The Use of Photomultipliers in SNO

Richard Keith Taplin
Hertford College, Oxford

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

Michaelmas Term 1995

The Use of Photomultipliers in SNO

Richard Keith Taplin
Hertford College, Oxford

Thesis submitted for the degree of Doctor of Philosophy

Michaelmas Term 1995

Abstract

The comparison of the predicted and measured fluxes of neutrinos from the sun has now been carried out by four independent experiments. All four have seen a deficit between the rate measured and the rate predicted from complex, iterative models of the sun. This Solar Neutrino Problem has been a source of interest for around 20 years, with possible explanations ranging from astrophysics to new particle physics. The Sudbury Neutrino Observatory (SNO) aims to resolve the problem by the direct comparison of the flux of electron neutrinos to the flux of all flavours of neutrinos, and hence check various models which suggest that the neutrinos are oscillating from electron to muon or tau flavour on their passage from the core of the sun to the earth.

SNO itself is a 1 kilo-tonne heavy water Čerenkov detector, which can measure the electron neutrino flux by the break-up of deuterium via a W boson, the CC reaction; and can measure the flux of all neutrinos with the weak disintegration of the deuteron via a Z boson, the NC reaction. In order to make the comparison between the two channels, the backgrounds on each channel must be very well known. The main sources of background are the long lived isotopes of thorium and uranium, and their subsequent decay chains, and to a lesser extent ^{40}K. The cleanliness requirements for the detector have meant that new regimes of contamination have had to be considered, and new techniques for assay developed.

In order to assay to the levels required, new counters have been developed by the author which use the fast β-α coincidences at the ends of the uranium and thorium chains. The counters have a high efficiency, source/counter separation, and at present have a low enough background to assay the heavy water in SNO to the required levels. The development of these counters has gone hand-in-hand with the investigation and calibration of the ion extraction process, seeded ultra-filtration, at the low levels required.

The author was also a CASE award holder with Thorn EMI, now renamed Electron Tubes Ltd., and development work on large area PMT's, in terms of efficiency, timing response and noise, is also included. Of particular note is an investigation of pulses arriving at the anode around 40 ns late.

Acknowledgements

I am perhaps unique in having been supervised by three people, and would like to thank, in chronological order: Rob Boardman, who gave me a good grounding in PMT's and their intricacies, before moving to ETL and contributing to the work in chapters 6 and 7; Martin Moorhead, who pulled me into the water group, and whose boundless enthusiasm usually overcame my black despair; and last, Neil Tanner, who can usually be relied on to have three new ideas per week, and be annoyingly correct on most of them.

Thanks are also extended to the rest of the SNO group at Oxford, who have helped in more ways than I would care to mention: Peter Trent, a dogged pragmatist, and who has read this thesis unbidden when others could not find the time (they know who they are); Nick Jelley, whose contributions from outside the closed world of the water group have proved invaluable; Barrie Knox, for day to day knowledge on chemical engineering; Will Locke, our resident chemist, for his patient corrections to my somewhat dusty A-level chemistry; Mike Lay, for his guiding hand through the wastelands of Fortran and PAW; Dave Wark, for his unceasing American-ness and fantastic knowledge of neutrino physics; my contemporary, Steve Brice, who has always been ready to explain himself, and in turn been ready to listen; George Doucas, who has always remained calm in sometimes stormy waters; and John Barton for his perspicacity and reliability in terms of radionuclide counting.

Thanks must also be extended to the Head of Department, Prof. Roger Cashmore, both for his support of, and interest in the SNO project. The lecture courses of Norman Booth and Wade Allison have been a great help over the three years, and a particular debt of thanks must be paid to the workshop staff. Among the many I could mention, many thanks are extended to Phil Evans, who built the β-α counters and came up with numerous mechanical evolutions of the main design, Peter Groves and Dave Boulter, who built the main elution rig from a series of somewhat sketchy drawings, and Mick Williams and Colin Graham, for constant advice and assistance. In line with departmental policy, it is also beholden on me to thank the Tea-Ladies, who have made the daily round slightly more lively.

The willingness and helpfulness of the staff at E.T.L. is acknowledged, with many thanks to Ron McAlpine for advice on the breakdown phenomenon, and his interest when other commitments were sometimes pressing.

My thanks are also extended to SERC, and latterly PPARC, for their funding of the research.

Finally, I would like to thank my friends and family, who have regularly brought me back down to earth when such an operation was necessary.

Contents

Chapter 1

The Solar Neutrino Problem

Neutrino physics has become more and more interesting since the observation in 1956 by Reines and Cowan of anti-neutrino scattering. With their original purpose of conserving energy in beta decay fully justified by the weak interaction theory, the properties of the neutrinos themselves became a source of great interest both to the particle and astrophysical communities. This is indicated by the large number of experiments that are currently detecting neutrinos, and the range of energies at which they are being studied.

The pervasiveness of the weak interaction has led to the neutrino appearing in many areas which would at first seem unlikely. One of the most well known is a problem in stellar evolution which has come to be known as the solar neutrino problem, or SNP. This thesis will begin by addressing the problem, and then go on to one of the experiments that will hopefully resolve the discrepancy between theory and measurement. We will begin by describing the theoretical models of the evolution of stars.

1.1 Stellar evolution

It is well known that the main energy generation mechanism within the solar core is the so-called 'pp-cycle'. A schematic of the scheme is shown in figure 1.1. The other well known process, the CNO cycle, contributes only around 1.5% of the sun's energy, and is usually discounted from discussion. The fusion process can proceed by many paths, but the final production of helium nuclei from four protons is the same for all. Thus the cycle boils down to a basic fusion process, also known as a termination:

$$4p + 2e^- \rightarrow \alpha + 2\nu_e + 4\gamma \qquad \mathrm{Q} = 26.7 \ \mathrm{MeV}$$

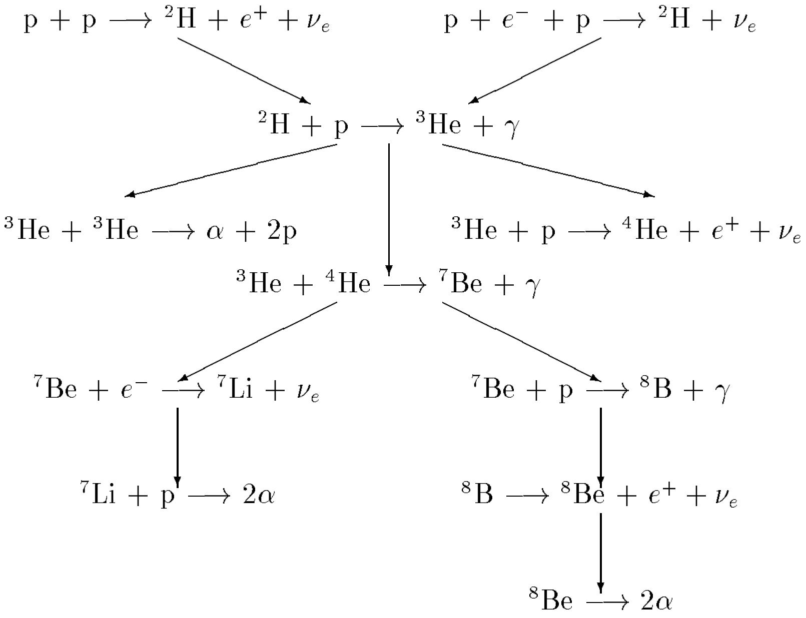

Reaction	Acronym	Termination (%)
$p + p \longrightarrow {}^2H + e^+ + \nu_e$	pp	100
$p + e^- + p \longrightarrow {}^2H + \nu_e$	pep	0.4
${}^2H + p \longrightarrow {}^3He + \gamma$		100
${}^3He + {}^3He \longrightarrow \alpha + 2p$		85
${}^3He + {}^4He \longrightarrow {}^7Be + \gamma$		15
${}^7Be + e^- \longrightarrow {}^7Li + \nu_e$	^{7}Be	15
${}^7Li + p \longrightarrow 2\alpha$		15
${}^7Be + p \longrightarrow {}^8B + \gamma$		0.02
${}^8B \longrightarrow {}^8Be + e^+ + \nu_e$	^{8}B	0.02
${}^8Be \longrightarrow 2\alpha$		0.02
${}^3He + p \longrightarrow {}^4He + e^+ + \nu_e$	hep	0.00002

Figure 1.1: The 'pp' cycle. The termination percentage is the fraction of terminations that include the particular reaction. Notice that those reactions which produce neutrinos are labelled with an acronym which will be used in the rest of the thesis.

The energy generated from these fusions diffuses away from the core and eventually provides sunlight. The transport of the energy proceeds via two distinct processes: photon diffusion, which dominates in the core of the sun; and convection, where the heated matter 'rises' from the centre of the sun: this latter process only becomes important in the last quarter of the sun's radius. Thus the observable light from the sun originates only in the outermost, convective part of the sun, and little information about the solar core can be obtained from optical observations.

However, when we go back to the pp cycle, we observe that neutrinos are also generated by the fusion process. As they interact only weakly, these neutrinos can pass through the star without scattering, and hence carry information about the fusion processes within the core. As an order of magnitude estimate of the numbers of these neutrinos, we can say:

$$2E_{Mean} = 26.7 - 2\overline{E}_\nu \approx 26 MeV$$

where $\overline{E}_\nu$ is the mean neutrino energy, and E_{Mean} is the energy contributed to the solar luminosity per neutrino generated. So:

$$\Phi_\nu = \frac{1}{E_{Mean}} \frac{L_\odot}{4\pi D^2} = 6.5 \times 10^{10} cm^{-2} s^{-1}$$

where Φ_ν is the neutrino flux and D is the Sun-Earth distance. This estimate however is of little practical use, as we have no direct measurement of the energy spectrum of these neutrinos. In order to obtain such information it is necessary to model the behaviour within the sun and deduce the rates of the reactions shown in the pp cycle. The features of these Standard Solar Models (SSMs) are summarised below.

1.2 The standard solar model

The first assumption that is advanced when constructing a solar model is that the system is in hydrostatic equilibrium, implying that the internal pressure due to the motions of the ions and electrons plus the radiation pressure matches the gravitational pull of the solar material, and such a system must have a pressure profile which goes as:

$$\frac{dP}{dr} = -\frac{GM(r)\rho(r)}{r^2}$$

Here, P is the pressure on a shell at radius r, with the local density at this radius given be $\rho(r)$, and the mass contained within the shell being M. Implicit in this equation is the assumption that the stellar material is not rotating about its centre. We know this to be untrue of course, but it has been shown that adding the sort of rotational speeds expected within the stellar core would not affect the final results.

The equation of state of the matter within the stellar core also gives a relation between the pressure, density and temperature at each point throughout the material.

As most of the material is completely ionised, the equation of state for an ideal gas is usually assumed, with corrections for degeneracy and incomplete ionisation being added.

As said before, the energy transport within the core is almost purely via photon diffusion, with convective cells occurring at around $0.75R_{\odot}$. The energy generation comes from the pp chain, with the rates of these reaction calculated using the energy dependant cross sections. These also have a Gamow penetration factor applied to them to take account of the tunneling required by the similar charge nuclei. The factors are important in that they produce the well known temperature dependencies of the neutrino rates.

The final assumptions made are of the initial elemental abundances in the primeval solar material. These abundances evolve through time as the hydrogen is burned, but the initial ratio of hydrogen to helium and the presence of heavier elements are further input parameters to these models. The actual modelling stage involves the iteration in time of the solar model up to the present age of the sun, which is usually calculated from radioactive aging of meteorites. The age obtained, 4.6×10^9 years, is generally accepted by all the different solar modellers.

The final stage in any modelling process is the comparison with observable quantities, most especially the solar radius and solar luminosity, the mass having already been fixed at the start of the model. The modelling process is repeated with differing input parameters until a satisfactory agreement with present data is found. It is at this point that the individual branching ratios of the different terminations, and hence the neutrino fluxes, can be extracted and compared with observation so as to confirm the model's accuracy within the stellar core.

In table 1.1 we list the results of this process for some of the better known groups in the field, along with the error on the theoretical estimates. In the work of Bahcall, the errors in the articles themselves are 3 sigma errors, but in the table have been divided by three for comparison with other groups.

The different authors listed in the table are: BU - Bahcall and Ulrich [2], TCCCD - Turck-Chieze et al [5], SBF - Sackmann, Boothroyd and Fowler [7], BP(92) - Bahcall and Pinsonneault(92) [3] and TCL(93) - Turck-Chieze and Lopez [6]. The final line, BP(95) [4], is the latest iteration by Bahcall and Pinsonneault and includes effects from the diffusion of the helium and heavy metals within the core of the sun. BP(95) figures

Acronym	pp	pep	^{7}Be	^{8}B
Group				
BU(88)	6.0±0.03	1.4±0.03	4.7±0.2	5.8±0.7
TCCCD(88)	5.98±0.2	1.3	4.18	3.8±1.1
SBF(90)	6.00	1.29	4.23	5.80
BP(92)	6.00±0.04	1.43±0.02	4.89±0.29	5.69±0.82
TCL(93)	6.03	1.39	4.22	4.43
BP(95)	5.91±0.06	1.40±0.03	5.15±0.36	6.67±1.13
Flux factor	10^{10}	10^8	10^9	10^6

Table 1.1: A selection of the standard solar model fluxes from different groups. The group acronyms are listed in the text, while the figures in brackets are the years of publication. The four fluxes shown are the readily observable fluxes for existing and proposed detectors. The fluxes are those predicted per cm^2 per second at the Earth when multiplied by the appropriate flux factor.

are assumed to be the most accurate numbers currently available, and are used in the rest of this chapter[1].

