriminant of the worst performing channels. The PCA is a ‘high’ rate calibration and so is the ECA (the problem of rate is discussed further in the next chapter). A cut which is more sensitive to rate or loading effects is likely to be more powerful (the data for these cuts could be stored in separate bank, which in SNO terminology would be called an ANXX bank). That said, the availability of a PCA cut cannot hurt (except in terms of development time) and should be considered as an extension to the current PCA.

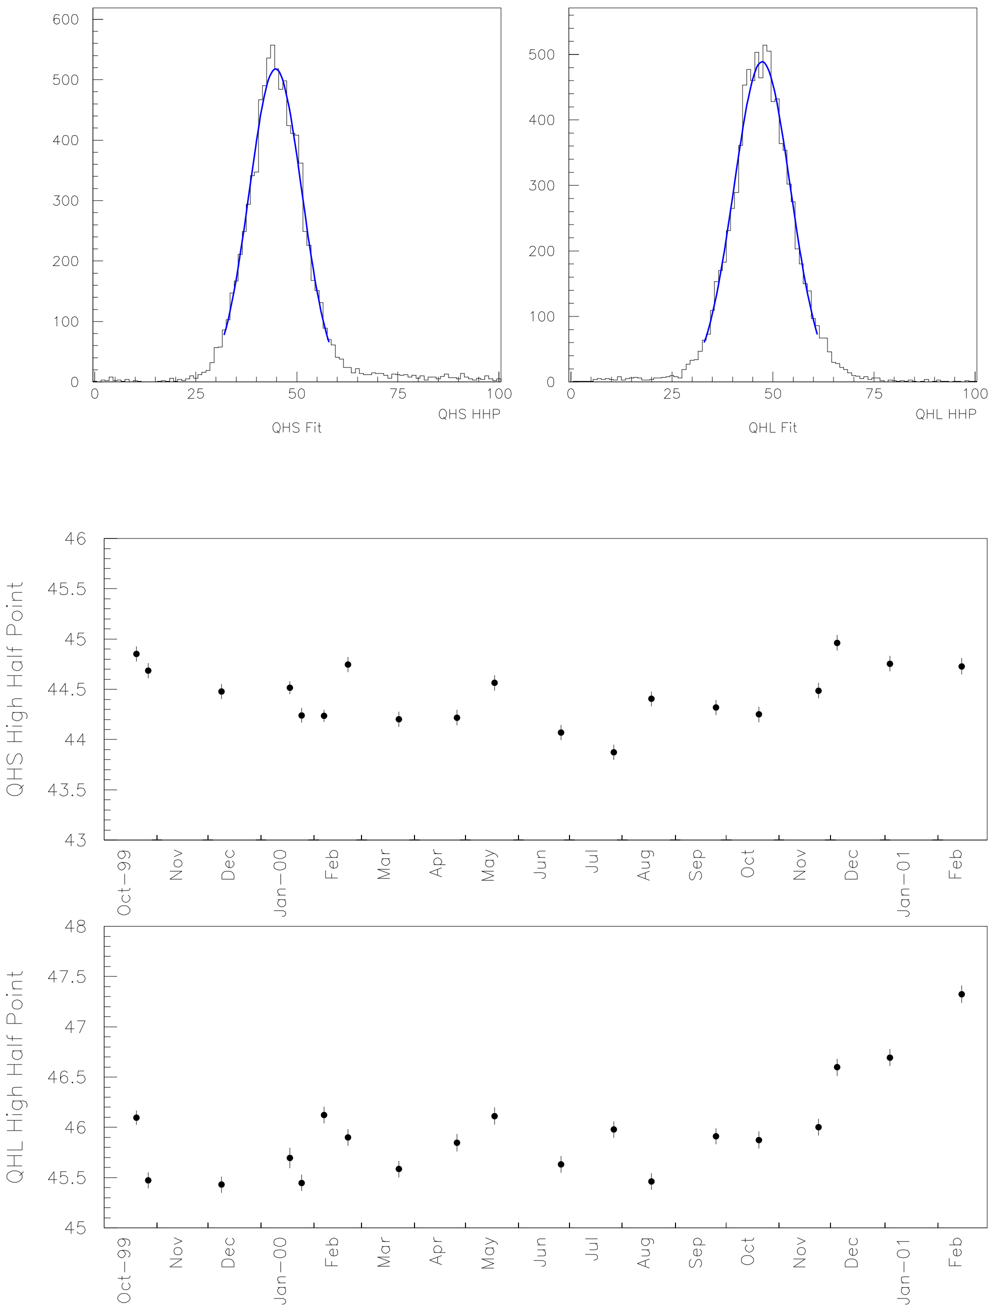

Figure 5.35: The upper plot shows the QHS and QHL high half points for a particular PCA. The peak was fitted with a Gaussian to find a typical value. The lower plot shows the result of the fit to all PCAs from the high voltage rise until mid-February 2001.

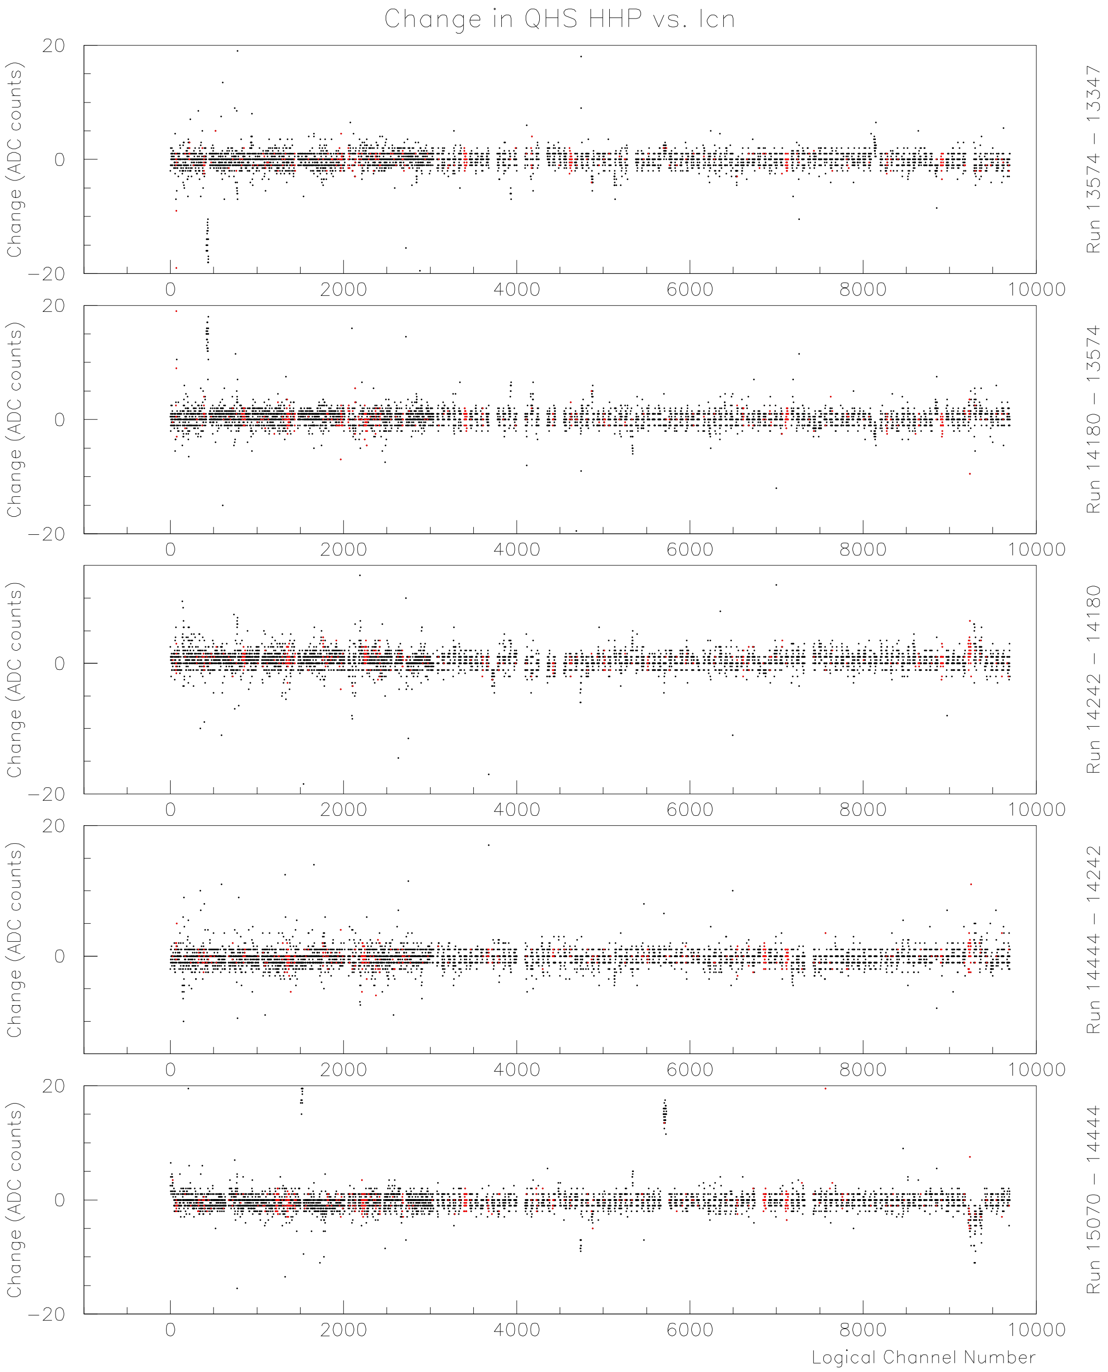

Figure 5.36: The change in the QHS high half point against logical channel number for six PCAs cutting at ±20 counts.

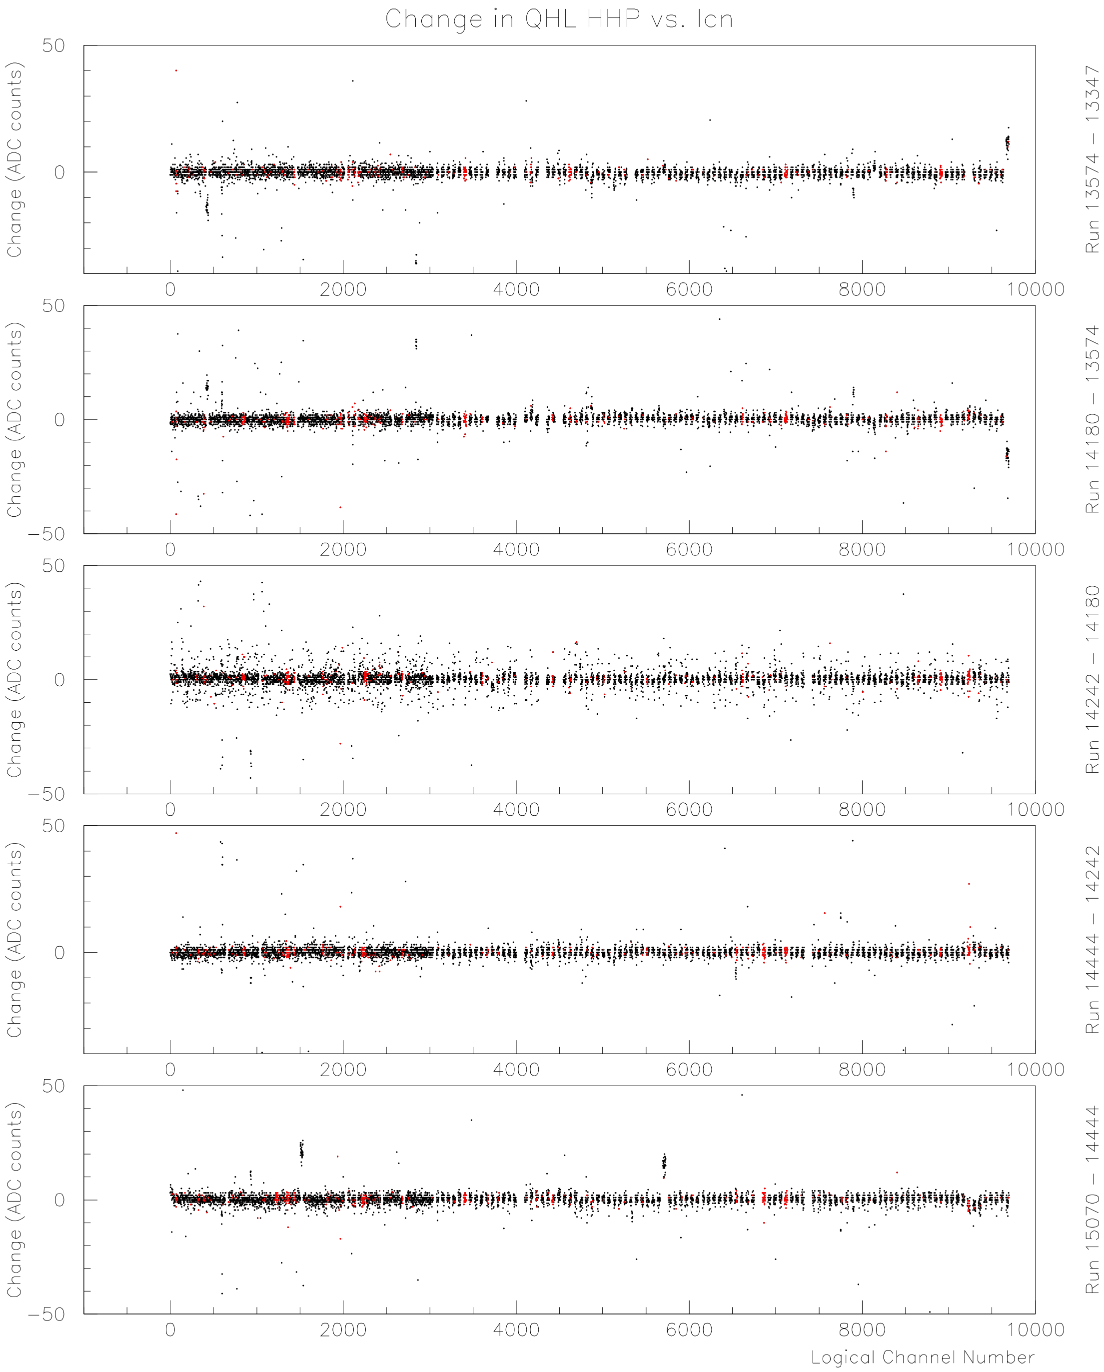

Figure 5.37: The change in the QHL high half point against logical channel number for six PCAs cutting at ±20 counts.

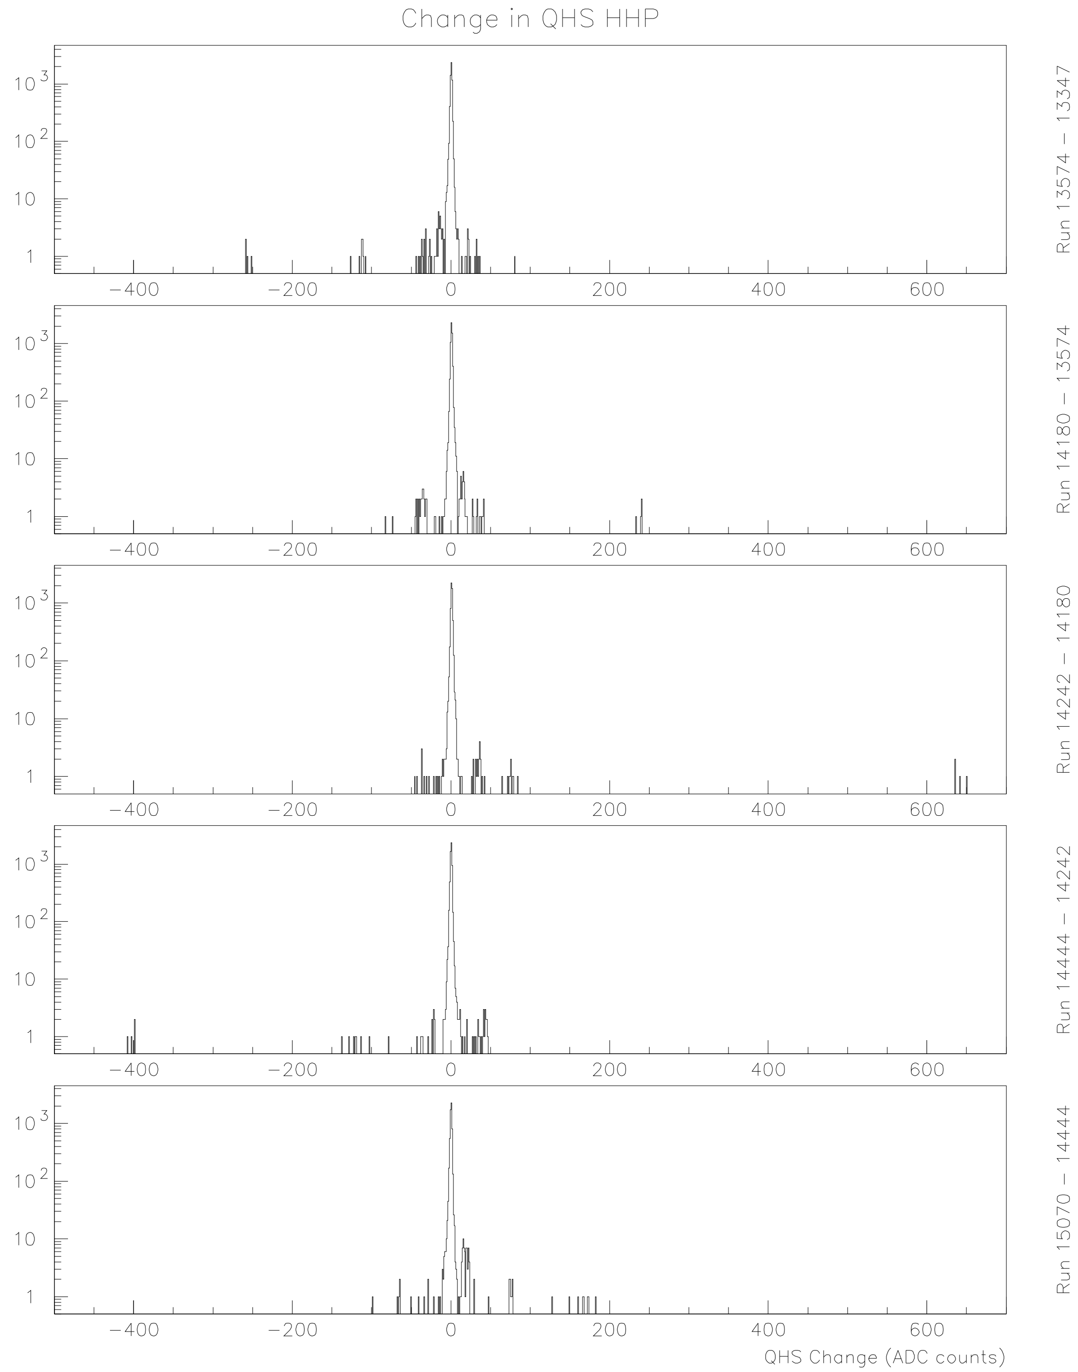

Figure 5.38: Histogram of the change in QHS for six different PCAs only cutting channels whose board IDs have changed.

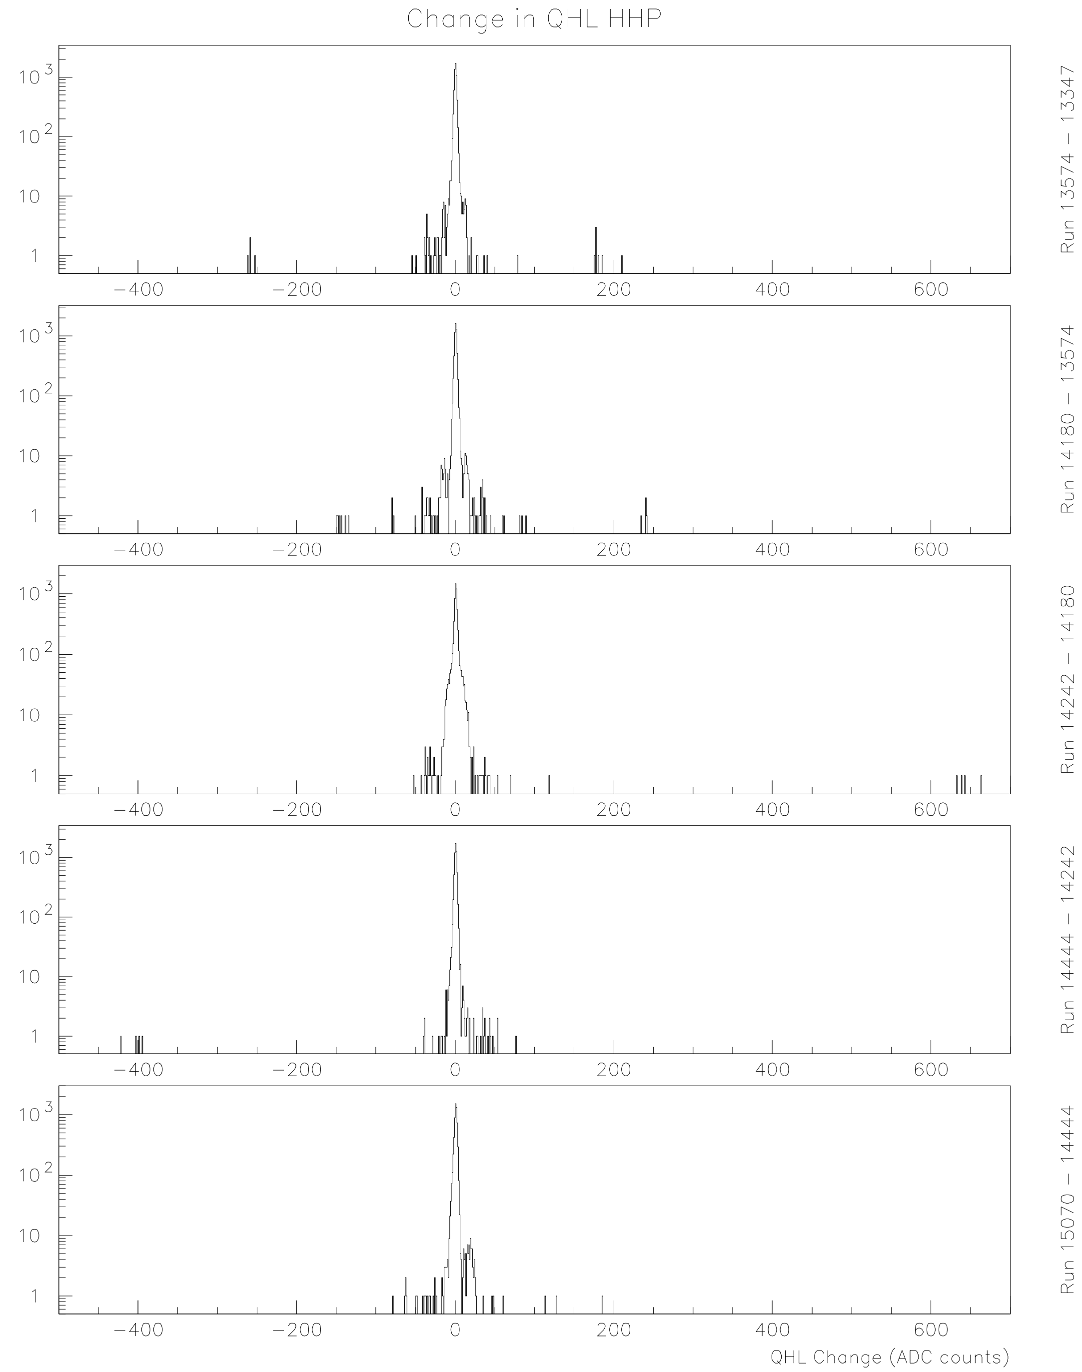

Figure 5.39: Histogram of the change in QHL for six different PCAs only cutting channels whose board IDs have changed.

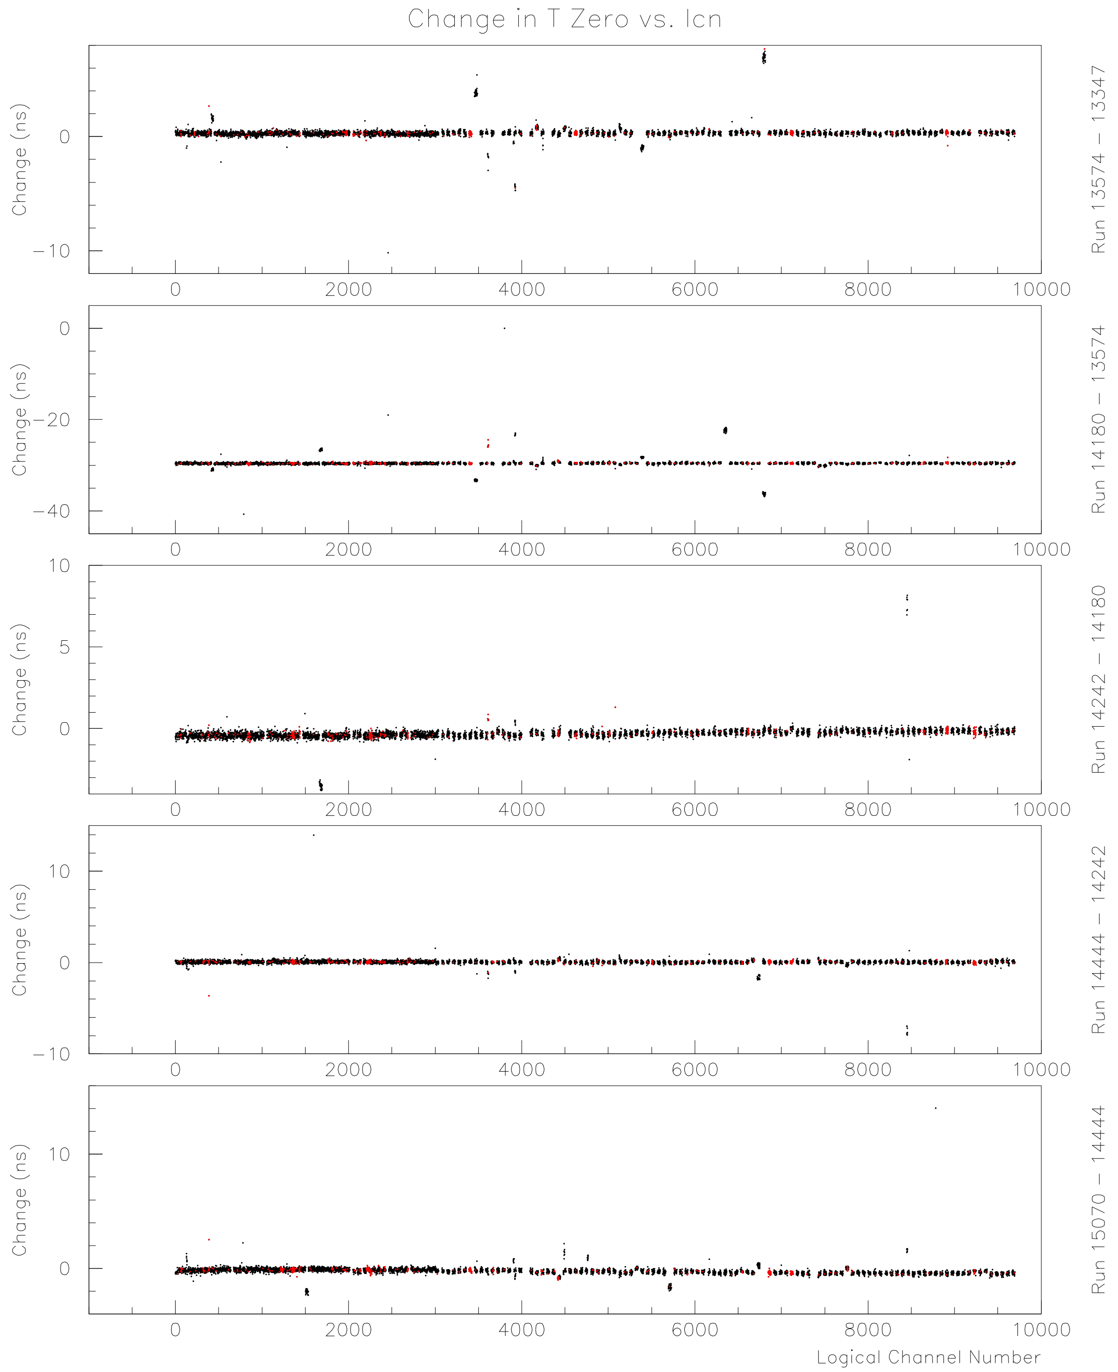

Figure 5.40: The difference in the cable delay time as a function of logical channel number for six different PCAs. The cable delay time was defined as the PCA correction at 100 counts above pedestal.

Chapter 6

The High χ^2 Problem

It was, very reasonably, anticipated that the ECA and the PCA would be the only two calibrations needed to analyse SNO data. Unfortunately it was discovered that the timing calibration did not work as well for low rate data as for high rate data. This observation came to be known as the high χ^2 problem. The next section provides an overview of the study of the high χ^2 problem with links to the contributions made by the author which will be described in the subsequent sections. The final section describes the current status of the high χ^2 problem and its correction.