Although the similarity between different calculations for some of the fluxes is obvious, the most striking differences are in the ^{7}Be and ^{8}B fluxes. These variations are an indication of the major sources of uncertainty in the standard solar model:

- The cross sections for the scattering processes in the pp cycle are measured at a few hundred keV, while the core temperature of the sun ($\approx 1.5 \times 10^7$ K) sets the energy scale at the order of 1 keV. The cross sections are extrapolated from these higher energies where the Gamow penetration factor is not dominant. These problems are particularly prevalent in the S_{34} (^{3}He(^{4}He$,\gamma)^7$Be) cross section and the S_{17} (^{7}Be$(p,\gamma)^8$B) cross section, which affect the ^{7}Be and ^{8}B fluxes directly.

- The calculation of the radiative opacity for a given mixture of elements is a very difficult process, although there has been considerable progress thanks to the OPAL group at Lawrence Livermore[8].

- The elemental abundances of the primordial solar material are also a subject of considerable discussion, with differences between the meteoritic content and the observed atomic lines in the solar photosphere only recently being resolved with better experimental determinations.

[1] The reader is directed to this latter paper for a good example of the current state of the standard solar model and numerous further references.

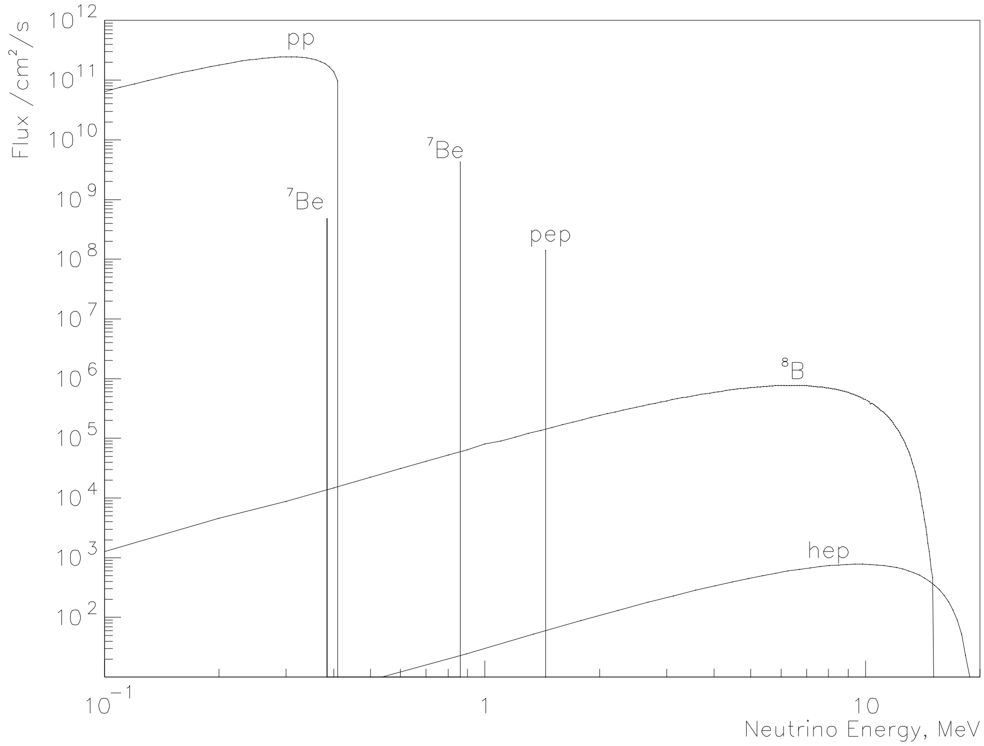

Figure 1.2: The energy spectrum of solar neutrinos. The different distributions are normalised to the fluxes given in [2]. The line sources are in units of number per cm^2 per second, while the continuum sources are in units of number per cm^2 per second per MeV.

- Finally, there has been much work of late on more precise measurements of the solar luminosity and the solar age. These were long assumed to be fairly well known, but recent measurements have reduced the errors associated with these parameters.

If the flux calculations are applied to the energy spectra expected for the neutrinos, we obtain the familiar spectrum of figure 1.2.

It should be noted that figure 1.2 is a log-log plot, and hence the overwhelming majority of the neutrinos are produced at the low energy end from the pp reaction, with the endpoint for this process being at 0.4 MeV. The line sources associated with ^{7}Be and pep have a width due to thermal motion of the order of 1 keV. The boron-8 spectrum extends up to around 14 MeV, although it is much smaller in absolute terms than the pp flux. The very highest energy neutrinos come from the hep (the $^3He(p, e^+\nu_e)^4He$ reaction) flux which has an endpoint of 19.8 MeV.

Of course, the standard solar models produce many more results than these neutrinos fluxes. The distributions of temperature, energy production and composition are also available, and the results for the outermost parts of the model, the convective zone, conform quantitatively with results from helioseismology – the study of small acoustic oscillations in the sun. However, four experiments have now measured solar neutrino fluxes, and all four have answers that disagree at the many sigma level with these predictions. This constitutes the solar neutrino problem, and we will now consider these experiments individually.

1.3 The Homestake experiment.

The first experiment that measured solar neutrinos was the Homestake chlorine solar neutrino detector, also referred to as Homestake and the Chlorine experiment. It is a radiochemical detector consisting of 615 tonnes of perchloroethylene (C_2Cl_4) in a tank at a depth of 4850 ft in the Homestake gold mine in South Dakota. The chlorine contained in the tank is sensitive to neutrinos via the reaction:

$$\nu_e + ^{37}\mathrm{Cl} \rightarrow ^{37}\mathrm{Ar} + e^- \qquad \mathrm{Q} = \text{-}0.81\ \mathrm{MeV}$$

Thus the production of argon is directly proportional to the integrated solar neutrino flux above this value. This energy range thus includes both the ^7Be and ^8B fluxes, as well as small contributions from the CNO cycle neutrinos. The ^{37}Cl makes up around 133 tons of the total tank weight.

As well as the C_2Cl_4 in the tank, 5% of the tank volume is made up of helium gas. This is used to flush out the argon after a run. At the end of the previous extraction, around 0.1cc of isotopically pure argon (either ^{36}Ar or ^{38}Ar) is added to the tank as a carrier. At the end of a run, typically lasting two months, the helium is bubbled through the tank, bringing with it the carrier argon and any neutrino induced argon. The helium is pumped through an activated carbon trap which collects the argon.

After the whole tank has been processed, the argon is removed from the charcoal and purified by gas chromatography before being mixed with tritium free methane. This argon methane mixture is then pumped into small proportional chambers which are operated for around 1 year so as to fit the decays to the argon half life plus a flat background. After this, the proportional counter is vented and the amount of argon compared with the amount added as the carrier. This gives the mean extraction efficiency for the gas. The Homestake collaboration also quotes a processing efficiency which measures the amount of argon in the charcoal trap and compares it with the amount inside the proportional chamber. Both the extraction and processing efficiencies are around 95%.

Since the chlorine experiment has shown a considerable deficit when compared with the SSM, there have been many experimental checks on the radiochemistry involved. Perhaps the most useful test was when ^{36}Cl was added to the detector. This decays via a beta to ^{36}Ar with approximately the same nuclear recoil as is produced in the neutrino capture reaction. Thus full recovery of this ^{36}Ar would imply that argon atoms produced with small recoil energies would not be trapped in the parent molecule. The most recent results[12] from these tests have shown a recovery of 102±4%, thus indicating that the neutrino-induced ^{37}Ar should not be trapped in the remnant molecule.

The background for the detector can be divided into two separate areas: firstly, the production of spurious ^{37}Ar atoms within the tank itself; and secondly, backgrounds within the proportional chambers used to count the gas. The first background comes mainly from cosmic ray muons producing spallation protons within the detector volume. These can capture on the the chlorine in $^{37}\mathrm{Cl}(p,n)^{37}\mathrm{Ar}$ reactions. The muon flux was measured using smaller C_2Cl_4 tanks nearer the surface, and the rate was extrapolated to the detector depth. Another check was achieved using a potassium loaded detector at the same depth as the tank, with both methods yielding consistent results for the induced ^{37}Ar background. Further checks on the background were made by measuring the neutron flux in the cavity and combining this with the effective cross section for (n,p) reactions followed by the $^{37}\mathrm{Cl}(p,n)^{37}\mathrm{Ar}$ reaction to give a neutron induced background. The cumulative background is given as 0.08 argon atoms produced per day.

The proportional counters look for the electron capture decay of the ^{37}Ar which produces 2.82 keV (capture from K shell) in 90% of the decays: 90% of these decays produce detectable Auger electrons (as opposed to undetectable soft X-rays). The detection efficiency is found using known spikes of ^{37}Ar, while the calibration of the devices is periodically checked with an external ^{56}Fe source. The main background events in the counters are through-going electrons from beta decay, but many of these

Neutrino source	σ_{Cl} $10^{-46}\mathrm{cm}^{-2}$	BP(95) SNU	TCL(93) SNU	BPML(93) SNU
pp	0	–	–	–
pep	16	0.22	0.22	0.22
^{7}Be	2.38	1.24	1.10	1.06
^{8}B	1.06×10^{4}	7.36	4.63	5.73
^{13}N	1.66	0.11	0.063	0.095
^{15}O	6.61	0.37	0.21	0.318
Total	–	9.3±1.4	6.36	7.43±2.7
Measured	2.55±0.25 SNU			

Table 1.2: The predicted SNU capture rates for the Homestake chlorine experiment. The three groups are: BP(95)[4], TCL(93)[6] and BPML(93)[9]. The cross sections are taken from the latter paper. The CNO cycle neutrinos are included as they make a small but significant contribution to the total theoretical rate. The systematic and statistical errors have been combined quadratically in the measured rate.

are eliminated by cuts on the pulse rise-time and height.

Results from Homestake.

The detector has now been running for around 25 years, and the measured neutrino flux, after a background subtraction, is given as[12]:

$$\Sigma(\Phi_i \sigma_i) = 2.55 \pm 0.17 \pm 0.18 \quad \mathrm{SNU}$$

where the sum is over the neutrino fluxes and cross sections, the first error is statistical and the second systematic. The rate is given in SNU, or Solar Neutrino Units. This is a convenient measure, with one SNU being equivalent to 1 event per 10^{36} target atoms per second. In the Homestake detector therefore, 5.35 SNU is equivalent to the production of one ^{37}Ar atom per day. The background to the measurement from the sources mentioned above is said to be 0.43 SNU. The comparison of this figure with the theoretical predictions is given in table 1.2.

As can be seen, there is a major discrepancy at around the 3σ level between the theory and experiment. Considering the make-up of the theoretical rate, it can be seen that by far the largest contribution comes from the ^{8}B termination, which is the most difficult to predict. Further work by Bahcall after the publication of the discrepancy led to the determination of the temperature dependencies of the various neutrino fluxes.

As the temperature affects both the luminosity calculations and the barrier-penetration factors in the cross sections, it was felt necessary to see if the data could be 'fitted' by changing the core temperature slightly. The temperature dependencies were found by running 1000 similar solar models with slightly differing input values, and plotting the neutrino rates against core temperature. From these plots, approximate power distributions were drawn up as follows:

$$\phi(^8\mathrm{B}) \propto T_c^{18}$$

$$\phi(^7\mathrm{Be}) \propto T_c^{8}$$

$$\phi(pp) \propto T_c^{-1.2}$$

Thus, it can be seen that were the core temperature to drop by around 10%, then the discrepancy is removed[2]. Of course, some mechanism has to be advanced to bring about this drop in temperature, and some of the more well known are discussed by Bahcall in [1] under the heading of non-standard solar models. However, with the advent of the results from Kamiokande, it was realised that such hypotheses were in trouble.

1.4 The Kamiokande experiment.

The Kamiokande solar neutrino experiment[13] was a large light water Čerenkov detector at the Kamioka mine in Japan. The mine cavity is situated at a depth of 1km of rock, this being equivalent to 2.7km of water. The detector consisted of a cylindrical tank containing 2140 tonnes of water, this volume being observed by 948 20 inch photomultiplier tubes (PMTs). Outside the 2.14 kilo-tonne volume was a further 1.8 kilo-tonnes of light water viewed by 123 PMTs to act as an anti-coincidence counter. The two volumes were optically separated by a black plastic liner. This detector is referred to in the literature as Kamiokande II, as it was an upgraded version of the original proton decay experiment, Kamiokande I. There was a further upgrade in 1990 to Kamiokande III, when the detector was drained, and dead PMT's were removed. Another electronic upgrade was also implemented at this point. This process is detailed in the above reference.

[2]The variation of temperature has subsequently been shown to be a sufficient condition for varying the fluxes without changing any of the other input parameters, see [10] for the basis of this finding.

Although the original Kamiokande experiment was built to look for proton decay (hence the name – Kamiokande is an acronym for the Kamioka Nucleon Decay Experiment), it was soon realised that it could also be used for the detection of solar neutrinos by their elastic scattering from electrons in the water tank. The recoil electron would then be detected by its production of Čerenkov light. The maximum energy of these electrons would be around 15 MeV, whereas the detector had been optimised for electrons of the order of half the proton rest-mass. The search for solar neutrinos thus began with an accurate calibration and background check for these relatively low energy scales.