6.1 Overview

One of the standard event fitters in SNOMAN is the time fitter. It is a simple, fast fitter which performs a least squares fit to the position and time of the event using the calibrated PMT times. Once an initial fit is made, the residual path corrected time is calculated for each PMT. The largest outlier is repeatedly thrown away and the event re-fit until all path corrected times are within -8 to +4 ns of the fitted event time. If too many PMTs are thrown away the fit is failed. There are 4 fitted parameters so the χ^2 per degree of freedom of the final fit is given by:

$$\chi^2/\text{dof} = \sum_{i=1}^{N_{\text{PMT}}} \left[\frac{|\underline{r}_i - \underline{r}_{fit}|/c - (t_i - t_{fit})}{\sigma_i} \right] / (N_{\text{PMT}} - 4)$$

where σ_i is taken to be 1.61 (this number was derived from the testing of the PMTs). For a good fit the χ^2/dof should be close to unity (Despite having thrown away the

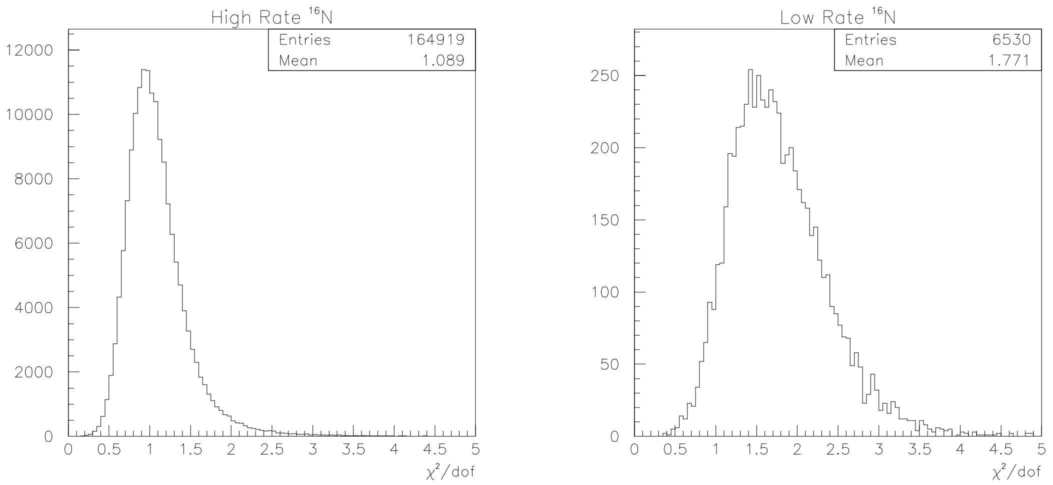

Figure 6.1: The χ^2/dof for high rate ^{16}N (run 12144, event rate $\approx$ 150 Hz) and low rate ^{16}N (run 12147, event rate $\approx$ 3 Hz). This is the high χ^2 problem.

largest outliers a good χ^2 is not guaranteed because the allowed time window is much larger than the 1.5 ns timing resolution of the PMTs). The χ^2/dof is indeed close to unity for the high rate Čerenkov data from the ^{16}N calibration source but when Xin Chen studied background events in neutrino runs he discovered that the χ^2/dof was suspiciously high. The ^{16}N is normally run at high rate in order to collect good statistics while neutrino data is necessarily taken at low rate. To see if the high χ^2 could be attributed to the detector rate, some low rate ^{16}N data was collected. Figure 6.1 shows the χ^2/dof for high and low rate ^{16}N. The difference is obvious and potentially poses a serious problem for neutrino analysis. The source is identical for low and high rate ^{16}N so there must be some detector-based effect which corrupts the timing of low rate data.

After a good deal of investigation by a number of people, a correlation was found between the PMT time and the time elapsed since the PMT was last read out. This time shift occurs over a period of tens of seconds. The PMTs fire at a rate of hundreds of Hertz but are read out at a much lower rate, depending on the source and trigger settings. Although the effect should strictly be called the time since last read out effect, it is normally referred to as the time since last hit (TSLH) effect. At first it was hoped that the TSLH correction might be similar for all channels so a global TSLH correction was attempted, see section 6.2.

Experiments where pedestals were fired at various rates soon showed that the channels have a longer memory than just the time of the previous hit. There is also some sensitivity to the overall rate at which the channel is being fired. A useful model is that each hit introduces an offset which then decays away with time. If the overall rate is high compared to one over the typical TSLH decay time then a pure rate correction might not perform too badly, section 6.3 describes an attempt at such a correction.

While it was clearly important to find a correction for data which was already taken it was also obviously desirable to eliminate the problem for new data. Two possible solutions to the high χ^2 problem were proposed: interleaving pedestals in the data or turning off the TAC twiddle bits. Section 6.4 examines the degree to which these two proposals reduce the high χ^2 problem.

The steps taken to reduce the high χ^2 problem and the current status of the correction in SNOMAN are given in section 6.5.

6.2 Global Time Since Last Hit Correction

If the size of the time since last hit (TSLH) effect were similar for all channels it might be hoped that a global TSLH correction could be made. This section describes an attempt to do just that. The TSLH data was collected from a high and low rate laserball run. The approximate number of PMTs which fire in an event (the Nhit) and pulse rates of these runs were:

Run No.	Nhit	Pulse Rate
12119	$\sim$600	38 Hz
12123	$\sim$45	4 Hz

The most natural place to make a TSLH correction in SNOMAN is before the ECA. The ECA removes minor wobbles in the TAC slope so ideally the number of

ADC counts should be corrected so that the TAC shape at low rate is converted to look like the TAC shape at high rate before applying the ECA correction. This correction converts low rate data to look like high rate data. The correction could equally well be done the other way around, although it would require re-running the PCA and other high rate calibrations using the modified calibrator.

Figure 6.2 shows profile histograms of ADC time and pedestal corrected charge against TSLH, collected from the low rate laserball run, with fitted functions shown in red. The TAC histogram accepted all times between 1600 and 2000 ADC counts, the QHS and QHL histograms accepted charges between -20 and 70 counts above pedestal while the QLX histogram accepted charges between -15 and 20 counts above pedestal. The blue points on the charge histograms show the intercept of the high rate laserball run with zero TSLH. If TSLH were the only effect then the low rate data should intersect with the high rate data at zero TSLH, which they do not. The high rate intercept for the TAC was 1940 ADC counts, perilously close to the cut off at 2000 ADC counts, and markedly different from the low rate intercept. The functional forms used to fit the low rate data are given in table 6.1 along with the fitted parameters.

Fit	Function
TAC	$\mathrm{P1} + \mathrm{P2}\exp(\mathrm{P3}\,x)$
QHS	$\mathrm{P1} + \mathrm{P2}\,x + \mathrm{P3}\exp(-0.5(x/\mathrm{P4})^2)$
QHL	$\mathrm{P1} + \mathrm{P2}\,x + \mathrm{P3}\exp(-0.5(x/\mathrm{P4})^2)$
QLX	$\mathrm{P1} + \mathrm{P2}\exp(-0.5((x-\mathrm{P3})/\mathrm{P4})^2)$

Fit	P1	P2	P3	P4
TAC	1802.0 ± 0.3	-8.7 ± 0.3	-0.018 ± 0.001	–
QHS	24.34 ± 0.05	-0.0041 ± 0.0004	3.12 ± 0.05	19.0 ± 0.3
QHL	26.8 ± 0.1	-0.002 ± 0.001	1.70 ± 0.09	19.6 ± 0.7
QLX	1.28 ± 0.01	2.56 ± 0.01	12.74 ± 0.08	18.8 ± 0.2

Table 6.1: Results

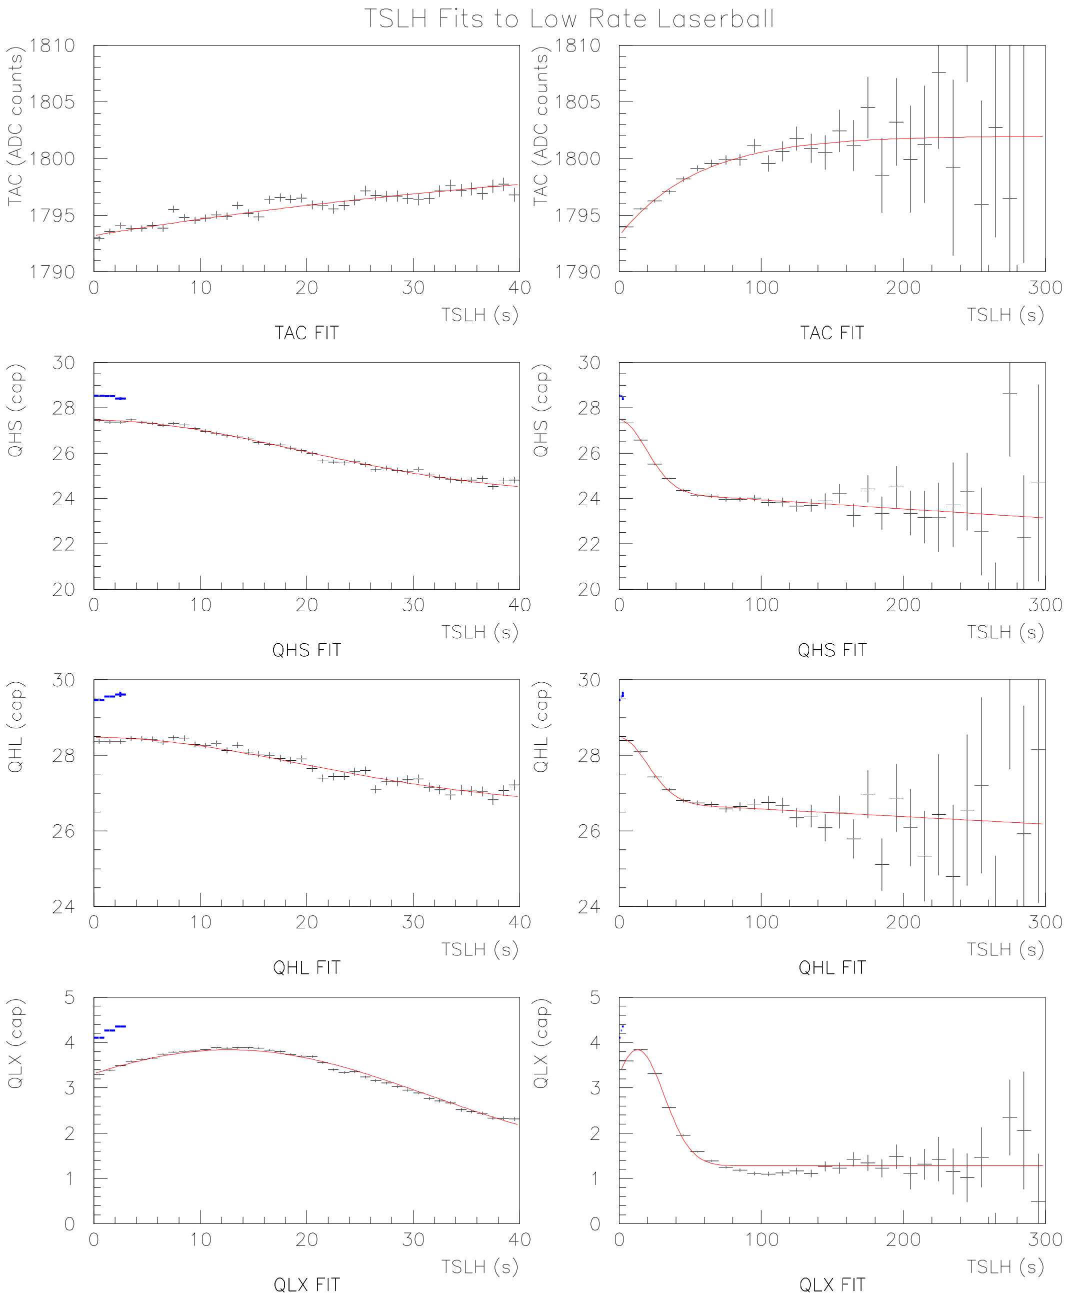

Figure 6.2: Profile histograms of ADC time and pedestal corrected charge against time since last hit (TSLH) for low rate laserball data. Fitted functions are shown in red. The blue points on the charge histograms show the intercept of high rate laserball data at zero TSLH, peculiarly they are 1 ADC count higher than the intercept of low rate data.

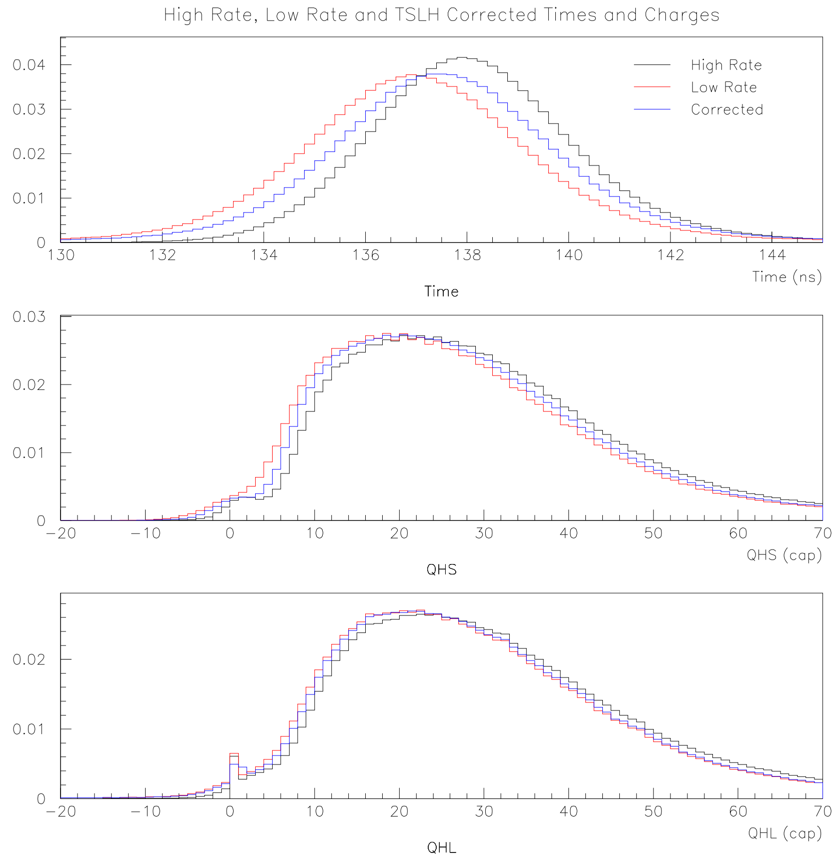

Figure 6.3: Time and charge histograms for high rate laserball data, low rate laserball data and TSLH corrected low rate laserball data.

The fitted functions were used to make a TSLH calibration. The TSLH of each hit was calculated and the time and charge adjusted to the value expected for zero TSLH. This effectively converts all data to high rate. Figure 6.3 shows PCA calibrated time and QHS and QHL histograms for high rate laserball data, low rate laserball data and TSLH corrected low rate laserball data. The situation has been improved but there is still a mismatch between the high rate and corrected low rate data. The χ^2 of the TSLH corrected low rate ^{16}N data is still poor, see figure 6.4.

^{16}N data taken with pedestals pulses interleaved in the run provides another piece of evidence that the high χ^2 problem cannot only be a function of TSLH. Two low rate ^{16}N runs where taken with detector wide pedestals being taken at a rate of

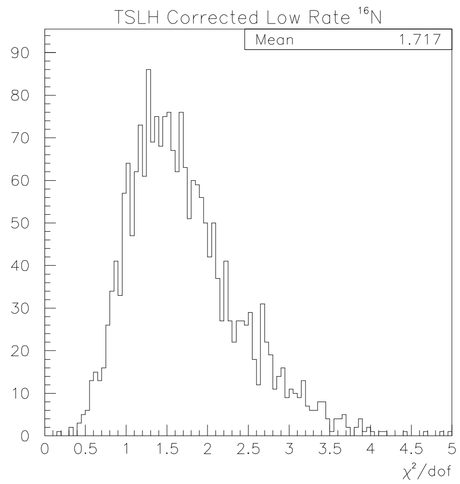

Figure 6.4: The χ^2/dof for TSLH corrected low rate ^{16}N (part of run 12147).

0.1 and 0.2 Hz. The pedestals force all channels to be read out at high rate but the ^{16}N event rate remains low. Figure 6.5 shows the χ^2/dof of the low rate ^{16}N against the time since the last *pedestal* (TSLP). The overall χ^2/dof for this data is clearly much better than low rate ^{16}N without pedestals. The χ^2/dof of a high rate ^{16}N run is shown at zero TSLP. The data from the three runs should coincide at zero TSLP if the high χ^2 problem were simply a function of TSLH because all channels are fired at zero TSLP.

6.3 Rate Correction

A good model of the behaviour of the channel pedestals is that each hit introduces an offset which then decays away over time. If the decay time is long compared to the typical time between hits then a simple rate correction might remove the bulk of the effect with the advantage that the detailed history of the PMT does not need to be known. All that will be needed to calibrate an event is a knowledge of the overall rate of the detector. This means that files containing selected events could be generated without having to record the detailed history of all the channels involved.

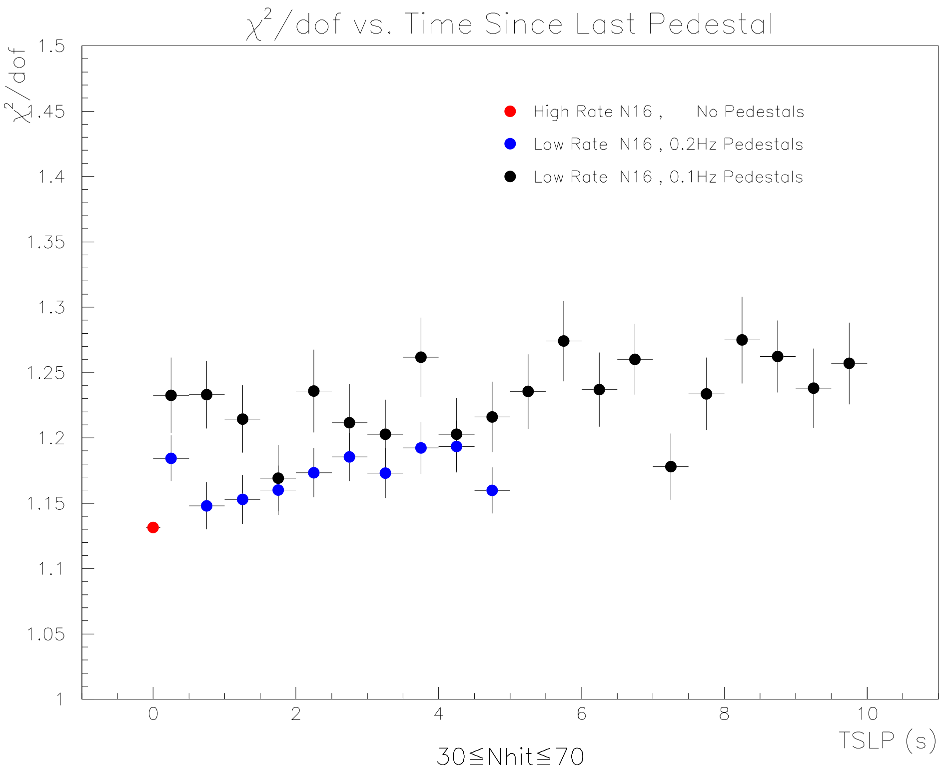

Figure 6.5: The χ^2/dof of ^{16}N data against the time since the last *pedestal* (TSL*P*). The χ^2/dof of high rate ^{16}N run 12280 is shown at zero TSL*P*. The low rate ^{16}N with 0.1 and 0.2 Hz pedestals are from runs 12286 and 12287 respectively. The data from the three runs should coincide at zero TSL*P* if the high χ^2 problem were simply a function of TSL*H* because all channels are fired at zero TSL*P*.

The value of the rate correction was taken from the difference in the average time and charge of each channel between a low and high rate laserball run. The laserball runs used were:

Runno	Nhit	Pulse Rate	Source Position
12251	∼500	35 Hz	0, -21.59, -0.17
12303	∼500	2 Hz	0, -21.59, -0.21

Data was only collected if the PMT was online and had the same electronic board IDs in both runs. First the average charge of every channel was found for both laserball runs. The charge was averaged over the range

$$\text{threshold} - (\text{hhp} - \text{threshold}) \quad \text{to} \quad \text{hhp} + (\text{hhp} - \text{threshold})$$

where the threshold and high half point (hhp) are given by the gain calibration. The gain calibration is done using high rate data so a weakness of this method is that it will be blind to very large changes in the charge pedestal. The typical change in QHS is three or four counts and there were only 32 channels where the change in QHS was more than 20 counts.

The most accurate way to measure the change in the time pedestal is to work with PCA corrected time. The ECA lines up the cells and the PCA removes the cable delay and the walk making the data much less spread and therefore reducing the error on the average. The PCA uses charge information so the charge was first corrected using a charge only rate correction. The average channel time was found by first locating the position of the prompt peak and then averaging all data within ± 7 ns.

Two rate calibrations were made from the collected data. One was the full time and charge correction and one was a time only correction. For the time only calibration, no rate correction was made to the charge when doing the PCA calibration. The purpose of the time only calibration was to get an idea of how much the charge pedestal correction improves the χ^2/dof. Figure 6.6 shows the χ^2/dof for the two calibrations. The rate correction clearly performs much better than the TSLH calibration.

6.4 To Ping or Not To Twiddle.

Two possible fixes to the high χ^2 problem were proposed, firing pedestals for every channel in the detector at a rate of around 0.1 Hz (pinging) or turning off the TAC twiddle bits. The pedestal pinging would raise the channel read out rate to make low rate data look like high rate data. The TAC twiddle bits adjust TAC durations to be very similar for all channels, in the range 380 to 400 ns. Turning off the TAC twiddle bits increases the range of the duration of the TAC ramp to between 340 and 400 ns, but also reduces the size of the TSLH effect. The favoured explanation

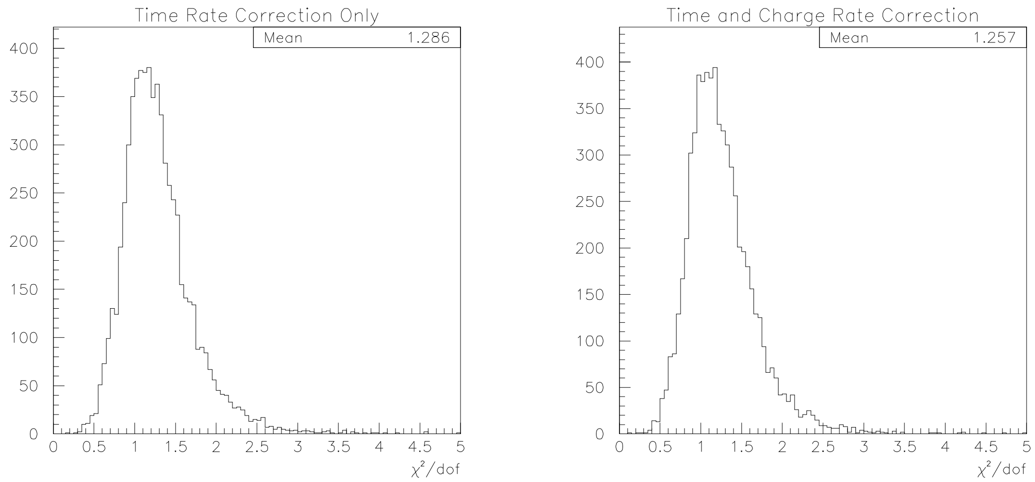

Figure 6.6: The χ^2/dof for rate corrected low rate ^{16}N (run 12147).

for this improvement is that the TAC twiddle bits draw current and heat up the chips. It has been demonstrated that cooling the chips with fans reduces the size of the high χ^2 problem. This section compares the degree to which these two changes reduce the high χ^2 problem.

Four pairs of laserball runs were compared:

- `No Fix:` Run 12251, High rate laserball, TB on
  `&` Run 12303, Low rate laserball, TB on
- `Pinging:` Run 12251, High rate laserball, TB on
  `&` Run 12299, Low rate laserball with 0.1 Hz pedestals, TB on
- `TB A:` Run 12477, High rate laserball, TB off
  `&` Run 12487, Low rate laserball, TB off
- `TB B:` Run 12485, High rate laserball, TB off
  `&` Run 12487, Low rate laserball, TB off

TB A and TB B are two different comparisons of the same thing, i.e. the difference between high and low rate laserball data with the twiddle bits turned off. Table 6.2

gives additional run details. The average charge and average distance corrected time was collected for every channel for the two runs being compared. The strategy was the same as was used in the rate calibration, i.e. the charge was averaged over a range defined by the gain threshold and high half point and the time was averaged over ± 7 ns from the prompt peak. Unlike the rate calibration the charge was not corrected before applying the PCA to the time. Only normal online channels were compared but the electronics board ID check was omitted.

Figure 6.7 shows the difference in the average charge against logical channel number for each of the four pairs of laserball runs. Figure 6.8 shows the error on the differences and figure 6.9 shows histograms of the differences. Pinging reduces the difference between the charge at high and low rates to near zero. Turning off the twiddle bits seems to have little effect. The widths of the distributions are mostly down to the error in the difference. Figures 6.10, 6.11 and 6.12 show the equivalent plots for the time averages. Neither pinging nor turning off the twiddle bits completely solve the problem but both reduce it. Overall pinging is the superior correction.

Despite the fact that pinging the detector reduces the problem more than turning of the twiddle bits it was not the chosen solution. Pinging the detector would have greatly increased the data rate and made the detector more prone to filling the circular buffer (where PMT data is held until it is united with its associated trigger information) during high rate bursts. The orphan rate, the rate at which

Run	Twiddle Bits	Pedestals	Pulse Rate	Nhit	Position
12251	TB on	—	35 Hz	~500	0, -21.59, -0.17
12299	TB on	0.1 Hz	2 Hz	~60	0, -21.59, 400.07
12303	TB on	—	2 Hz	~60	0, -21.59, -0.21
12477	TB off	—	35 Hz	~800	0, -21.59, -0.47
12485	TB off	—	45 Hz	~60	0, -21.59, 400.17
12487	TB off	—	2 Hz	~60	0, -21.59, 400.17

Table 6.2: The laserball runs used to investigate the effect of firing pedestals and turning off the TAC twiddle bits.

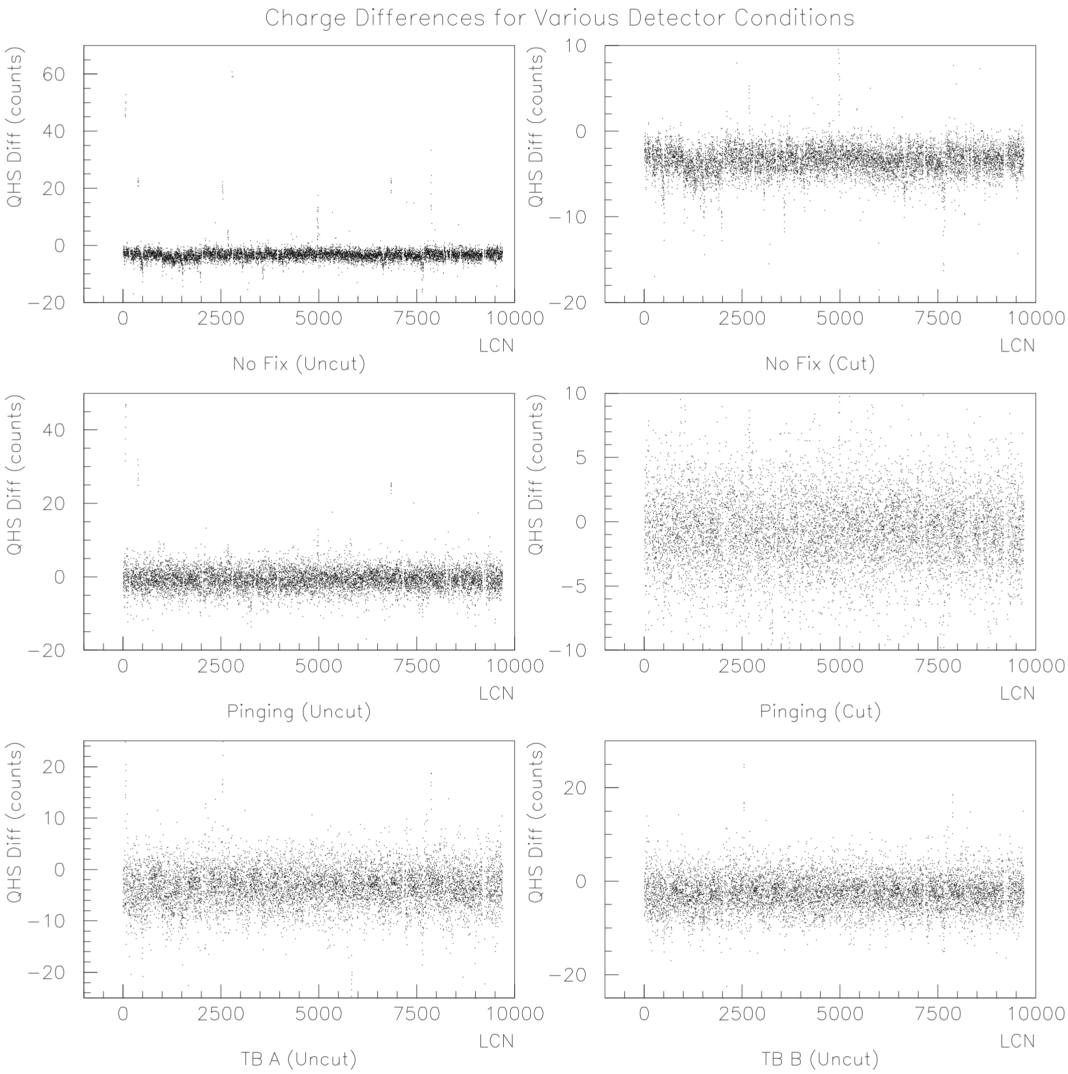

Figure 6.7: The difference in the average charge against logical channel number for four pairs of laserball runs. The plots for 'no fix' and for 'pinging' are shown for both all channels and with large outliers cut. The plots for 'TB A' and 'TB B' do not have any large outliers. There is no visible correlation with logical channel number in the pinging and twiddle bit data.

the builder fails to combine all PMT bundles with their associated trigger, would probably also have increased. The twiddle bits were turned off on 17th August 2000.