The energy response of the detector was calibrated in a number of ways. The first source utilized a californium 252 source at the centre of a water-filled nickel can. The neutrons emitted by the ^{252}Cf were moderated by the water and thermally captured on the nickel producing 9 MeV gamma rays via Ni(n,γ)Ni*. As the ^{252}Cf also produces gamma rays of much smaller energies, the source was also run without a nickel can, and the two data sets were subtracted to give the spectrum due to the 9 MeV gammas alone. The Monte Carlo of the detector was tuned so as to reproduce the data, and the source was then moved around within the detector volume, and the Monte Carlo prediction was shown to be within 2% of the data. The source was also run periodically over the time of the experiment so as to check the stability of the detector with time, and again the data was within 2% of the expectation.

For higher energies, the decay spectrum of stopped muons in the tank was also used to calibrate. The fit to the Michel spectrum was good to the 1% level, although the energy range sampled by this source was mostly well above the solar neutrino-induced events. As a final check, the energy spectrum of spallation products was used as another source. The time fit of the muon-induced spallation products showed that the main isotopes produced were ^{12}B and ^{12}N in the ratio 0.85:0.15. Monte Carlo events of this nature were generated and then compared with observation. Again, an agreement to within 2% was seen. A timing calibration was also achieved using a pulsed nitrogen laser, the light from which was sent along an optical fibre to a diffusing ball at the centre of the detector. Finally, the angular resolution was checked by placing a collimator on the nickel gamma source, and comparing the Monte Carlo data with the fitted gamma direction. From all these checks, the final calibration was:

$$\frac{\sigma(E_e)}{E_e} = \frac{19\%}{\sqrt{E_e/10MeV}}$$

$$\frac{\sigma(\theta_e)}{\theta_e} = 28^\circ(E_e = 10MeV)$$

$$\frac{\sigma(r)}{r} = 1.7\text{m}(E_e = 10MeV)$$

where E_e is the scattered electron's kinetic energy. The third line shows the position resolution for the event. This was important in the discrimination of background events.

The main background in the detector was found to be due to ^{222}Rn in the water. This decays via a series of isotopes (of which much more later), with the most problematical being the beta decay of ^{214}Bi. Although the endpoint energy of this decay is below 5 MeV, the poor energy resolution of the detector at these low energies meant that the removal of the radionuclide was paramount. With a sophisticated water purification system, the level of radon was brought down to around 0.04 Bq per tonne. This activity allowed the threshold for recoil electrons to be set at around 9 MeV. Further radioactive contamination of the surrounding rock, leading to gammas producing Compton scattered electrons in the detector volume, was reduced firstly by the anti-counter and subsequently by the adoption of a fiducial volume of 680 m^3 within the 2.14 kilo-tonne tank.

The final background considered was that due to the spallation products from cosmic rays. As stated above, this source of background was actually used to some extent as a calibration source, and by vetoing the detector for 50 seconds after a cosmic ray event, the background was reduced to an acceptable level. Perhaps the most obvious background rejection technique was to use the directionality of the neutrino-electron scattering. At order 10 MeV, the electron's direction is strongly forward peaked away from the sun. As all the background sources are isotropic with respect to the solar direction, a further reduction in background of around 10 could be achieved. Such a plot of angular distribution produces the well known plot of figure 1.3.

The significance of this plot is sometimes lost: this is the first direct observation of solar neutrinos. However, when comparing the flux measured with the predictions from the standard models, it was found that there was yet another discrepancy. The ratio of predicted to measured is[14]:

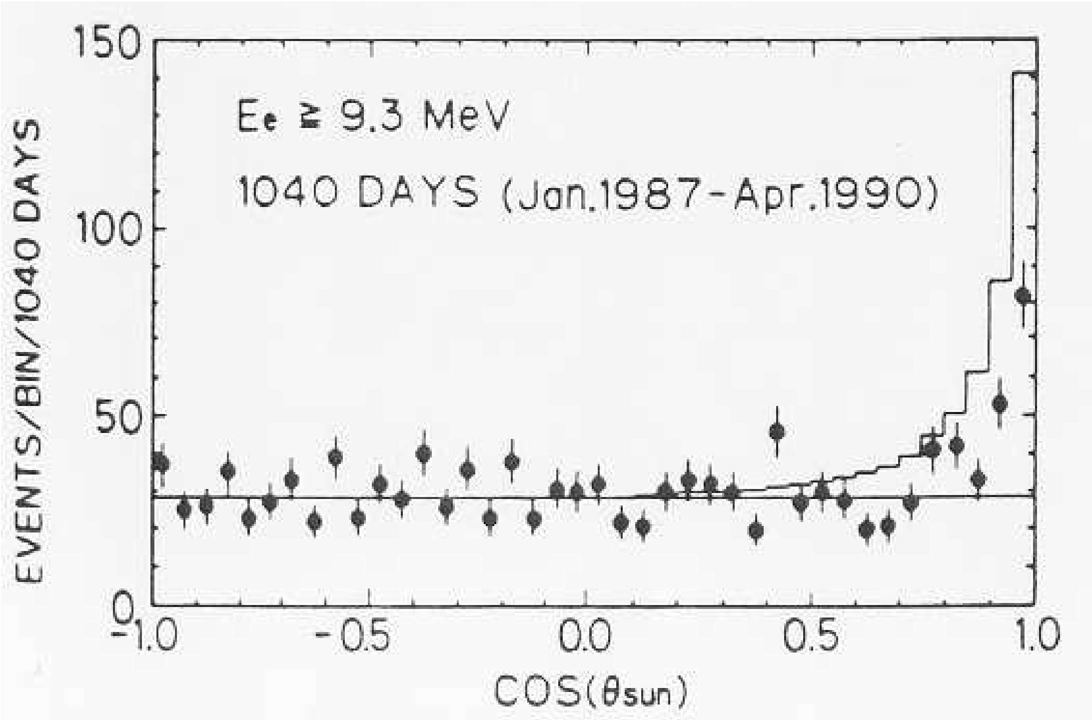

Figure 1.3: The angular distribution of electrons in the Kamioka detector above 7.5 MeV measured energy. The peak away form the sun demonstrates the observation of solar neutrinos. The solid line is the expected distribution for the solar model of BP(92)[3] with a flat background.

$$\frac{Observed}{Predicited} = 0.51 \pm 0.04 \pm 0.06$$

where the first uncertainty is the statistical error, and the second is the quadratic combination of the systematic error on the measurement and the error quoted for the SSM flux, as taken from BP(95)[4].

Thus, the observed flux is only around half that expected. Just about the only assumption when this comparison is made is that the energy spectrum of the neutrinos at the detector is the same as the normal ^{8}B decay spectrum. Bahcall discussed this point in [11], and came to the conclusion that the major correction of the shape would be due to the neutrinos moving out of the gravitational potential well of the sun. This would change the spectrum by around 1 part in 10^5.

It was at this point that the solar neutrino problem became much more serious. It was argued up until the Kamioka result that there was something wrong either with the Homestake experiment or with the standard solar model. However the fact that the theoretical prediction for Kamiokande has a higher dependency on the temperature than that for Homestake, while Homestake sees a larger deficit than Kamiokande tells us that varying the temperature in the core of the sun cannot help us. This situation is further reviewed later.

These two experiments were still viewed with some suspicion: the Homestake experiment has a fairly low rate, while the Kamiokande experiment was never really designed for the observation of solar neutrinos. Also, both experiments had fairly high thresholds, and hence could not see the largest flux of neutrinos produced by the pp reaction, the rate of which can be inferred from the solar luminosity independent of the solar model. It was from this basis that the gallium experiments began.

1.5 The Gallium Experiments.

The attractiveness of gallium in a solar neutrino detector is the very low Q value associated with its conversion into germanium:

$$^{71}\mathrm{Ga} + \nu_e \rightarrow ^{71}\mathrm{Ge} + e^- \qquad \mathrm{E}_{Thres} = 0.2332\ \mathrm{MeV}$$

This value of 233 keV is important as it allows pp neutrinos to be detected. This flux is by far the largest in the standard solar model, and hence good statistics can be obtained with comparatively small detectors. Also, the neutrino-induced ^{71}Ge decays via electron capture with a half-life of 11.43 days, thus allowing the gallium to be exposed for a few weeks before the germanium is recovered. Of course, the gallium is also sensitive to all the other neutrino fluxes, and with a cross section that rises steeply with energy, around one third of the predicted flux comes about from the ^{7}Be and ^{8}B neutrinos which had already been measured by Homestake and Kamiokande. These expected capture rates are shown for different models in table 1.3.

The large dependence on the pp rate should be noted as the majority of the sun's luminosity is generated by the pp reaction, and hence a strict lower limit on the number of neutrinos can be set using the observed luminosity. This is known as the minimal standard solar model: essentially, the cross-section for $^3\mathrm{He}(\alpha,\gamma)^7\mathrm{Be}$ is set to zero, with the only energy generation mechanism being the pp-I termination:

$$p + p \rightarrow ^2\mathrm{H} + e^+ + \nu_e$$

Neutrino source	σ_{Ga} $10^{-46}\mathrm{cm}^{-2}$	BP(95) SNU	TCL(93) SNU	BPML(93) SNU
pp	11.8	69.7	71.1	70.0
pep	214	3.0	2.99	3.0
^{7}Be	7.27	37.7	30.9	32.5
^{8}B	2.43×10^4	16.1	10.77	13.1
^{13}N	61.8	3.8	2.36	3.53
^{15}O	116	6.3	3.66	5.58
Total	–	137±8	122.5	127.7±20

Table 1.3: The predicted SNU capture rates for the gallium experiments. The three groups are: BP(95)[4], TCL(93)[6] and BPML(93)[9]. The cross sections are taken from the latter paper. The CNO cycle neutrinos are included as they make a small but significant contribution to the total theoretical rate. Note that due to the definition of the SNU, the predicted rates for both SAGE and GALLEX are the same.

$$p +^2 \mathrm{H} \rightarrow^3 \mathrm{He} + \gamma$$

$$^3\mathrm{He} +^3 \mathrm{He} \rightarrow^4 \mathrm{He} + p + p$$

This means that only pp (and to a small extent pep) neutrinos are produced. Given the solar luminosity, it is possible to calculate that the gallium experiments must see at least 77 SNU. This is known as the minimal model as it is the minimum number of neutrinos that ensure that the sun still shines. Of course, the fact that Kamiokande has seen neutrinos from the ^{8}B spectrum would imply that this model is immediately wrong, but as said above, this is only used as an artificial lower limit below which new neutrino physics must be invoked.

We will now detail the two gallium experiments, starting with SAGE:

1.5.1 The SAGE experiment.

SAGE is an acronym for the Soviet[3]-American Gallium Experiment[17] and consists of 8 Teflon lined vessels, each of which contain 7 tonnes of pure gallium metal. The vessels are heated to 30°C so as to keep the gallium in liquid form. The site of the experiment is a cavity under a mountain in the Northern Caucasus mountains of Russia, with an effective depth of 4.72 km water equivalent. After an exposure time of around 4 weeks, the liquid metal is treated with various chemicals to achieve a solvent/solvent

[3]Although the experiment is now referred to as the Russian–American Gallium Experiment, a proposal to change the name to RAGE was thought inappropriate.

extraction of the germanium. The germanium is then further concentrated and finally used in the synthesis of the gas germane, GeH_4, which is mixed with an amount of xenon and placed in a miniature proportional chamber. Before the start of the next exposure, 0.1 mg of natural germanium (i.e. containing no ^{71}Ge) is added to each of the tanks to act as a carrier for the next extraction. The extraction efficiency of the process is measured by checking the amount of the carrier that is recovered, and usually lies at around 80%.

The counting is carried out by the proportional chambers, with the ^{71}Ge decaying via electron capture with the subsequent production of detectable Auger electrons. The two largest energy captures, K and L shell, are both available for counting, although the detection of the L capture only become feasible after an electronics upgrade carried out in the summer of 1992 added pulse shape discrimination to the signals from the counters. These counters are calibrated over the energy range in question with a number of external sources.

The backgrounds for the experiment are split into two separate sections, much as with the Homestake experiment: those associated with the production of ^{71}Ge within the tank from sources other than the neutrinos, and the backgrounds associated with the proportional chambers. The first source is associated with the production of free protons which then scatter in the $^{71}\mathrm{Ga}(p,n)^{71}\mathrm{Ge}$ reaction. The protons come from three distinct areas: external neutrons, cosmic ray spallation and internal radioactivity. The contributions of these sources have been calculated from measurements within the SAGE cavity, and contribute only around 1% of the expected rate from neutrinos. The background in the counters is well controlled by tight cuts on both the energy and risetime of the pulses, as well as the fitting of the counts from the proportional chambers to an exponential (with the 11.4 day half life) plus flat background.

Other tests of the detector included many of the processes carried out in the Chlorine experiment, with perhaps the most obvious being checks on whether the recoil of the ^{71}Ge could form some unextractable complex molecule. By neutron activation of the stable isotopes of gallium and their subsequent beta decay to stable isotopes of germanium, it was shown that there should be no large effect. Also, a known spike of ^{71}Ge was added to one of the tanks and the extraction of this radioactive germanium was compared to the extraction of the stable germanium carrier with the result that the two germanium isotopes seemed to follow each other almost exactly.