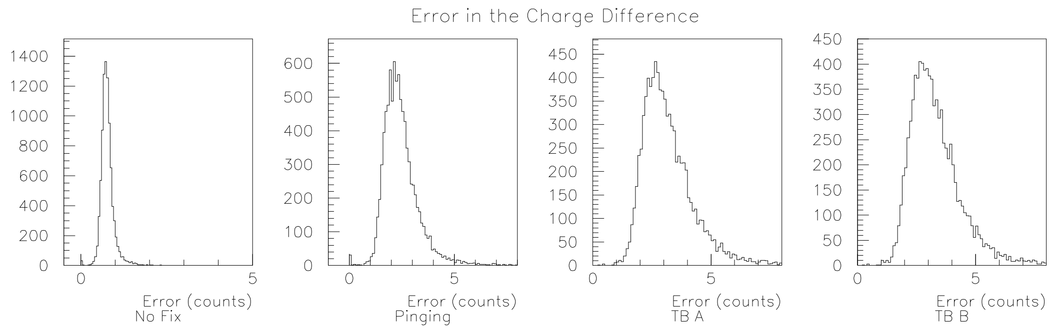

Figure 6.8: The error in the difference of the average charges for four pairs of laserball runs.

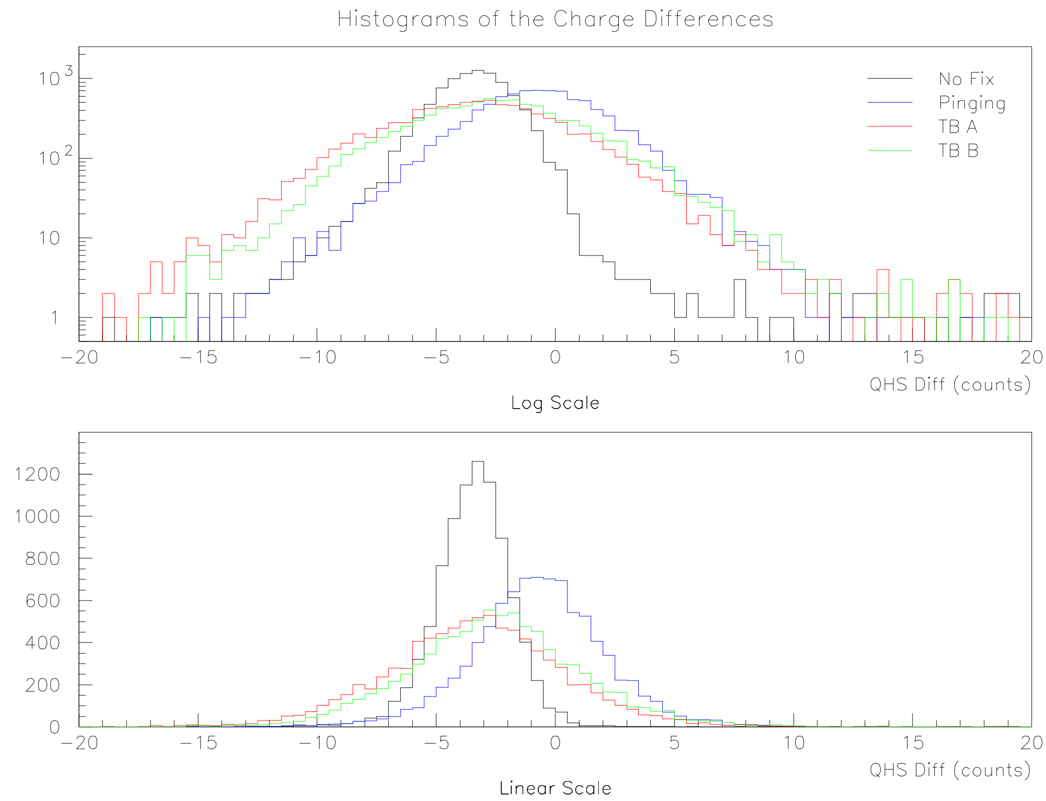

Figure 6.9: Histograms of the difference in the average charge for four pairs of laserball runs. The widths of the peaks are largely dictated by the error on the difference rather than by channel to channel variations. In other words the overall shift is significant but the spread is not. It can be seen that pinging cures the shift in charge while turning of the TAC twiddle bits has no discernible effect, as might be expected.

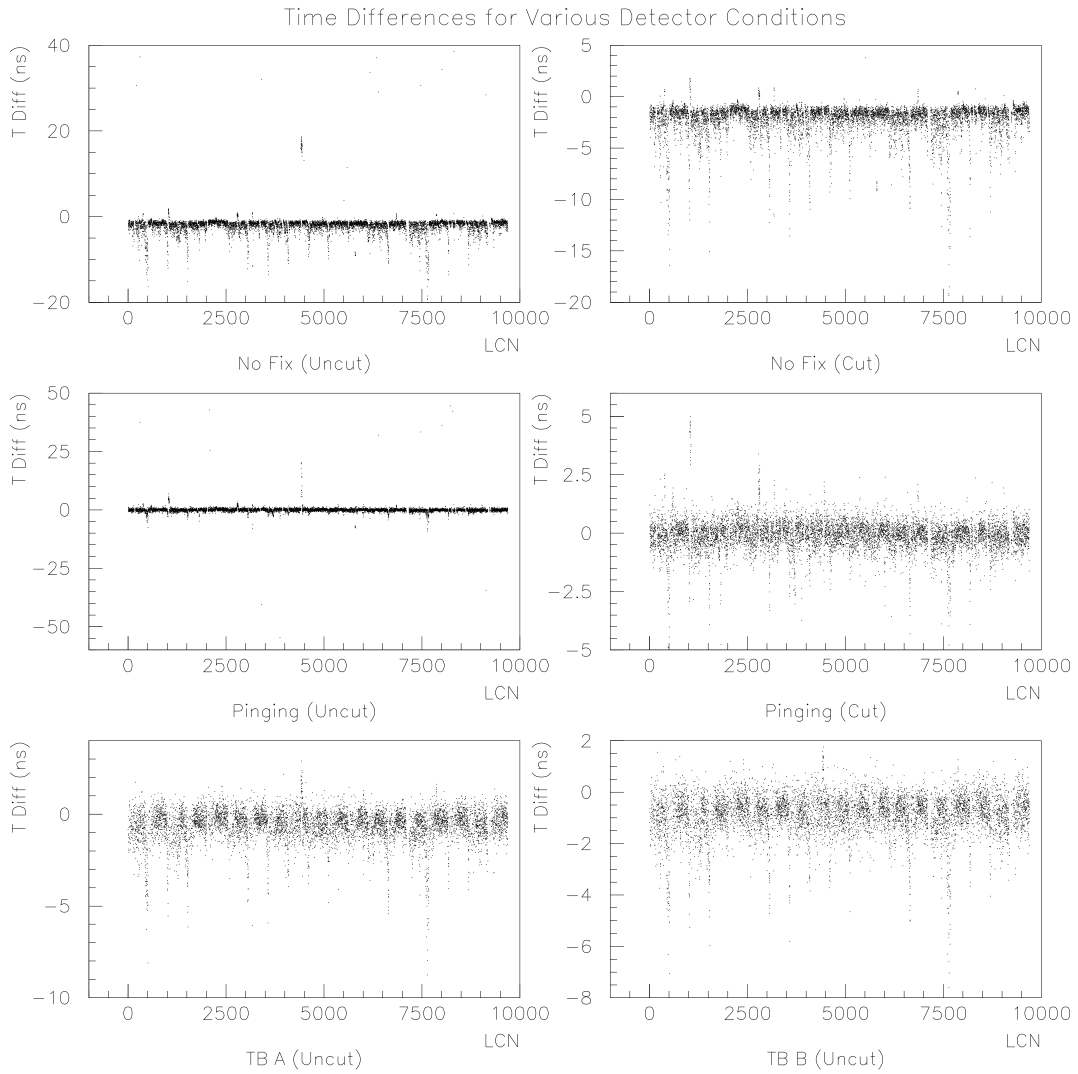

Figure 6.10: The difference in the average time against logical channel number for four pairs of laserball runs. The plots for 'no fix' and for 'pinging' are shown for both all channels and with large outliers cut. The plots for 'TB A' and 'TB B' do not have any large outliers. The regular spikes of outliers correspond to card 15s. Both pinging and turning off the twiddle bits still suffer from outliers but at a reduced level.

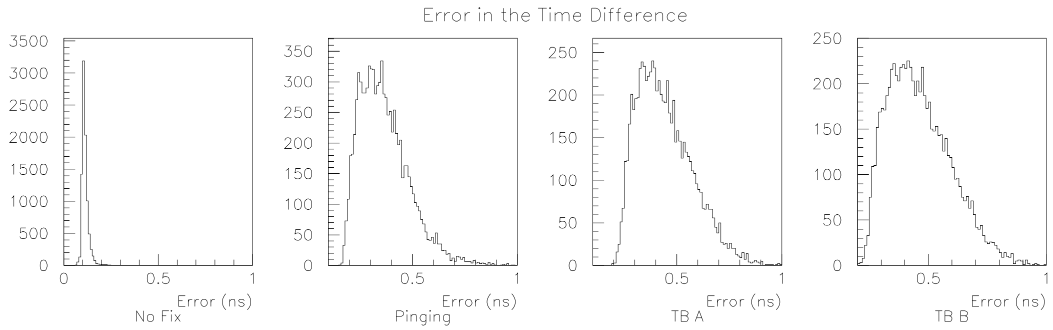

Figure 6.11: The error in the difference of the average times for four pairs of laserball runs.

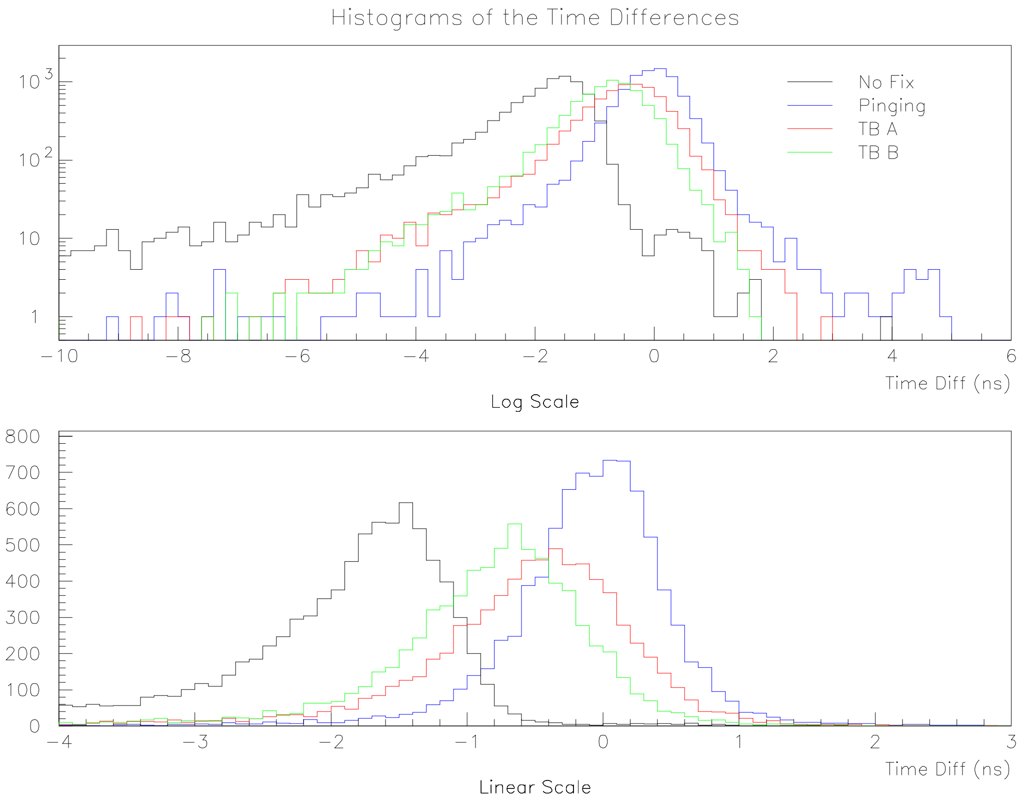

Figure 6.12: Histograms of the difference in the average time for four pairs of laserball runs. Both pinging and turning off the twiddle bits reduce the difference but pinging is superior.

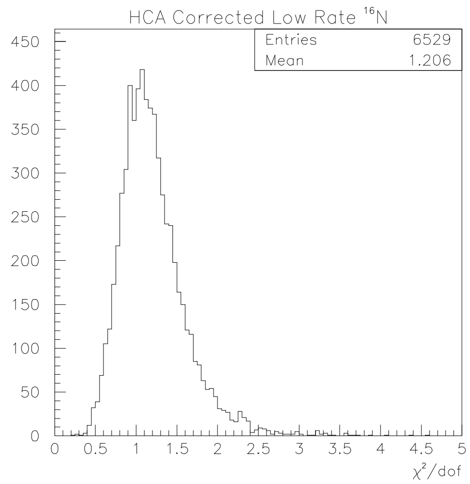

Figure 6.13: The χ^2/dof for HCA corrected low rate ^{16}N (run 12147).

6.5 Current Status

This chapter has described how the charge and time pedestals are sensitive to the time since last hit and the overall rate. As investigations continued it was found that the change in the time pedestal was not a constant for all TAC values but was actually a slope change of the TAC and the two TACs need to be treated separately. Another discovery was that the channels which were most cooled by fans were the least afflicted with pedestal shifts. For this reason more fans were installed to cool all channels equally. This took effect from 1st December 2000.

The correction of these effects in SNOMAN is done by the History Calibration processor, the HCA, written by Gordon McGregor [50]. The HCA currently only corrects the time. From the history of the PMT it estimates the amount by which the TAC slope will have changed. The correction is derived from special variable rate pedestal runs. Figure 6.13 shows the χ^2/dof of HCA corrected low rate ^{16}N.