The experiment has now been running for 5 years, with various interruptions for upgrading and calibration, and the combined result of the exposures so far gives a flux of[16]:

$$Observed = 74 \pm^{13}_{12} \pm^{5}_{7} \qquad \text{SNU}$$

where the first error is statistical, and the second systematic. This result is around 3 sigma away from the latest Bahcall flux calculations. We now go onto the other gallium experiment, GALLEX.

1.5.2 The GALLEX experiment.

The GALLEX[17] collaboration has built and is using a gallium radiochemical detector consisting of a 70m^3 tank which contains an 8 Molar aqueous solution of gallium chloride $GaCl_3$, this volume being further acidified with HCl: the total mass of ^{71}Ga in the tank comes to 12 tonnes. The experiment is located in the Gran Sasso facility, which has an overburden of 3.3 km water equivalent, hence reducing cosmic ray backgrounds. The extractions began in 1991, and around 3 and a half years of data have now been taken, with around 150 events having been detected.

At the start of each exposure, a stable germanium carrier (either ^{72}Ge, ^{74}Ge or ^{76}Ge) is added, with the particular isotope being changed between runs so as to monitor the amount of carry–over from one run to the next. The germanium is relatively easy to extract as germanium tetrachloride is very volatile, and so is easily swept out of the tank using a flushing gas stream of nitrogen. Measurements of the amount of carrier recovered show that the extraction efficiency is of order 80%. The germanium tetrachloride is collected in a water bath and then processed further to produce the counting gas Germane, GeH_4. As in SAGE, this is mixed with xenon and placed in miniature proportional counters. Unlike SAGE, the counters are low enough in rate to measure the L shell electron capture of ^{71}Ge and this increases the counting efficiency by around a factor of two to around 70%. Again, the rate of events from these counters is fitted to the exponential decay of the ^{71}Ge plus a flat background.

Unsurprisingly, the backgrounds in the GALLEX detector are much the same as those for the SAGE experiment discussed above, and tests of the extractability of the recoiling ^{71}Ge atoms have been checked in much the same fashion, and indeed with the same result. Blank runs, where the time of exposure is kept as small as possible, carried out by GALLEX are perhaps the best check of any backgrounds associated with the instrumentation, and have yielded a result consistent with zero.

The result for the data taking is[18]:

$$Observed = 79 \pm 10 \pm 6 \quad \text{SNU}$$

Again, the first error is statistical and the second systematic. The agreement with the SAGE result is obvious, and indeed a good confirmation, as it is difficult to imagine a situation in which both experiments could be wrong other than some miscalculation of the ^{71}Ge production cross-section. Having said this, the best confirmation that both experiments could have would be a direct calibration with a strong neutrino source, and in the last year or two, both have used such a source.

SAGE and GALLEX neutrino calibrations.

The calibration of the two gallium experiments has been the final confirmation of the solar neutrino problem. As with many radiochemical experiments, there are sufficient degrees of uncertainty so as to make absolute calibration vital. With both gallium experiments reporting numbers that were only just compatible with the minimal standard solar model, and both reporting that their extraction and counting efficiencies were well known, the only other area of doubt was the neutrino cross section associated with the inverse beta-decay $^{71}\mathrm{Ga}(\nu_e, e^-)^{71}\mathrm{Ge}$.

The isotope used to calibrate both detectors was ^{51}Cr. This radioisotope has a half-life of 27.7 days, and has therefore to be produced by neutron activation of the stable ^{50}Cr in a nuclear reactor. Its primary advantage is that decays via electron capture with a well known neutrino spectrum, 90% of the decays giving neutrinos at 746 keV, and 10% at 426 keV. This should be compared with the endpoint of the pp neutrino flux at 420 keV and the ^{7}Be lines at 383 and 861 keV.

Both experiments have now reported on the calibration runs, with the results shown below:

$$GALLEX : \frac{Obs.}{Theo.} = 1.04 \pm 0.12$$

$$SAGE : \frac{Obs.}{Theo.} = 0.93 \pm 0.17$$

where the GALLEX result is taken from [19] and the SAGE result from [20]. These numbers include all the extraction procedures and include any of the unspecified chemical effects that were thought to be the major source of uncertainty. Hence, with both calibrations consistent with unity, we have a strong confirmation of the measurements of the fluxes in both experiments, and indeed a strong confirmation of the solar neutrino problem.

1.6 Confronting the Solar Neutrino Problem.

We now have four experiments, all of which have seen a deficit when compared to the standard solar model. The magnitudes of these deficits vary, as does the sensitivity of various experiments to various components of the solar neutrino flux. In order to reconcile the differences, it is necessary to consider the most uncertain parts of the model, and see if we can change them within acceptable bounds so as to fit the data. We will begin with those solutions which essentially lower the core temperature of the sun. We know that:

$$\Phi(^7\mathrm{Be}) \propto T_c^8$$

$$\Phi(^8\mathrm{B}) \propto T_c^{18}$$

When we consider the measurements from Kamioka and Homestake with respect to the predictions:

$$Ratio, R = \frac{Obs.}{Pred.}$$

$$R_{Homestake} = 0.27 \pm 0.05$$

$$R_{Kamiokande} = 0.51 \pm 0.07$$

we can see that temperature variation is not going to help us: Homestake sees neutrinos both from the ^{7}Be and ^{8}B fluxes, and hence has a smaller temperature dependance than Kamiokande. Hence to fit the two simultaneously with a common temperature is impossible. It is at this point that many authors display a prejudice for saying that Homestake is wrong, although little or no evidence is forwarded other than a feeling that radiochemical experiments are in general much more difficult than the somewhat simpler water Čerenkov detectors.

More suggestively, it is well known that the S_{17} cross section (ie proton capture on beryllium) is perhaps the most uncertain measured quantity used in the solar model: recent results for its measurement are reviewed in [21]. If we allow it to range freely so as to fit the Kamioka result, we can then predict the number of ^{8}B neutrinos that should be seen by Homestake. Of course, if we reduce this cross-section, then it will increase the number of ^{7}Be neutrinos produced by electron capture on the beryllium, but as the relative rates of the two are significantly different in the solar core this is a negligeable effect.

So, using the Kamiokande result to predict the ^{8}B flux:

$$\phi_{Kam} = 2.89 \pm 0.22 \pm 0.35 cm^{-2} s^{-1}$$

So that:

$$\phi_{Kam} \sigma_{Cl} = 3.21 \pm 0.46 \qquad \text{SNU}$$

where the σ_{Cl} is the cross section for the neutrino capture on chlorine in the Homestake experiment, and where the statistical and systematic errors have been combined in quadrature. If we subtract this from the observed rate seen by Homestake:

$$\phi\sigma(pep + \text{CNO} + ^7\text{Be}) = -0.66 \pm 0.52 \qquad \text{SNU}$$

So, we have a severe discrepancy, and this is only exacerbated when one considers that the observed Homestake rate for these three fluxes should be at around 2 SNU according to the SSM. We can carry out the same process on the gallium results, and again the best estimate of the predicted ^{7}Be flux in these detectors is negative.

Perhaps the most succinct paper on this issue is by Parke[22]. He derives the count rates of all four results with respect to the three prominent fluxes: pp, ^{7}Be and ^{8}B. Using the luminosity constraint removes the dependance on the pp flux, giving:

$$Rate_{Cl} = 6.2\phi_B + 1.2\phi_{Be} \qquad \text{SNU}$$

$$Rate_{Kam} = \phi_B \Phi_{SSM}$$

$$Rate_{Ga} = 14\phi_B + 30\phi_{Be} + 78 \qquad \text{SNU}$$

where the lower–case ϕ's represent the fluxes normalised to the SSM and the capital Φ is the flux from the SSM. The coefficients come about from the integration of the fluxes and cross sections for each experiment, while the figure of 78 SNU results from the minimal model. The gallium rate is the weighted mean of the SAGE and GALLEX results. By calculating χ^2 given the experimental values, it is possible to fit the best

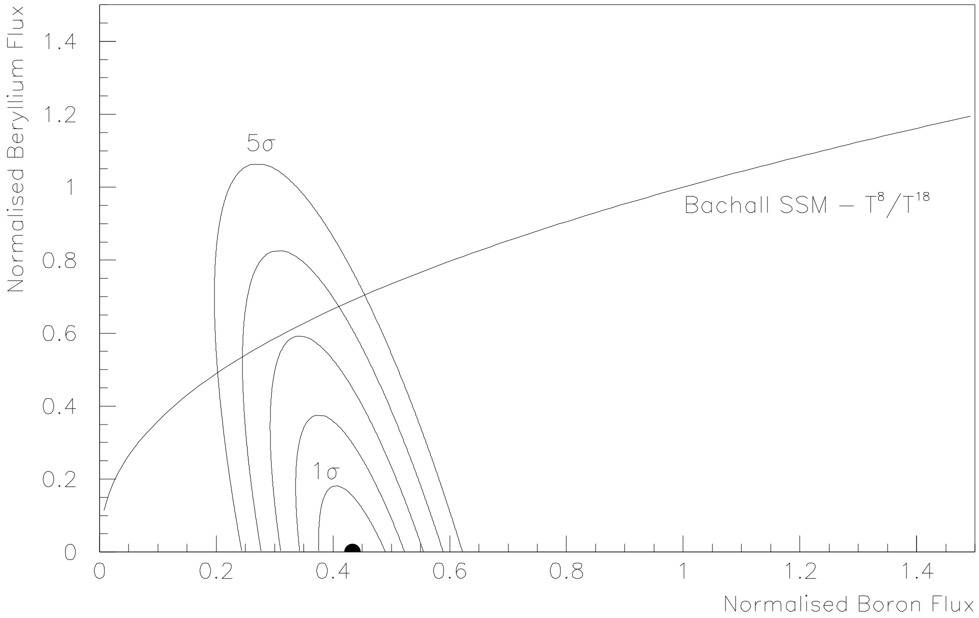

Figure 1.4: A constrained fit of the boron and beryllium flux. The fluxes are normalised to the values of the fluxes given in BP(92)[3], while the line shown is the expected fluxes for changes in the core temperature. The contours are successive sigma away from the best fit, shown as the dot on the x–axis.

values of the two fluxes. With the extra constraint that $\phi_{Be} > 0$, the χ^2 plot is shown in figure 1.4.

As can be seen, even with an artificial free reign on the variation of temperature, the SSM is still over 3σ away from the best value for all four experiments. It should be noted that the best fit is for negative ^{7}Be fluxes, with the constrained fit on the x–axis having a χ^2 value of 2.55.

Considering the equations again, it can be seen that we have three equations but only two unknowns, so that we can throw away one of the experimental results arbitrarily. Unsurprisingly, when this is done the best fit comes about when the Homestake result is removed: this is as expected as the gallium rate is the least sensitive to the ^{8}B flux, while the Kamioka result is made up only of this flux. However, when any of the experiments are excluded, the best constrained fit still occurs at $\phi_{Be} = 0$.

These results would seem to indicate a likely scenario, where the ^{7}Be flux goes to zero, while the ^{8}B flux drops to around half its value, and the pp flux stays the same. There is obviously no 'astrophysical' means by which this can happen, as the production of ^{8}B neutrinos immediately implies the production of ^{7}Be neutrinos in any solar model. Thus, if the solar physics can no longer help us, we must begin to think of ways of reducing the observed neutrino flux. At this point, we must begin to consider the hypothesis of massive neutrinos.

1.7 Massive Neutrinos.

The propagation of massive neutrinos and their subsequent mixing are well known constructs with analogies in the K^0 system. The idea of neutrino mixing was first suggested by Pontecorvo[23]. If the weak eigenstates are not the same as the mass eigenstates then we can say for the simplified case of two neutrino types:

$$|\nu_e>= \cos\theta_v \ |\nu_1> + \sin\theta_v \ |\nu_2>$$

$$|\nu_x>= -\sin\theta_v \ |\nu_1> + \cos\theta_v \ |\nu_2>$$

where e and x are the weak eigenstates and the subscripts 1 and 2 denote the mass eigenstates. The rotation angle, θ_v, is analogous to the Cabbibo angle in the quark sector. Thus the wavefunction of the neutrino is given by the arbitrary superposition of the two mass states such that:

$$|\nu(t)>= \alpha_1(t) \ |\nu_1> + \alpha_2(t) \ |\nu_2>$$

where alphas are time varying function that satisfy the Schrodinger equation. Hence:

$$\alpha_i = \alpha_i(0)e^{-iE_it}$$

Now, we know that at t=0, the neutrinos produced in the sun are purely ν_e, so that:

$$\alpha_1(0) = \cos\theta_v$$

$$\alpha_2(0) = \sin\theta_v$$

So:

$$|\,\nu(t)\,> = \cos\theta_v e^{-iE_1t}\,|\,\nu_1> + \sin\theta_v e^{-iE_2t}\,|\,\nu_2>$$

If we propogate these states with respect to time and then ask what the probability for the state still being ν_e is:

$$P(\nu_e \rightarrow \nu_e) = |\,<\nu_e\,|\,\nu(t)\,>\,|^2 = 1 - \sin^2 2\theta_v \sin^2[(E_2 - E_1)t/2]$$

Now, for particles with energies very much greater than their rest mass:

$$E_i = \sqrt{p^2 + m_i^2} \approx p + \frac{m_i^2}{2p}$$

So:

$$P(\nu_e \rightarrow \nu_e) = 1 - \sin^2 2\theta_v \sin^2(\pi L/L_V)$$

where $L = t$ and $L_V = 4\pi E/\Delta m^2$, and $\Delta m^2 = |m_1^2 - m_2^2|$. The quantity L_V is also known as the vacuum oscillation length. In units more applicable to the solar neutrino problem, we can say:

$$\frac{\pi L}{L_V} = 1.9 \times 10^{11} \left(\frac{1\ \mathrm{MeV}}{E_\nu}\right)\left(\frac{\Delta m^2}{\mathrm{eV}^2}\right)\left(\frac{L}{1\ \mathrm{AU}}\right)$$

Thus for Δm^2 larger than $5 \times 10^{-10} eV^2$, the fraction will run over many values of 2π over the energy spectrum of the pp and ^{8}B fluxes. For line neutrino spectra, changes in the phase come about via the thermal energy of the atoms in the solar core as well as the distribution of distances from which the neutrinos are produced. The average value of the term in the survival probability will become equal to one half as the phase goes through many cycles, and:

$$P(\nu_e \rightarrow \nu_e) = 1 - 0.5 \sin^2 2\theta_v$$

So, the Kamiokande result alone would seem to imply that the vacuum mixing angle is around 45°, and it can also be said that the gallium experiments are seeing of the order of one half of the signal they would expect. This case is known as maximal mixing, a particularly good review of which is given in [24]. However, this case does not allow for the Homestake result.