In the future the HCA will be extended to include a charge correction. This is an important extension, not only because it should further improve the timing of low rate data, but it may also result in slightly more channels passing the PCA stage in the calibrator. As was explained in section 4.6, the PCA has a charge threshold below which PMTs are not calibrated. At low rate the charge spectrum is lower than at high rate and so is closer to the PCA cut. The PCA cut is typically 3 counts below the DAQ threshold but might still result in a small fraction of the data being failed. In high rate N16 run 12144 the fraction of hits which fail the PCA charge threshold is around 1.4%, most of which is probably pickup noise. In low rate N16 run 12147 this figure rises to 3.5%. Some of the increase could be attributed to pickup noise which passed in the high rate configuration but fails in the low rate configuration, but probably not all of it.

Chapter 7

Conclusions

This thesis has discussed the purpose, implementation and performance of the Photomultiplier Tube Calibration. Some rate related problems were also investigated. The timing characteristics of the PMTs has been discussed at length but the impact of improved timing resolution on data analysis has not been considered. Section 7.1 briefly demonstrates the positive impact of the PCA calibration on ^{8}Li event reconstruction. In section 7.2 Dave's promise [51] finally comes good with a sneak look at neutrinos. Section 7.3 draws some final conclusions.

7.1 ^{8}Li Reconstruction

The ability to fit the location of events in the detector was the prime motivation for the time calibration. The benefit of improved timing is improved fitter performance. The simplest way to demonstrate improved fitter performance is using Čerenkov events with a known location. The ^{8}Li source provides such events with an nhit spectrum (the number of PMTs that fire in an event) that spans the nhit range of neutrino data.

Run 14348 is a long ^{8}Li run with the source close to the centre of the detector. The manipulator system reports the source position to be (-16.03, 23.50, 0.08). The distance between the source position and the fitted event vertex for four levels of calibration are shown in figure 7.1. The green lines show the fit position for data calibrated with the ECA plus HCA only (some of the ECA validation cuts were used, words 15 and 16 of the TCAL bank were set to 1, see [43][35]). The blue lines show the results when using the ECA, HCA and cable delay corrected times. The cable delay corrected times were found by modifying the calibrator to use the high charge linear fit rather than the full walk interpolation (the PCA charge cut for data points too far below the first interpolation point was left unaltered). The black lines show the results for the full calibration and the red dotted lines show the results for the full calibration but without the HCA correction. Note that the ^{8}Li is not a good source for testing the HCA because the run was taken in December 2000 after the twiddle bits were turned off and additional fan cooling was installed. Also the detector rate for ^{8}Li is somewhere between normal calibration (high) rate and data taking (low) rate. The fit performance of fully calibrated data and calibrated data without the HCA has been studied in detail using ^{16}N [52]. It is clear that the PCA greatly improves upon the fit performance of the ECA. Most of the improvement is due to correcting the cable delays but removing the discriminator walk also results in a noticeable change.

Runs 14358 and 14360 were taken with the ^{8}Li source at (-16.03, 23.50, 450.05). The distance between the source position and the fitted event vertex for the same levels calibration are shown in figure 7.2. Figure 7.3 shows the difference between the fitted X, Y and Z and the source X, Y and Z with the source at the centre of the detector and at Z = 450 cm. This plot shows that not only is the ECA plus HCA calibration very low resolution but it also suffers from a systematic vertex shift. There is no systematic vertex shift when just using the cable delay corrected times but the resolution is worse, see figure 7.4.

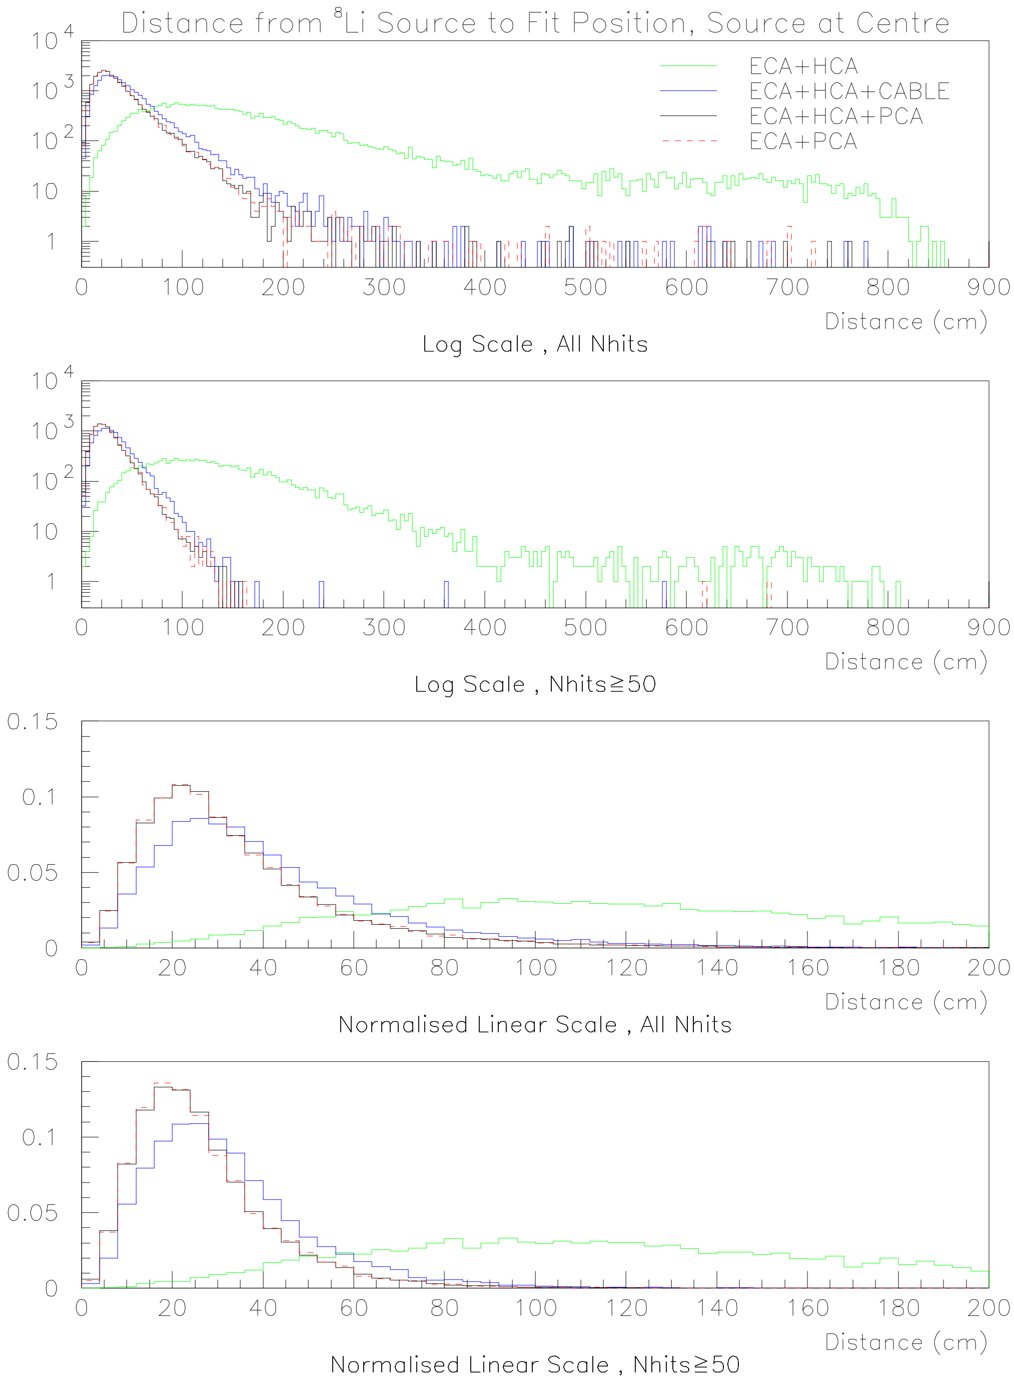

Figure 7.1: The distance from the ^{8}Li calibration source, located at the centre of the detector, to the fit position for various calibration levels. The path fitter was used to find the fit vertex and the source position was that reported by the manipulator system.

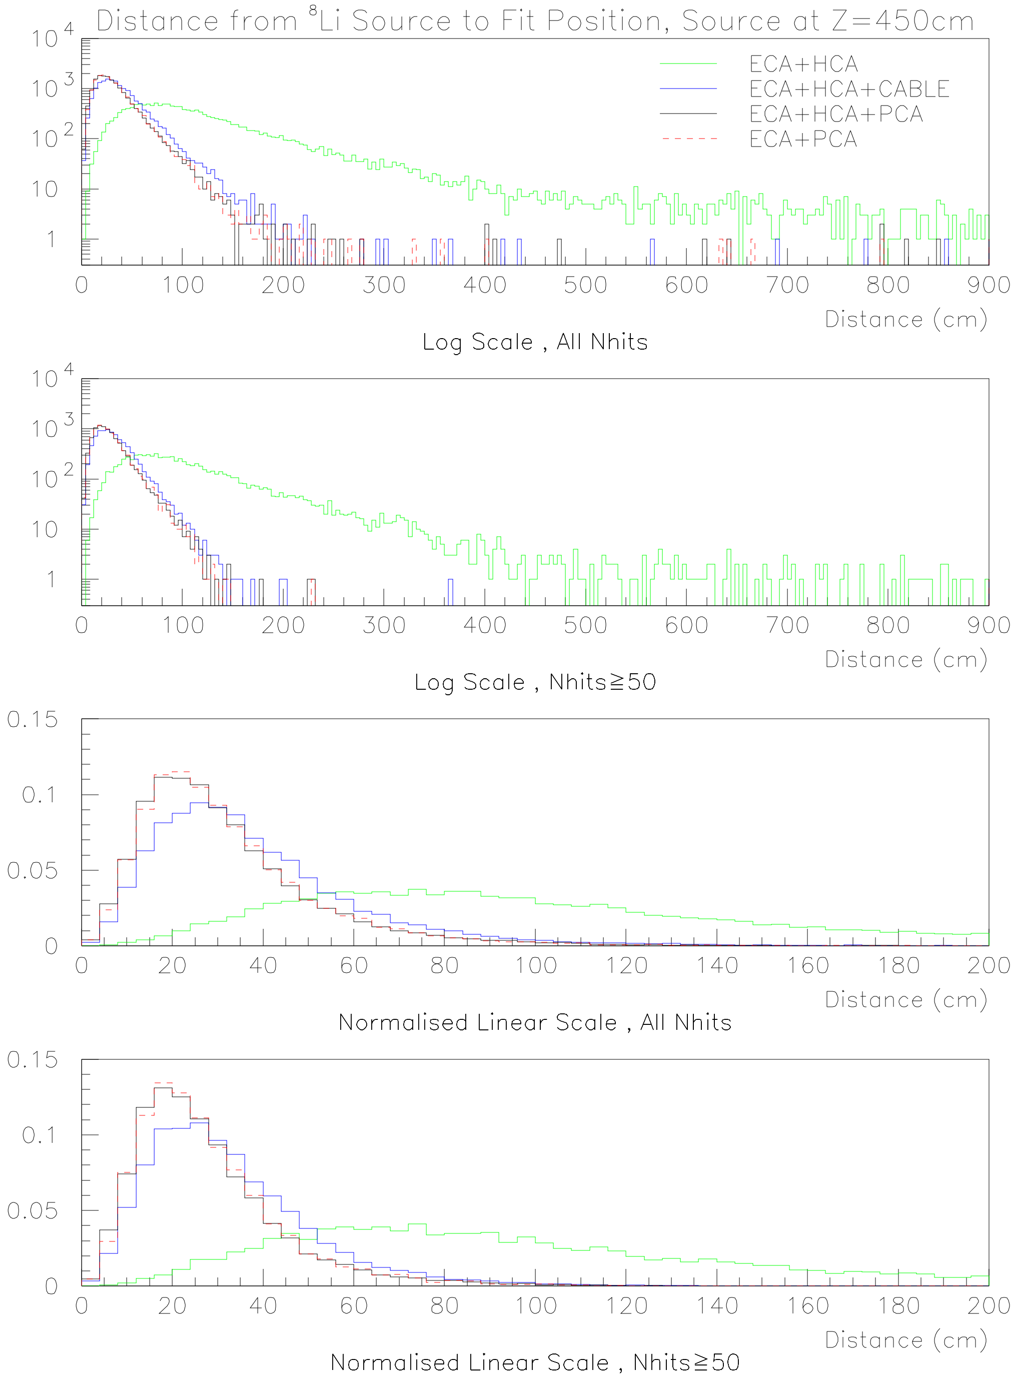

Figure 7.2: The distance from the ^{8}Li calibration source, located at approximately X=-16, Y=24, Z=450 cm, to the fit position for various calibration levels. The path fitter was used to find the fit vertex and the source position was that reported by the manipulator system.