Alternatively, the Δm^2 parameter can be fine-tuned so that the average value of the $\sin^2(\pi L/L_V)$ rises above 0.5 for the given energy spectra and Sun-Earth distance. This allows more theoretically favoured values (ie $\theta_v \to 0$) of the mixing angle to bring about the required drop in observed flux. This gives a value for Δm^2 of around $10^{-11} eV^2$. This fine-tuning of the parameters is somewhat disfavoured, and indeed should become apparent when the observed rate is compared with the slight changes in the Sun-Earth difference over the course of a year. At present, the data from the gallium experiments cannot confirm or deny this, although tests may be available with better statistics.

Implicit in all this work is that the ν_e's oscillate into either ν_μ's or possibly ν_τ. This means that were a neutral current measurement of the solar neutrino fluxes possible, it would demonstrate oscillations. Although Kamiokande does have some sensitivity to (ν_μ, e) scattering via Z^0's, the relative cross sections at the energies involved are around a factor of seven lower than the (ν_e, e) scattering, which can proceed via W boson exchange also.

This last point raises an interesting question. Since the sun contains electrons but not muons or tauons, then the scattering amplitudes for the ν_e's produced in the centre should be different from those for ν_x, where x is muon or tau flavour. This leads to the well-known MSW effect.

1.7.1 Matter enhanced oscillations – The MSW effect.

The algebra involved in the derivation of the propagation of the neutrino states in a constant density of electrons is in appendix A. The rotation angle which replaces the vacuum angle is given by:

$$\sin^2 2\theta_m = \frac{\sin^2 2\theta_v}{\sin^2 2\theta_v + (\alpha - \cos 2\theta_v)^2}$$

where:

$$\alpha = \frac{\sqrt{2} N_e G_F}{E_2 - E_1} = \frac{\sqrt{2} N_e G_F E}{\Delta m^2}$$

And, the electron density where $\sin^2 2\theta_m = 1$, and hence where the mixing becomes maximal, is:

$$N_{res} \approx 66 N_A \cos 2\theta_v \frac{\Delta m^2}{10^{-4}\mathrm{eV}^2} \frac{10\ \mathrm{MeV}}{E}$$

Now, the SSM gives a core electron density of around $100\mathrm{N}_A$, so at the centre we have a region where:

$$N_e \gg N_{res} \Rightarrow \sin^2 2\theta_m = \sin^2 2\theta_v \left(\frac{N_{res}}{N_e}\right)^2$$

Here, the matter oscillations are the same as the vacuum oscillations, but they are suppressed by the small ratio of electron densities. For regions where:

$$N_e = N_{res} \Rightarrow \sin^2 2\theta_m = 1$$

by definition, and the two states will flip over. Finally, where:

$$N_e \ll N_{res} \Rightarrow \sin^2 2\theta_m = \sin^2 2\theta_v$$

and the matter oscillations tend towards the vacuum oscillations (as indeed they should as $N_e \to 0$). This scheme was first proposed as a possible solution to the solar neutrino problem by Wolfenstein[25]. However, after further work by Mikheyev and Smirnov[26], it was realised that in a medium with varying electron density, these three situations could all be achieved as the neutrino passes out of the solar core. This leads to an energy dependant loss of electron neutrinos. This is now known as the MSW effect.

This energy dependance comes about in a number of ways. Firstly, the maximum value for the electron density at the solar core imposes a lower energy limit for the mixing for a given value of Δm^2. Also, if the local electron density changes slowly, the matter eigenstates are able to follow rotating mass eigenstate, so that ν_e is completely converted to ν_x as it propagates into the vacuum. This is known as the adiabatic condition, and is satisfied for a large range of energies above the minimum energy. However, where the electron density changes quickly, the mixing will be large, but will not have the chance to become resonant, and the probability of oscillation will be reduced from the maximal case. We can plot the transition probability as a function of the quantity $E/\Delta m^2$ in figure 1.5.

The form of this curve is perhaps the most exciting feature of the MSW effect. With an adjustment of the mass difference, it can be seen that the pp flux could be moved below the minimum energy, giving a complete survival of that flux. For vacuum mixing angles of the order of those shown in figure 1.5, this would mean that the ^{7}Be flux would be almost completely extinguished, and there would be a distortion of the

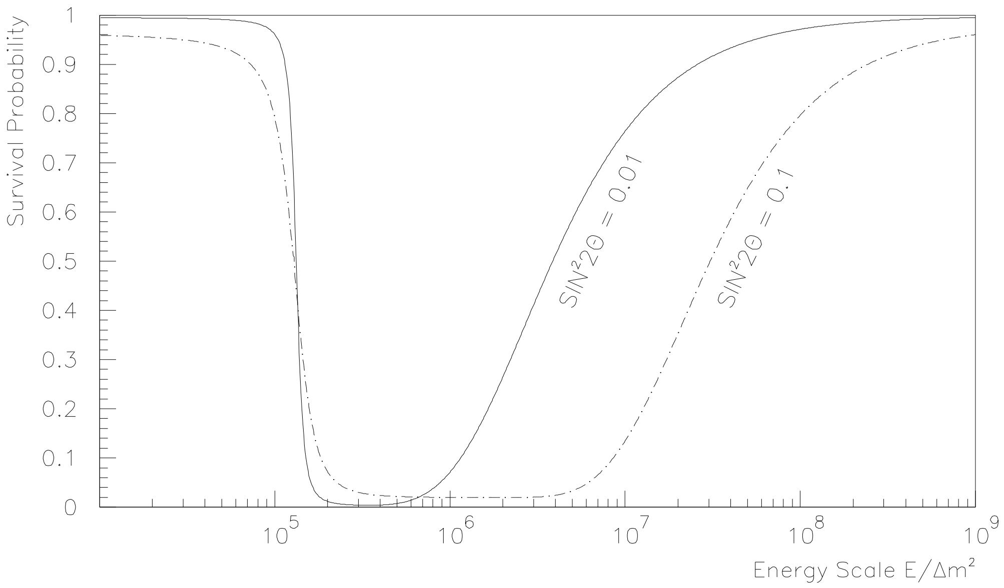

Figure 1.5: The survival probability per unit energy for electron neutrinos as a function of $E/\Delta m^2$. The two lines shown are for differing vacuum oscillation probabilities. This energy dependency on the mixing for a given mass difference is the most attractive feature of the MSW effect with respect to the solar neutrino problem.

^{8}B flux spectrum which would be sensitive to the exact value of the vacuum mixing angle.

With the MSW effect in mind, there are three requirements for new experiments to test the hypothesis: increased event rates, spectral measurements and neutral current measurements. At present there are three experiments which hope to meet this challenge, and they are detailed below:

1.8 New solar neutrino experiments.

The first of the new generation experiments that will come on line (in mid 96) is the Super Kamiokande experiment[27]. As the name suggests, this is a larger version of the previous Kamioka experiments, and indeed is situated in the same mine. With the fiducial volume of the light water increased from 680 tonnes to 22,000 tonnes, and the PMT coverage increased by a factor of around 2, this new experiment should see around 7000 events per year, compared with the 200 per year seen by Kamiokande

II. With the addition of an improved water purification system, it is hoped that the threshold in the detector will be at 5 MeV.

As with Kamiokande, the neutrinos scatter off the electrons in the light water, leading to a measured electron energy spectrum which will be a convolution of the ^{8}B neutrino spectrum with the normal ν-e scattering amplitudes. All the neutrino flavours will scatter in this process, although the cross section for electron-flavour neutrinos will be around 7 times higher in the 'solar' energy range due to the increased amplitude from W-boson exchange. The distortion of the spectral shape of the ^{8}B spectrum would indicate a 'non-astrophysical' solution to the solar neutrino problem, with the MSW effect obviously taking centre stage.

The Borexino detector will consist of 200 tonnes of scintillator enclosed in a spherical vessel, this being viewed by a few hundred PMT's. The use of the liquid scintillator means that the elastic scattering of neutrinos off electrons will be visible down to a projected threshold of 250 keV. This will allow the determination of the ^{7}Be flux directly and in real time: the much higher flux for these neutrinos is perhaps the driving force behind the experiment, with an expectation of around 50 events per day. The proposal has already been demonstrated to work in a test detector of 3 tonne volume, with the required levels of radioactive purity, at 10^{-16} g/g of Th and U, having been achieved.

The Sudbury Neutrino Observatory, or SNO, is the third major experiment. It is currently under construction, and consists of 1000 tonnes of heavy water, D_2O. This use of deuterium allows both a charge current and neutral current flux measurement independently and concurrently. With an expected threshold of 5 MeV, the detector should be able to both measure any shape distortion of the ^{8}B spectrum, and also the mixing of the neutrinos using the neutral current channel. The experiment is detailed in the next chapter.

Chapter 2

The Sudbury Neutrino Observatory.

From the previous chapter, we have a good idea what is required in the next evolution of solar neutrino detectors: low thresholds coupled with large event rates, good energy resolution, low backgrounds and neutrino-flavour differentiation. This last point is especially important: if muon or tau flavour neutrinos could be seen coming from the sun at any significant rate, it would prove that neutrino oscillations were occurring. The SNO detector[28] should meet a number of these requirements, with the only disadvantage really being the threshold of 5 MeV that will be imposed by the backgrounds. The presence of deuterium allows several new channels to be investigated, and in particular will allow the simultaneous measurement of both the electron neutrino flux via the charge current (CC) reaction, and the flux of all types of neutrinos, via the neutral current (NC) reaction. The two new channels, the CC and NC, are the essence of the SNO detector, and although the absolute cross sections are known only to around $\pm 5\%$, their relative value is known to around 0.5%[31]. A comparison of these two rates will thus test the neutrino mixing hypothesis independently of the solar model.

We will begin by describing the detector itself, and then go on to the interaction channels, with the final sections covering the calibration and running of the detector.

2.1 Detector construction.

The Sudbury Neutrino Observatory, or SNO, is a large heavy water Čerenkov detector for the neutrinos from the ^{8}B decay in the sun. A schematic diagram of the detector is shown in figure 2.1. The active central volume consists of 1000 tonnes of high purity heavy water, D_2O, held in an acrylic sphere of 6 metre radius (the AV), which sits in another 7 kilo-tonnes of light water, and is viewed by 9600 8 inch PMT's arranged on an 8.2 metre concentric spherical shell (the PSUP). The detector and a large experimental area are situated in a specially constructed cavity in the INCO nickel mine in Sudbury, Ontario: the site was selected both for its depth and for the very hard quality of the rock allowing such a large cavity to be excavated. The approximately barrel-shaped cavity is situated at a depth of 2.07 km, this being equivalent to 6.2 km of water, reducing the muon flux to around 100 per square metre per year. For comparison, the Super Kamiokande site will have a muon flux around 200 times higher.

The acrylic sphere is constructed from segments of order 2 metres by 2 metres and 51 mm thick, with 100mm thick pieces around the equator in order to accommodate grooves for supporting ropes. The acrylic used is not UV-stabilized, giving a lower cut-off in wavelength terms of around 330 nm, and hence more transmission of the Čerenkov light. Both the light and heavy water will be cleaned in purification plants detailed in the next chapter, and a slight overpressure of light water will be held on the inside of the PSUP so as to wash away any radioactive impurities that leach out of the photomultiplier glass. The D_2O itself is on loan from AECL, with the total cost of buying the heavy water being around twice the projected cost of the detector alone. This imposes a life-span on the detector, with the loan expiring in 2003.

Each of the photomultipliers will have a non-imaging concentrator[30] attached to it. These increase the effective area of light sensitivity by a factor of 1.6, while adding only around 10% to the cost of the PMT/concentrator combination. The opening angle of the concentrators is such that the whole of the D_2O and 1 metre of the surrounding H_2O is visible to the PMT, while neighboring PMT's are optically shielded from each other. This is important as radioactivity in the glass envelopes of the photomultipliers can create bursts of light and hence create cross-talk. The total coverage achieved by the PMT's with such a construction is around 65% of the area of the PSUP.

The photomultipliers themselves are Hamamatsu R1408 8 inch bialkali photocathode tubes, which although they have a poor charge response, are very 'fast' in that they have very small transit time spreads for such large area tubes. This should allow accurate position reconstruction of the event, which in turn should allow a fiducial volume to be defined within the D_2O away from the more heavily contaminated acrylic vessel. The signals from the PMT's are processed in custom ADC/TDC chips[29] situated

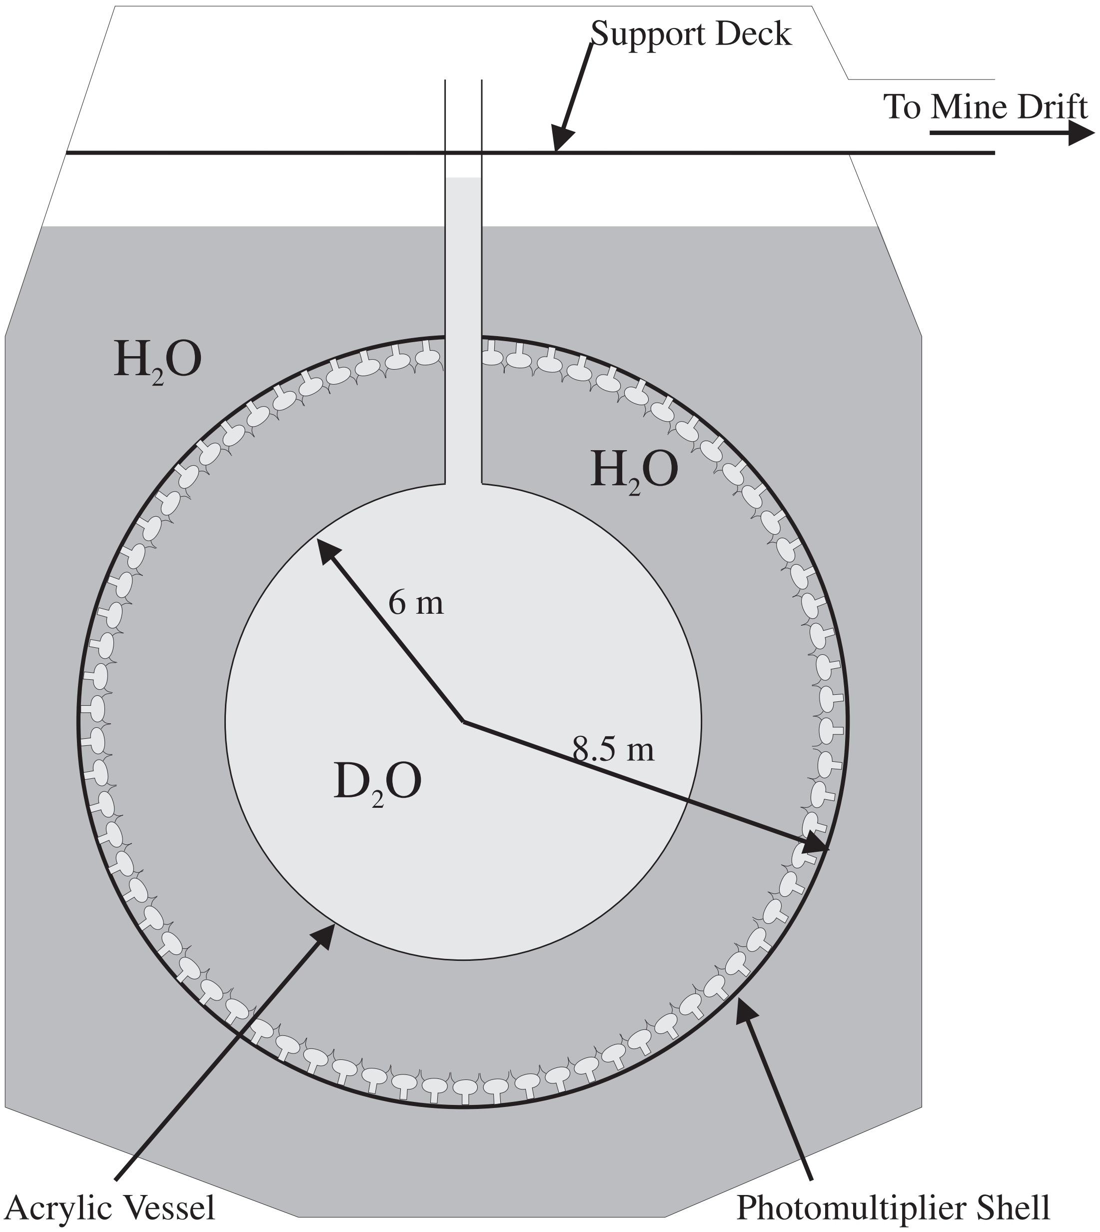

Figure 2.1: A schematic view of the SNO detector. It should be noted that the original design had a barrel-like central vessel, but this has been changed to a spherical volume for a number of reasons.

on the deck above the detector itself. Data will be written to tape when the number of hits within a 100 ns window exceeds a certain level; the exact level in number of hits (NHITS) will be decided by the speed with which the data can be written. The trigger is essentially dead-time less, thus allowing large bursts of data, possibly from a supernova, to be contained and subsequently recorded.

Reaction	Acronym	Q-value (MeV)	σ at 10 MeV /cm^2 ($\times 10^{44}$)
$\nu_e + d \to p + p + e^-$	CC	-1.44	270
$\nu_x + e^- \to \nu_x + e^-$	ES	–	9.2
$\nu_x + d \to n + p + \nu_x$	NC	-2.22	110
$\overline{\nu}_e + d \to n + n + e^+$	CC	-4.03	120
$\overline{\nu}_e + p \to n + e^+$	CC	-1.80	500

Table 2.1: A summary of the channels in the SNO detector. The first three are the most important for solar neutrinos, whereas the last two are for possible supernova detection, with the very last only occurring in the 7 ktonnes of light water.

Reaction	SSM (full)	SSM $\times$ 1/3	Vacuum mixing	MSW Adiabatic	MSW Non-adiabatic
CC	9750	3255	3260	3750	3750
ES	1100	360	480	645	488
NC	2800	933	2800	2800	2800

Table 2.2: The expected event rates per year in the SNO detector. Note that the neutrino mixing solutions use the full flux from the SSM. As can be seen, the flavour independence of the NC reaction differentiates between mixing-based solutions and astrophysical solutions. The adiabatic MSW case is where the sharp turn-on of mixing occurs at 9 MeV, whereas the non-adiabatic case is for the preferred solution of the existing experiments.

2.2 Interaction channels

The construction of such a relatively small water Čerenkov detector may seem a little strange when compared with the 50 kilo-tonne Super Kamiokande experiment mentioned in the previous chapter. However, as well as the detection of neutrinos via elastic scattering of electrons, the presence of deuterium allows new channels to be investigated. All the reactions are shown in table 2.1, the energy dependant cross sections for the reactions are in figure 2.2, while the expected numbers of events for some scenarios are shown in table 2.2. We will go on to describe the three major channels in more detail:

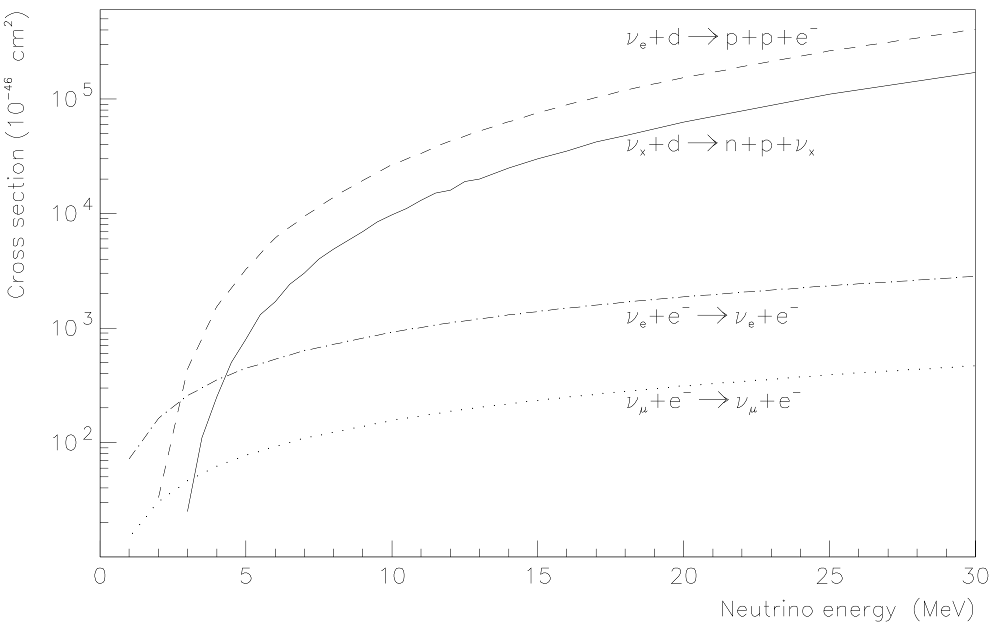

Figure 2.2: The cross sections against energy for the processes shown in the graph. Notice the log scale on the y-axis, with the prevalence of NC and CC reactions when compared with ES reactions at energies above 5 MeV. The deuterium reactions are taken from Kubodera et al.[31].

The charge current reaction.

The charge current (CC) reaction is the 'bread-and-butter' reaction in SNO, with the production of an electron with 2 protons. The mass ratio means that the recoil of these protons can be ignored, and the electron will have a kinetic energy equal to the neutrino energy minus 1.44 MeV. This is very important as it allows SNO to measure directly the differential energy spectrum of neutrinos, rather than the convoluted spectra of Kamiokande and Borexino. This spectral shape measurement is particularly important as the favoured MSW solution distorts the ^{8}B spectrum to an extent that is not consistent with any astrophysical solution of the solar neutrino problem.

However, the Q-value of the CC reaction also imposes a strict low energy threshold on the neutrinos of 1.44 MeV, with such a value excluding the ^{7}Be and pp fluxes. In practice, the backgrounds in the detector will impose a low energy limit of around 5 MeV 'electron' energy, leading to a low energy limit of around 6.5 MeV on the neutrinos. This means that this channel will only be sensitive to the highest energy neutrino fluxes, ^{8}B and hep.

Another feature of the CC reaction is that the break-up of the deuteron has some directionality due to the transition being Gamow-Teller, with an angular distribution:

$$I(\theta_{Sun}) = 1 - \frac{1}{3}\left(\frac{v}{c}\right)\cos\theta_{Sun} \approx 1 - \frac{1}{3}\cos\theta_{Sun}$$

This gives a 3 to 2 forward to backward asymmetry in terms of numbers of events in the forward and backward hemispheres, but it is not envisaged that this will be of any particular experimental interest due to the rather poor correlation.

The elastic scattering reaction.

The elastic scattering (ES) reaction is exactly that used in Kamiokande and Borexino. It is almost the opposite of the CC reaction, with the energy distribution of electrons being not strongly related to the initial neutrino energy, while the angular distribution is much more forward peaked. Of course, there is no theoretical low energy limit on the neutrinos, but again the backgrounds in the detector will mean that only neutrinos of 5 MeV and above will be detected. It is unlikely that this ES reaction will be widely used by SNO, although it could be used as a cross-calibration with data from the Kamiokande experiment.

The separation of ES from CC events is of course a problem, in that both produce a single energetic electron. However, the nature of the ES reaction means that the angular distribution of events will be forward peaked away from the sun. This is of course evidenced by the plot from the original Kamiokande experiment. It should be possible to achieve a systematic separation of the two channels by fitting Monte Carlo histograms of the $\cos\theta$ distributions to the data, and hence derive the rates of the two reactions. Also, were a 'clean' CC signal required, for example when comparing the energy spectrum with the MSW distortions, those events with angles close to the solar angle could be discarded with a consequential, and theoretically calculable, loss in rate.

The neutral current reaction.

The neutral current (NC) reaction is the most important feature of the SNO detector. The neutrino-disintegration of the deuteron via a Z boson is flavour independent. Thus the rate of NC reactions gives a neutrino flux measurement that is insensitive to any flavour oscillation that occurs between neutrino types, and hence, when compared to the CC rate, a test of neutrino mixing without reference to the solar model.

The NC reaction itself, with a threshold of 2.2 MeV, will be sensitive only to the ^{8}B and hep neutrinos, and is signalled by the production of a free neutron within the detector. This neutron, after having thermalized, can be detected by one of two methods:

The first and default option is to dissolve 2 tonnes of $MgCl_2$ in the D_2O itself. The thermal (n, γ) capture on the ^{35}Cl produces a γ-cascade with a total energy of 8.6 MeV, and an average of 2.6 γ-rays. These will Compton scatter off electrons, with these electrons producing Čerenkov light. This method is complicated by the fact that it produces a signal within the detector which is similar to that produced by the CC reaction. The slight difference in hit patterns may allow a statistical separation using neural net techniques[33], but at present the default method is to cycle between salt-in (NC+CC) and salt-out (CC) runs so as to calculate the two rates.

The second method involves the use of discrete proportional counters [32] placed within the D_2O. The neutron thermalises and then captures on the ^{3}He in the helium/xenon mixture within the detectors to produce a characteristic signal in the central wire of the proportional counter. These neutral current detectors (NCD's) have the added advantage that the detection method is entirely separate from the CC reaction, leading to well separated CC and NC rates.