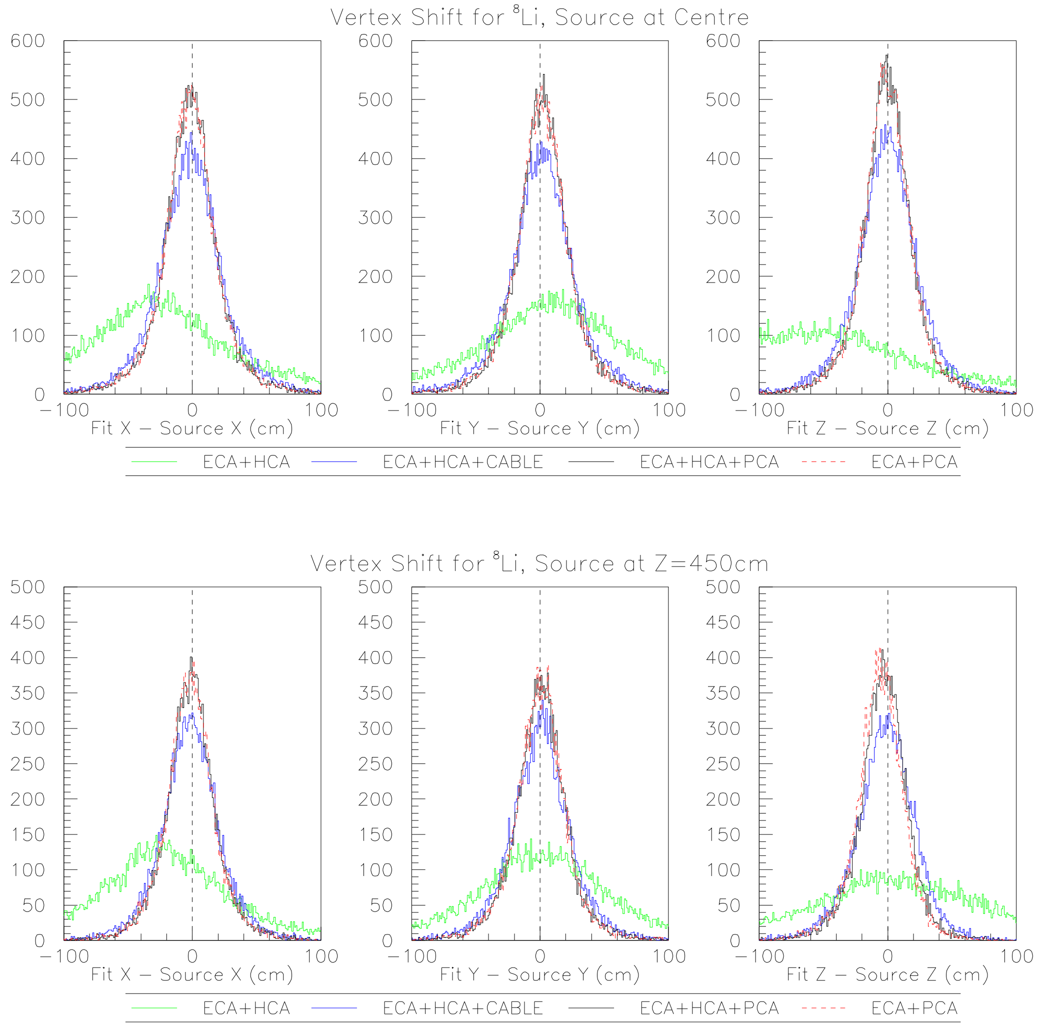

Figure 7.3: The difference between the fitted event X, Y and Z positions and the source X, Y and Z for various calibration levels. The ^{8}Li source was at the centre of the detector for the upper plot and at approximately X=-16, Y=24, Z=450 cm for the lower plot.

7.2 Neutrino Reconstruction

The same cleaned data set used in the first publication of SNO results, see section 1.6.5 and [28], was used to look at the effect of different levels of calibration. The data was recorded between 2nd November 1999 and 15th January 2001, and has a live time of 240.95 days. The data was calibrated to the desired level and the events fitted. Any event following a muon by less than 20 seconds was cut. None of

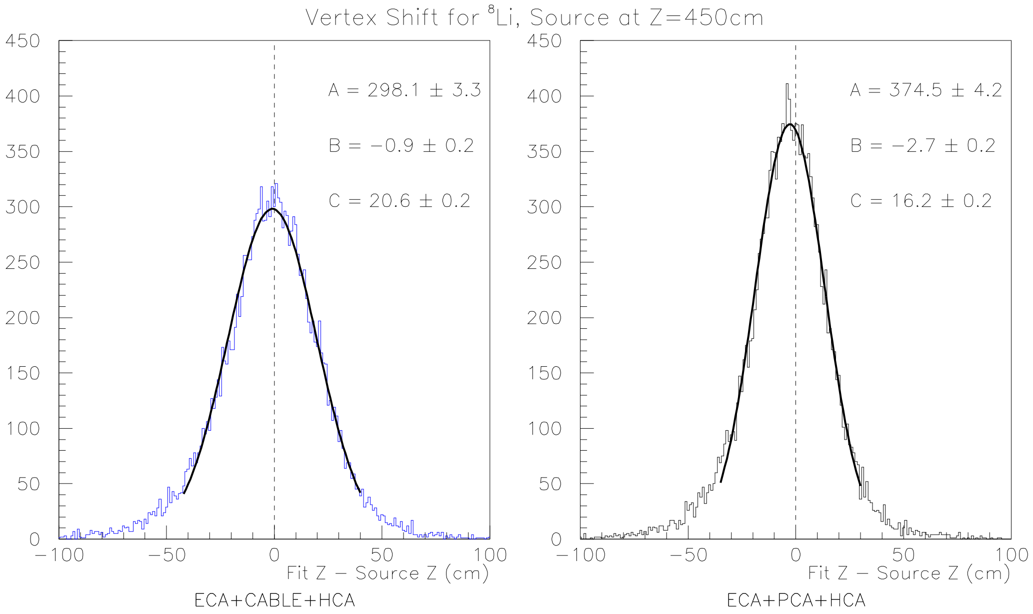

Figure 7.4: The difference between the fitted event Z positions and the source Z position for cable delay corrected time and fully PCA calibrated time with the source at X=-16, Y=24, Z=450 cm. A, B and C give the values of the height, position and width of the fitted Gaussians. There is a difference of 2.7 cm between the most common fitted Z position and the manipulator position when using the fully calibrated time. This discrepancy is probably within the error in the manipulator position. There could also be a systematic effect from the source being close to the neck of the acrylic vessel.

the high level cuts applied to the published results were used.

Figure 7.5 shows the nhit spectrum of data which reconstructs inside a radius of 550 cm. The events were fit using the grid fitter [53]. The spectrum is very similar at high nhit but at low nhit the lowest level calibration, ECA plus HCA, clearly produces many more events. Figure 7.6 shows the radial distribution of the reconstructed events in units of acrylic vessel radii cubed for different calibration levels and different nhit ranges. Plotting the cube of the radius makes each bin correspond to an equal volume of water. The reason for the large number of events reconstructing inside 550 cm at low nhit when using the ECA plus HCA calibration is now obvious. Radioactive background from outside the D_2O is being wrongly

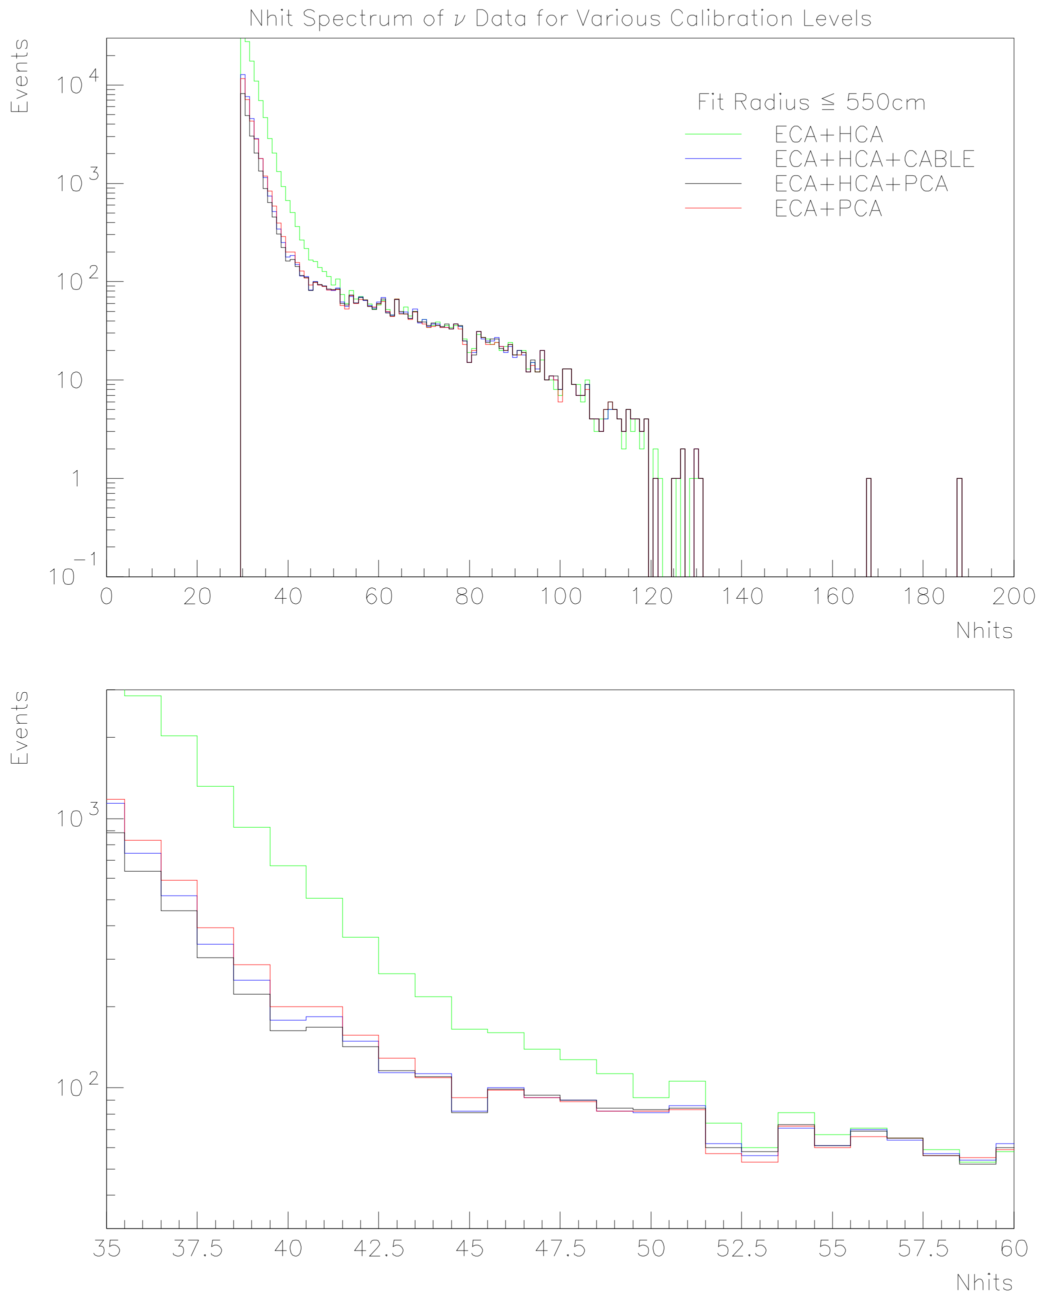

Figure 7.5: The nhit spectrum of neutrino data (without high levels cuts) which reconstructs, using the grid fitter, inside 550 cm for various levels of calibration, see the key. The lower plot is simply a close up of the upper plot for a restricted nhit range.

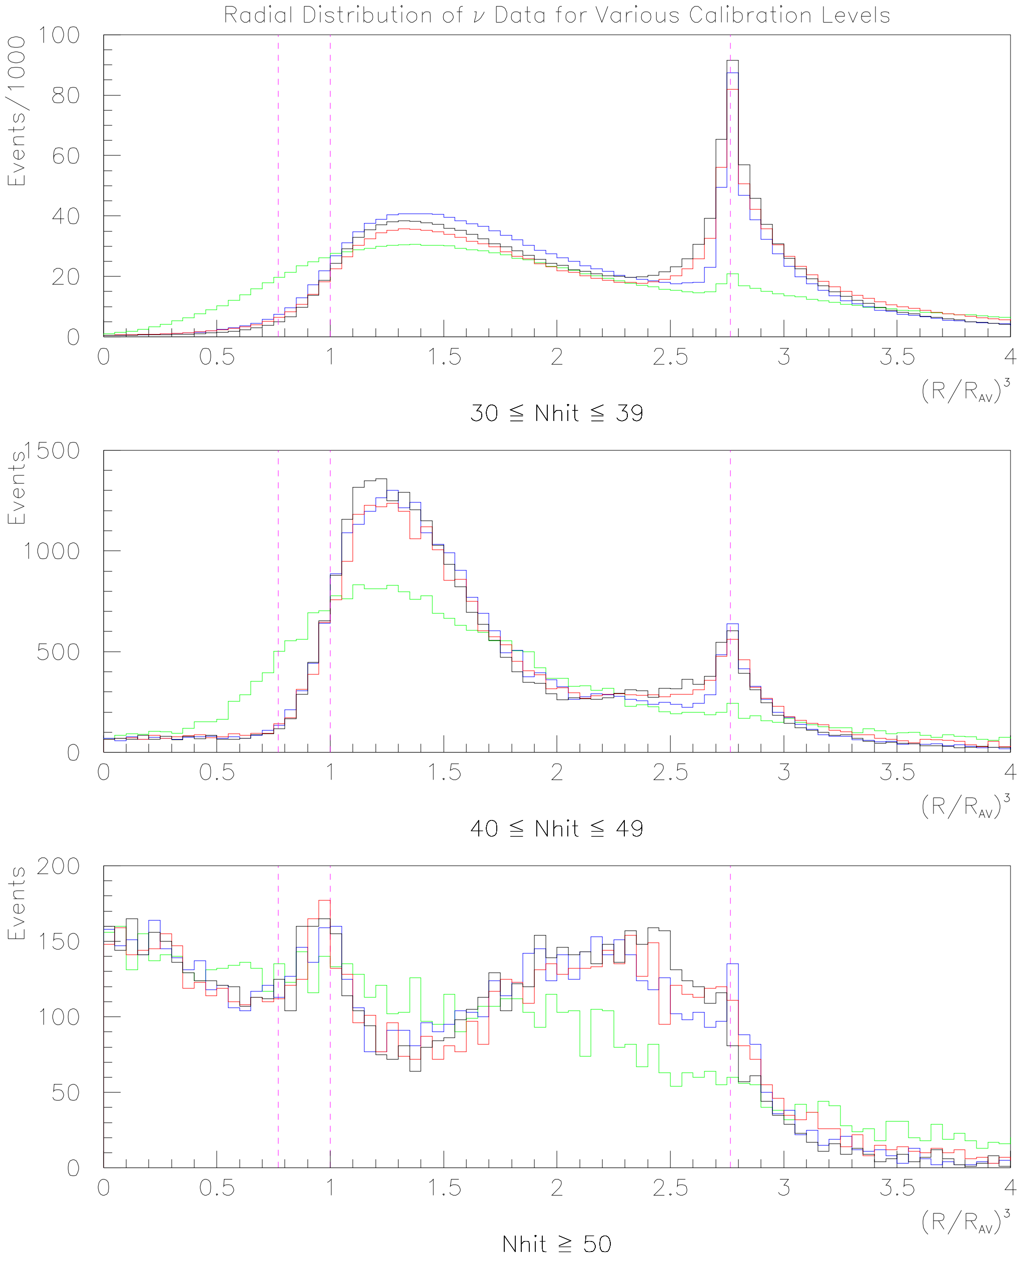

Figure 7.6: The radial distribution of grid fitted neutrino data (without high level cuts) for different calibration levels and nhit ranges. The green line corresponds to the lowest level of calibration, ECA plus HCA. The blue line is ECA, HCA and cable delay calibrated data. The black line is the full ECA, HCA and PCA calibrated data. The red line is the full calibration without the HCA correction. The three vertical dotted magenta lines are, from left to right, 550 cm, the acrylic vessel radius and the PSUP radius.