2.3 Detector response

The detector's response to energetic electrons is being simulated using a Monte Carlo code known as SNOMAN[34]. This uses the EGS 4 code to track electrons and gammas through the detector, with various additions having been made by the SNO collaboration for Čerenkov light generation[35] and neutron transport. At present the code shows that around 10 hits are expected per MeV of initial electron energy, with this relationship being linear over the energy range of interest. With the threshold at 5 MeV, this represents an energy resolution of 14% using root N errors on the number of hits, which should be the dominant statistic. The expected differential energy spectra are shown in figure 2.3. The radioactive backgrounds, which will be further discussed in chapter 3, are not shown here but will impose a lower limit for event detection at around 50 NHITS, or 5 MeV.

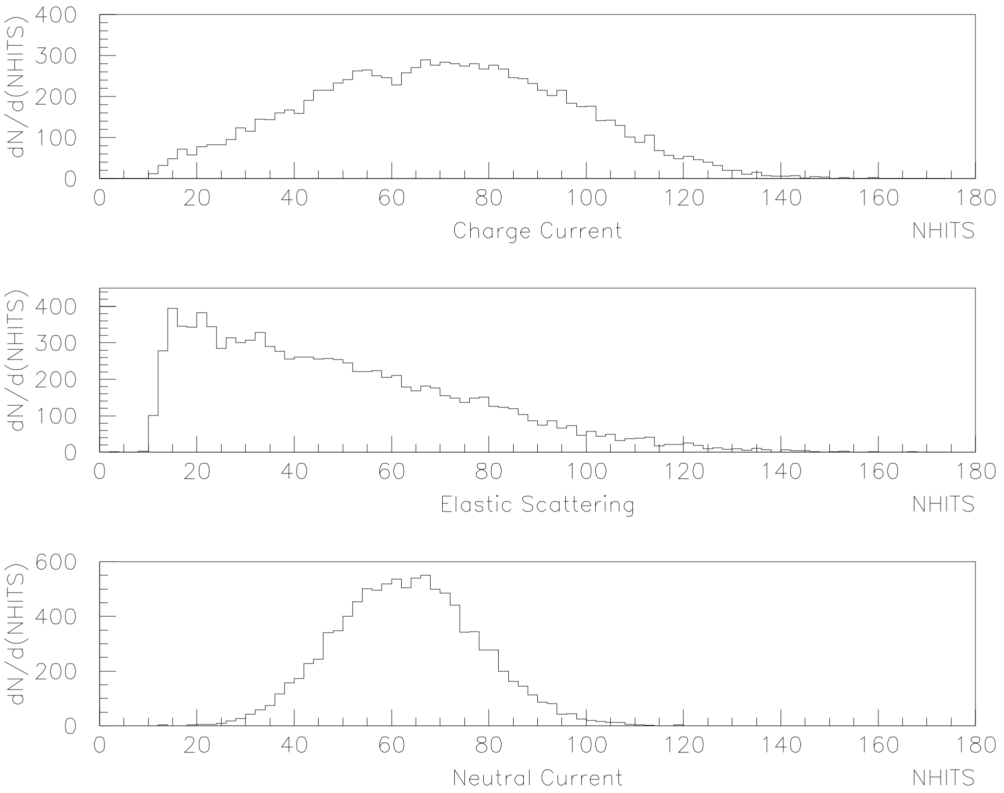

Figure 2.3: The differential energy spectra for the CC, ES and NC channels in the SNO detector, with 10000 events in each histogram. The data comes from the SNO Monte Carlo and the x-axis is shown in NHITS, the number of PMT's firing within a 100 ns window of the event. As stated in the text, the rough calibration is 10 NHITS per MeV equivalent electron energy. The threshold is expected to lie at around 50 hits.

The detectors directional and positional resolution is of course much more dependent on the event reconstruction algorithm used in the Monte Carlo, whereas the energy is a calorimetric measurement. At present, position and direction resolutions of 25 cm and 20° respectively are expected for a 5 MeV electron. These come about both from the transit time spread ($\approx$1.5 ns) of the PMT's and the multiple scattering of the electron along its track which tends to smear the Čerenkov ring projected onto the PMT shell.

The detector will of course need to be calibrated, at least in so far as to adjust the Monte Carlo to the data. The range of sources required for a full calibration is fairly wide:

- The first source is to check the optical properties of the detector and consists of a pulsed nitrogen dye laser at 384 nm, shining into a diffusing ball. This 'laser ball' will be moved around the inside of the acrylic vessel and the response of the photomultipliers to the light will be checked against the Monte Carlo. It may also be possible to calibrate out the slight differences in the transmission of the acrylic panels that make up the AV, and measure the absolute efficiencies of the PMT's.

- In order to calibrate the neutron transport code another source will be formed by enclosing a sample of ^{252}Cf in a block of acrylic. The californium itself is mixed with a scintillator before being inserted into a small well in the acrylic; the fission of the californium will thus be tagged by the scintillation light, with an average of three neutrons then undergoing capture within the detector.

- There will be three sources to calibrate the response to energetic electrons: a ^{16}N γ-ray source of energy 6 MeV should allow a calibration of the detector near the threshold; a ^{3}H(p,γ)^{4}He γ-ray source of energy 20 MeV will encompass the entire energy range of interest for solar neutrinos; and finally a tagged low energy source using neutron activated sodium iodide. The ^{24}Na produced in the crystal decays with the production of two low energy gammas, so allowing a check on the detector's response to electrons below threshold.

2.3.1 Running modes

Although SNO aims to resolve the SNP by measuring the CC to NC ratio, it is also known that for the most likely MSW solution, there will be a significant distortion of the ^{8}B energy spectrum. This change will be most notable at the lower energy region of the spectrum, where both the ES and the background can also produce distortions. If, as proposed, the first year of operation in SNO is carried out without a neutral–current detection scheme, the free neutrons will still be generated. Around half of these, based on Monte Carlo results, will travel out of the D_2O and will be captured either on hydrogen in the acrylic itself, or in the light water. However, the other half of the neutrons will capture on the deuterium, producing 6.25 MeV γ's. These will contribute to a further background at the low end of the energy spectrum. It is for this reason that a neutron poison has been suggested.

The most probable candidate for such a poison would be to dissolve around 100 kilos of $Li(NO_3)$ in the detector, with the lithium being isotopically enriched ^{6}Li. This would reduce the deuterium captures down to the few percent level, and the capture of the thermal neutrons on the lithium would produce only tritium and an alpha, with no significant light output. This CC–only running mode is currently being modelled, while the cleaning of the lithium with respect to impurities is considered to be a minor addition to the rig that will be used to clean the $MgCl_2$.

Another running mode is the full mode: here the neutrons are detected either by the chlorine additive or the NCD's. As mentioned before, the chlorine salt may be cycled in and out of the D_2O so as to compare the CC rate to the CC+NC rate. This will not be an option for the NCD's, as they are difficult to install and remove, but their response to neutrons is completely distinct from the PMT's response to electrons, and so no cycling is necessary.

Finally, the central volume of the detector will be filled with light water. At first, it was proposed that this would be done at the start of the experiment, but it was realised that the light water would diffuse into the acrylic vessel, and when the light water was replaced by the heavy water, the degradation of the D_2O would exceed the prescribed limits. Thus it is now thought that the light-water fill will be done when the D_2O loan runs out.

The purpose of the light water fill is to investigate the backgrounds with no CC or NC interference. Although elastic scattering will still occur, this will be subtracted away to give a pure background run of the detector. This rather complicated procedure is required as we have no control as to whether the source is on or off.

2.4 Supernova detection

To end the chapter, we will briefly discuss the abilities of SNO as a supernova detector. After the detection of SN1987A by IMB and Kamiokande[36], there has been a certain amount of work on what SNO could expect to see were a supernova to occur in the vicinity of the sun. After the paper of Burrows et al.[37], the number of events expected in SNO are shown in table 2.3.

As can be seen, the burst phase, where the iron core starts its infall and the protons

Reaction	Target	Events from ν_e burst	Events from cooling phase
$d + \nu_e \to p + p + e^-$	D_2O	10	82
$e + \nu_x \to e + \nu_x$	D_2O/H_2O	1	22
$d + \nu_x \to n + p + \nu_x$	D_2O	6	272
$d + \overline{\nu}_e \to n + n + e^+$	D_2O	0	67
$p + \overline{\nu}_e \to n + e^+$	H_2O	0	330

Table 2.3: The expected number of events in SNO for a supernova at 10 kiloparsecs from the sun. The expected rate of such supernovae is around 1 per 10 years. The large NC number in the cooling phase may allow time-of-flight measurements to be done to measure the neutrino mass directly.

and electrons fuse to form neutrons, produces a few events in the detector. These should be distinguishable from the solar neutrino signal not only by their close proximity in time (the burst phase is thought to last around 10 ms), but also by their high energy of around 15 MeV. The cooling phase, as the protoneutron star loses energy by emitting neutrinos, is though to last around 10 seconds, and involves the emission of all neutrino and anti-neutrino flavours. This offers perhaps the most exciting prospect for the SNO supernova detection: mass measurement by time of flight.

It is envisaged that mass differences in the range 50 eV to 100 KeV will be accessible. Below 50 eV the spread between the three neutrino bursts would be affected by the long cooling phase time, whereas above 100 KeV, the bursts would have large separations in time and could not be distinguishable from the solar neutrino background. At present however, the mass limits are 270 KeV for ν_μ and 31 MeV for ν_τ[38], so this range would be a huge improvement. These mass limits also depend on the distance of the supernova. If it occurs too close, the neutrinos will not be spread-out enough in time to resolve the three bursts expected, whereas if the supernova is too far away, the small total number of events will be difficult to pick out of the solar neutrino signal.

Although the abilities of SNO, both as a solar neutrino and supernova detector are beyond question, it should be noted that in the above text, there has been no talk on the backgrounds. The SNO detector is particularly susceptible to natural radioactivity, and the issues deserve a separate chapter.

Chapter 3

Backgrounds in SNO

SNO, like so many 'low' energy experiments, is subject to backgrounds that arise from naturally occurring radioactivity. The main isotopes of interest are those that produce high energy electrons and those that give γ-rays of energies greater than 2.2 MeV. The electrons produce signals which are indistinguishable from the CC events, and hence lead to a lower limit on the energy spectrum; the γ-rays lead to photo-disintegration of the deuterons, which then mimic the NC signal. Of those nuclides which decay in these fashions, we really need only consider those with long half-lives, and which are present in relatively high abundances in the materials to be used in SNO. This limits the discussion to two heavy nuclides and their subsequent decay chains: ^{232}Th and ^{238}U.

These two heavy heavy nuclides decay via a number of radioactive daughters, with many alphas, betas and gammas being emitted over the chain of decays. The two chains are shown in figures 3.1 and 3.2. The data for these diagrams is taken from [39].

The majority of the decays in these chains are alphas of 4-6 MeV. These do not pose a major problem, as they do not produce any Čerenkov light. Of a more serious nature are the β-decays followed by the prompt (i.e. of the order of ns) emission of γ-rays from excited states of the daughter, these also being known as β-γ decays. Near the end of the thorium chain we have the ^{208}Tl$\rightarrow$^{208}Pb transition; while in the uranium chain we have the ^{214}Bi$\rightarrow$^{214}Po transition (the 0.021% branch to ^{210}Tl makes it uninteresting). These two decays are shown in more detail in figure 3.3.

These two transitions have a total (β+γ) energy of up to 5 MeV, and hence their presence in the D_2O can set the lower energy limit on the observed electron energy spectrum induced by solar neutrinos via the charge current interaction. These β-γ

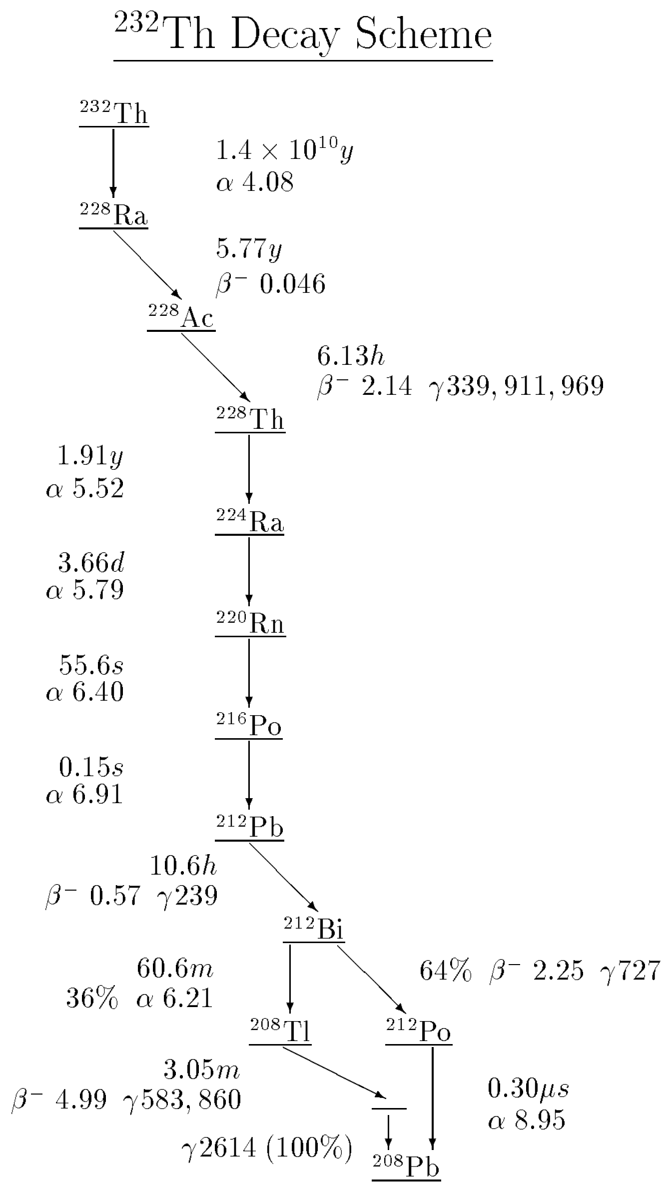

Figure 3.1: The thorium decay chain. All half-lives are shown, with the Q values of beta and alpha decays in MeV, and gamma rays in keV.