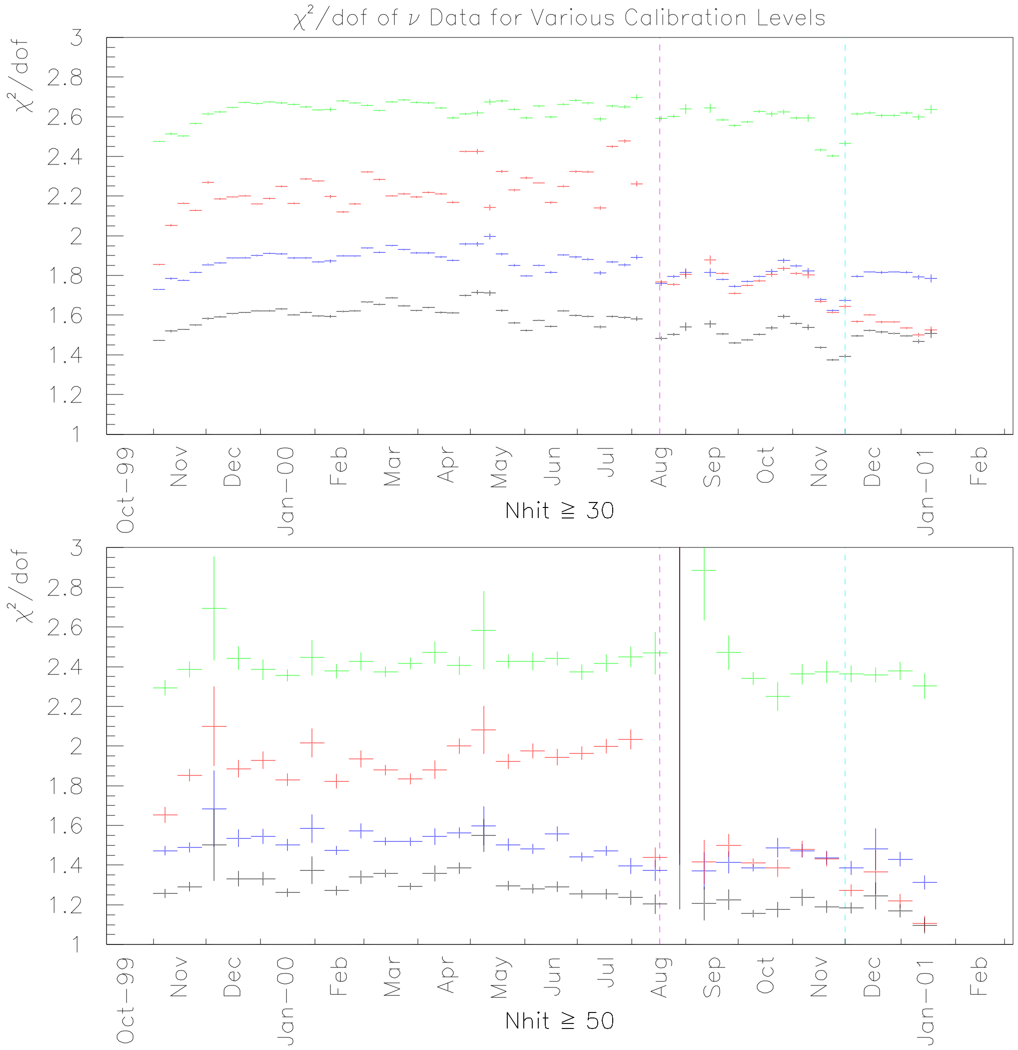

Figure 7.7: The χ^2/dof of neutrino data, using the time fitter, against date for various calibration levels. Note that the y axis runs from 1 to 3. The green line is for ECA and HCA calibrated data. The blue line is for ECA, HCA and cable delay calibrated data. The black line is for fully calibrated data and the red line is for fully calibrated data but without the HCA correction. The nhit $\geq$ 30 histogram has 7 day bins while the nhit $\geq$ 50 histogram has 14 day bins. The vertical dotted magenta line indicates the date when the TAC twiddle bits were turned off and the vertical dotted cyan line indicates the date when additional fan cooling was installed.

reconstructed inside 550 cm. The agreement at high nhit is because there is very little background at high nhit and not because the fits are improved at high nhit. The uppermost plot, corresponding to the 30–39 nhit range, shows some interesting variations between the other calibrations. The full calibration (black) reconstructs the least number of events inside the acrylic vessel and the most number of events to the PMT support structure (PSUP). The cable delay calibration (blue) reconstructs less events to the PSUP and more in the H_2O and D_2O regions. The full calibration without the HCA correction (red) results in more events reconstructed in the D_2O but less events in the H_2O. This calibration results in the most number events reconstructed in the unphysical region outside the detector with the exception of the ECA plus HCA calibration.

Figure 7.7 shows the χ^2/dof of neutrino data against date for the various calibration levels. The vertical dotted lines indicate the dates that hardware improvements were made to reduce the high χ^2 problem. The resulting drop in the χ^2/dof of the red points testifies to the success of those measures.

7.3 Final Conclusions

This thesis has described the development and testing of the Photomultiplier Calibration of the Sudbury Neutrino Observatory. The PMT calibration has two main parts: the walk calibration and the gain calibration. The purpose of the walk calibration is to correct the timing of the PMTs for the effects of differing cable delays and discriminator walk. The purpose of the gain calibration is to convert the integrated charge of all the PMTs into the same units.

Chapter 3 describes procedures by which the walk and gain calibrations can meet their goals. Chapter 4 gave a somewhat lengthy account of how these methods were implemented in SNOMAN. The salient points of how the resulting calibration constants are applied to neutrino data were also covered. The chapter concludes with a description of how badly behaving channels are identified and cut.

Chapter 5 presented a detailed investigation of the performance of the PMT calibration. It was demonstrated that the walk calibration restores the PMT timing performance to the extent possible. The optical path correction used by the walk calibration was also investigated and found to be adequate. A method for fitting the position of the laserball, independent of the manipulator system, was described and the results for all PCA laserball runs given. The effect of different light intensities on the walk and gain calibrations was shown to be negligible for the intensity regimes spanned by the PMT calibration and neutrino data. Finally, the stability of the calibration constants was discussed. Chapter 6 contains some studies on the problems caused by the rate sensitivity of the electronics.

This chapter concludes the thesis by demonstrating the beneficial impact of the PMT calibration on event reconstruction. The bottom line can be seen by comparing the green and the black curves in figure 7.1. Event reconstruction is essential for the separation of the different types of neutrino interactions occurring in SNO, and no meaningful SNO solar neutrino analysis would be possible without the work described in this thesis. On a broader level, all analyses of SNO data will use the PCA timing calibration, and many will also use the gain. The successful development of the PMT calibration is therefore a crucial contribution to SNO.

References

[1] J. Bonn et al. *The Mainz Neutrino Mass Experiment.* In *Neutrino 2000, Proceedings of the XIXth International Conference on Neutrino Physics and Astrophysics*, editors J. Law, R. Ollerhead and J. Simpson, Nuclear Physics B (Proc. Suppl.) **91**, 273 (2001).

[2] V. Lobashev et al. *Direct Search for Neutrino Mass and Anomaly in the Tritium Beta-Spectrum: Status of "Troitsk Neutrino Mass" Experiment.* In *Neutrino 2000, Proceedings of the XIXth International Conference on Neutrino Physics and Astrophysics*, editors J. Law, R. Ollerhead and J. Simpson, Nuclear Physics B (Proc. Suppl.) **91**, 280 (2001).

[3] D.E. Groom and others. *Review of Particle Physics.* The European Physics Journal C **15**, 1 (2000).

[4] J. Roney. *Review of the Tau Neutrino Mass.* In *Neutrino 2000, Proceedings of the XIXth International Conference on Neutrino Physics and Astrophysics*, editors J. Law, R. Ollerhead and J. Simpson, Nuclear Physics B (Proc. Suppl.) **91**, 287 (2001).

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

[6] J. N. Bahcall, M. Pinsonneault and S. Basu. *Solar Models: current epoch and time dependences, neutrinos, and helioseimological properties.* Astrophysics Journal **555**, 990 (2001). astro-ph/0010346.

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

[8] B. Cleveland et al. *Measurement of the solar electron neutrino flux with the Homestake chlorine detector.* Astrophysics Journal **496**, 505 (1998).

[9] J. Abdurashitov et al. *Measurement of the solar neutrino capture rate with gallium metal.* Physical Review C **60**, 055801 (1999).

[10] V. Gavrin et al. *Solar neutrino results from SAGE.* In *Neutrino 2000, Proceedings of the XIXth International Conference on Neutrino Physics and Astrophysics*, editors J. Law, R. Ollerhead and J. Simpson, Nuclear Physics B (Proc. Suppl.) **91**, 36 (2001).

[11] W. Hampel et al. *GALLEX solar neutrino observations: results for GALLEX IV.* Physics Letters B **447**, 127 (1999).

[12] M. Altmann et al. *GNO solar neutrino observations: results for GNO I.* Physics Letters B **490**, 16 (2000).

[13] E. Bellotti et al. *First Results from GNO.* In *Neutrino 2000, Proceedings of the XIXth International Conference on Neutrino Physics and Astrophysics*, editors J. Law, R. Ollerhead and J. Simpson, Nuclear Physics B (Proc. Suppl.) **91**, 44 (2001).

[14] Y. Fukuda et al. *Solar Neutrino Data Covering Solar Cycle 22.* Physical Review Letters **77**(9), 1683 (1996).

[15] S. Fukuda et al. *Solar ^8B and hep Neutrino Measurements from 1258 Days of Super-Kamiokande Data.* Physical Review Letters **86**, 5651 (2001). hep-ex/0103032.

[16] N. Hata and P. Langacker. *Solutions to the solar neutrino anomaly.* Physical Review D **56**(10), 6107 (1997). hep-ph/9705339.

[17] B. Pontecorvo. *Neutrino experiments and the problem of conservation of leptonic charge.* Soviet Physics JETP **26**(5), 984 (1968).

[18] V. Gribov and B. Pontecorvo. *Neutrino Astronomy and Lepton Charge.* Physics Letters B **28**(7), 493 (1969).

[19] B. Kayser. *On the quantum mechanics of neutrino oscillation.* Physical Review D **24**(1), 110 (1981).

[20] L. Wolfenstein. *Neutrino oscillations in matter.* Physical Review D **17**(9), 2369 (1978).

[21] S. Mikheyev and A. Smirnov. *Resonance enhancement of oscillations in matter and solar neutrino spectroscopy.* Soviet Journal of Nuclear Physics **42**(6), 913 (1985).

[22] S. Mikheyev and A. Smirnov. *Resonant Amplifications of ν-Oscillations in Matter and Solar-Neutrino Spectroscopy.* Il Nuovo Cimento **9C**, 17 (1986).

[23] G. Fogli et al. *Model-dependent and independent implications of the first Sudbury Neutrino Observatory results.* hep-ph/0106247.

[24] M. Apollonio et al. *Limits on neutrino oscillations from the CHOOZ experiment.* Physics Letters B **466**, 415 (1999).

[25] H. H. Chen. *Direct Approach to Resolve the Solar-Neutrino Problem.* Physical Review Letters **55**(14), 1534 (1985).

[26] J. Boger et al. *The Sudbury Neutrino Observatory.* Nuclear Instruments and Methods in Physics Research A **449**(1–2), 172 (2000). nucl-ex/9910016.

[27] P. Vogel and J. Beacom. *The angular distribution of the reaction* $\bar{\nu}_e + p \rightarrow e^+ + n$. Physical Review D **60**, 053003 (1999). hep-ph/9903554.

[28] Q. Ahmad et al. *Measurement of the rate of* $\nu_e + d \rightarrow p + p + e^-$ *interactions produced by* ^8B *solar neutrinos at the Sudbury Neutrino Observatory.* Physical Review Letters **87**, 071301 (2001). nucl-ex/0106015.

[29] A. Hime. *Flavor Content Analysis of the* ^8B *Solar Neutrino Flux.* SNO Internal Report (2001).

[30] R. Taplin. *The Use of Photomultipliers in SNO.* Ph.D. thesis, Oxford University (1995).

[31] J. Cameron et al. *The Gain and the Pain: A Study of Detector Response and Instrumental Backgrounds After the High Voltage Raise.* SNO-STR-99-039 (1999).

[32] E. Beier et al. *Performance of the Electronics for the Sudbury Neutrino Observatory.* SNO-STR-98-002 (1998).

[33] D. Cowen et al. *The Sudbury Neutrino Observatory Electronics Chain* (2000).

[34] J. Klein et al. *The SNO Trigger System.* SNO-STR-97-035 (1997).

[35] S. Biller et al. *SNO Electronic Calibration Constants.* SNO-STR-01-005 (2001).

[36] R. J. Ford. *Calibration of SNO for the Detection of* ^8B *Neutrinos.* Ph.D. thesis, Queen's University at Kingston (1998).

[37] B. A. Moffat. *The Optical Calibration of the Sudbury Neutrino Observatory.* Ph.D. thesis, Queen's University at Kingston (2001).

[38] M. R. Dragowsky. *Sudbury Neutrino Observatory Energy Calibration using Gamma-Ray Sources.* Ph.D. thesis, Oregon State University (1999).

[39] N. J. Tagg. *The* ^8Li *Calibration Source and Through-Going Muon Analysis in the Sudbury Neutrino Observatory.* Ph.D. thesis, The University of Guelph (2001).

[40] N. McCauley. *Sacrifice update to include recent* ^8Li *data.* SNO Internal Report (2001).

[41] The SNO Collaboration. *The SNOMAN User's Manual.* Available with the latest SNOMAN release.

[42] The SNO Collaboration. *The SNOMAN Programmer's Manual.* Available with the latest SNOMAN release.

[43] The SNO Collaboration. *The SNOMAN Companion*. Available with the latest SNOMAN release.

[44] *PAW – Physics Analysis Workstation*. CERN Program Library Long Writeups Q121 (1999).

[45] *ZEBRA*. CERN Program Library Long Writeups Q100/Q101 (1995).

[46] D. Cowen et al. *The SNO Database: SNODB v3_06_05* (1998). available at `http://www.hep.upenn.edu/SNO/`.

[47] M. Neubauer. *Evidence for Electron Neutrino Flavour Change Through Measurement of the ^8B Solar Neutrino Flux at the Sudbury Neutrino Observatory*. Ph.D. thesis, University of Pennsylvania (2001).

[48] J. Klein and M. Neubauer. *Using Time and Angle Information in Reconstruction and the SNOMAN path fitter*. SNO-STR-00-022 (2000).

[49] W.H.Press et al. *Numerical Recipes in Fortran 77: The Art of Scientific Computing* (Cambridge University Press, 1999), 2nd edn.

[50] G. McGregor. Private communication.

[51] D. Wark. *"You will have solar neutrino data for your thesis"*. Private communication.

[52] SNO Collaboration. *Analysis for First Solar Neutrino Paper*. SNO Internal Report (2000).

[53] X. Chen. *An Improved Grid Fitter*. SNO-STR-00-004 (2000).