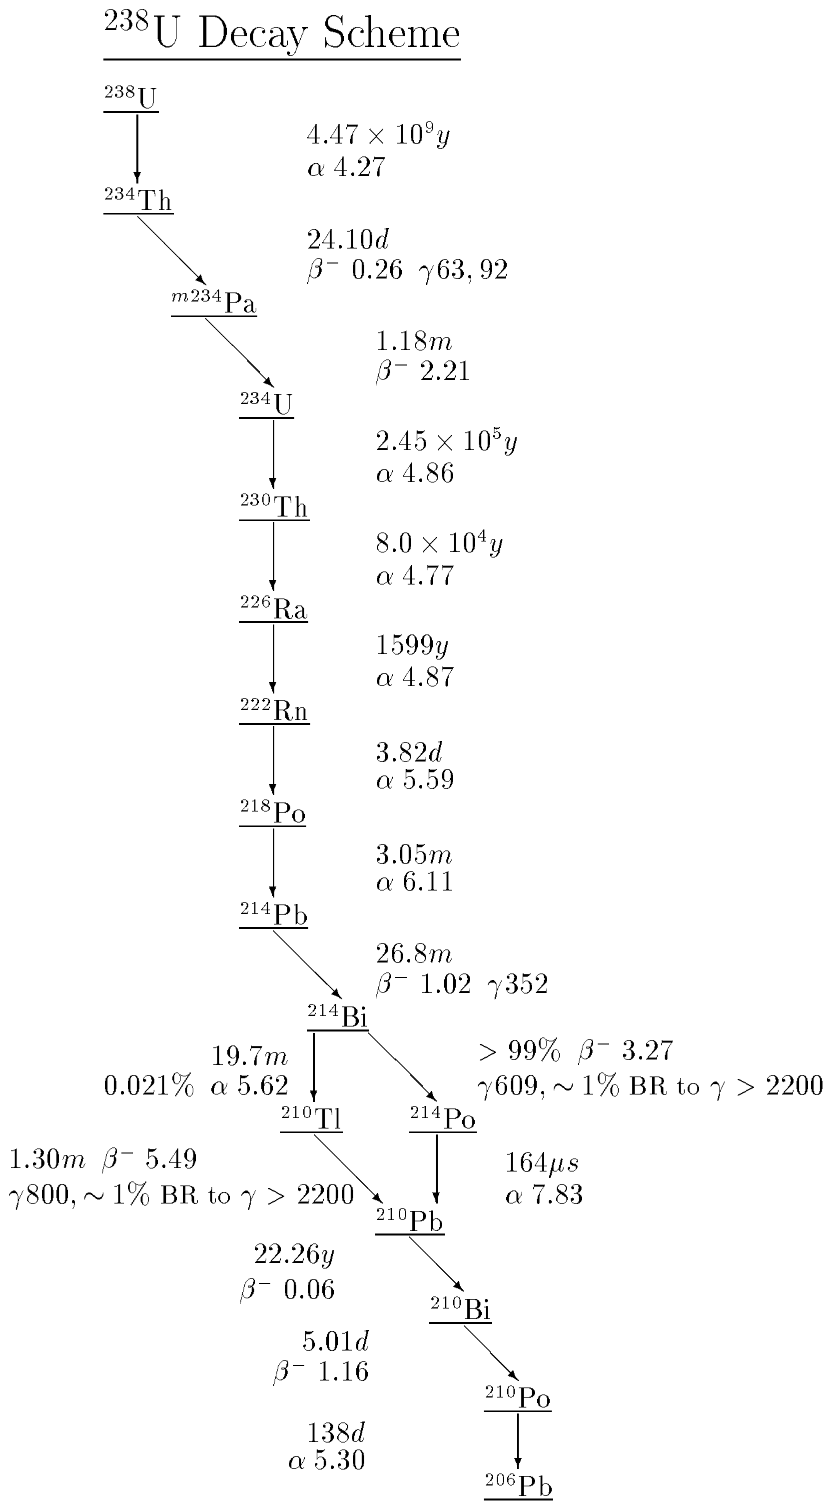

Figure 3.2: The uranium decay chain. All half-lives are shown, with the Q values of beta and alpha decays in MeV, and gamma rays in keV.

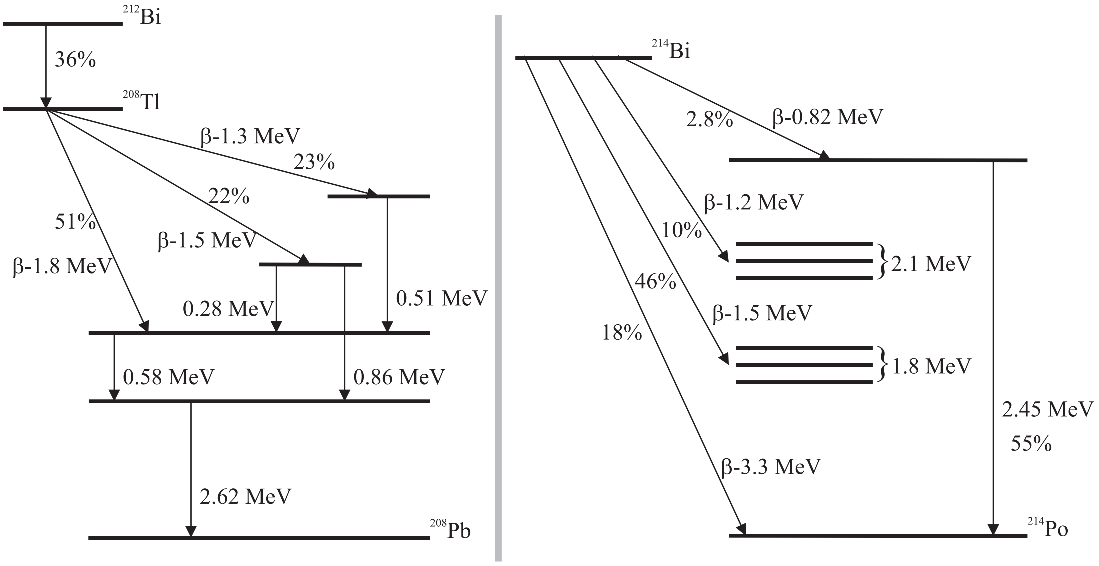

Figure 3.3: The decay schemes of ^{208}Tl and ^{214}Bi. The thallium decay always produces the 2.62 MeV gamma, although the branching fraction to the thallium is around one third. The bismuth decay has a one percent (2.8% $\times$ 55%) branch to the 2.45 MeV gamma.

decays will be produced throughout the detector and are classed into two separate groups: internal β-γ decays are those which occur in the D_2O, the acrylic and the H_2O; external β-γ decays are those which come from the PMT's and the PMT support structure, the PSUP. Another class of background events is the high energy (>5 MeV) gammas from the cavity and the PSUP. The reconstructed energy spectra of these decays is shown in figure 3.4.

In this figure, the relative normalisations of the distributions are scaled for the radioactive contamination levels presently expected for the detector. The Monte Carlo generated data are for those events reconstructed within the D_2O and are simulated for the situation when the central vessel is filled with D_2O and 2 tonnes of $MgCl_2$.

The CC spectrum that is included for reference demonstrates that the threshold for solar-neutrino-induced electrons will be at around 5 MeV. Below this energy, the background rises steeply, producing what is known as the background 'wall'. The activity contained within the active volume itself, the internal β-γ decays, will be the major contributor to the background. The raw rate of external β-γś is much higher than the internal β-γś, but the chance that they will be reconstructed within the D_2O is much lower.

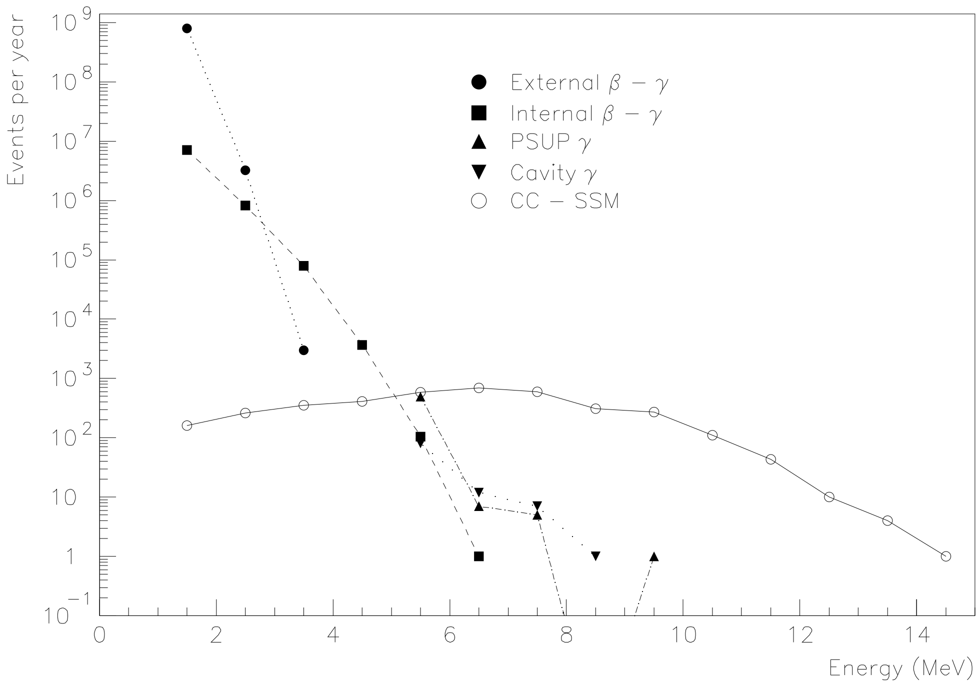

Figure 3.4: The constituents of the background contribution in the electron energy spectrum. The y-scale is the number of events that reconstruct within the acrylic vessel per year. The CC rate is shown for the SSM for comparison. Note that the external β-γ events are scaled as the Monte Carlo models them more slowly than they occur within the detector.

The steep fall in the external β-γ spectrum is due in part to the energy spectrum of the events, but is also a consequence of the reconstruction of the event. As the number of hit PMT's rises, the reconstruction accuracy also rises: as well as there being more hits, the relative contribution from the uncorrelated noise rate due to the dark current in the PMT's is reduced. Thus the chance of reconstructing (wrongly) within the acrylic vessel falls with increasing energy.

The cavity-γ and PSUP-γ backgrounds shown come about via γ rays produced in $(\alpha, p\gamma)$ or $(\alpha, n\gamma)$ captures in the rock and the PSUP structure. These can produce very high energy gammas, ranging up to around 10 MeV, and obviously if one such gamma scatters off an electron within the D_2O it will be indistinguishable from a solar-induced electron. This problem seems to be fairly well under control given recent background measurements in the SNO cavity itself. As can be seen, the two distributions are small, and it is not anticipated that this will be a large concern for the SNO experiment.

There has been some worry as to whether radionuclides could be leached out of the PMT glass into the light water over the ten year period of the experiment. In order to counter this the light water will be pumped into the inside of the PSUP, so that there will be a net flow of water out through small holes moulded in the plastic of the concentrators. This flow outwards means that any β emitters leached out should not approach the acrylic vessel.

3.1 Neutral Current Background

The most important decays within the thorium and uranium chains are the gamma rays of energy above 2.2 MeV, the deuteron binding energy. If such decays occur and photodisintegrate deuterium, free neutrons will be produced. Since the only method of detection in SNO is thermal neutron capture, no measurement of the neutron's initial energy is possible, and hence these neutrons will be indistinguishable from neutrons produced by the neutrino dissociation. Such a background is crucial in that SNO's importance is due to a precise determination of the neutral current rate.

The number of these gammas we can allow per day has been calculated and in essence sets the limiting activity for the Th and U in the heavy water. After the generation of the 2.6 MeV gamma in the Th chain, it has a deuteron disintegration cross section of around 12 mb, while the electron scattering cross section (mainly made up of Compton scattering) is 125 mb. Also, we have two deuterons for every ten electrons in a molecule of D_2O, and this reduces the effective cross section still further to the level that one gamma in 470 will produce a free neutron. In the uranium chain, the lower energy of the gamma means that one in every 750 of these gammas will lead to a free neutron.

When the specific activities of the thorium and uranium are factored in, we can estimate the total mass of thorium and uranium allowed within the D_2O volume. Assuming a rate of around 10 solar-neutrino-induced neutrons per day, if we want the neutral current background to contribute only one neutron per day, we require $3.7 * 10^{-15}$ g/g of ^{232}Th and $4.5 * 10^{-14}$ g/g of ^{238}U, where the specific activities have been taken from [39]. To compare these levels to everyday concentrations is a little cavalier, but the conventional statement is that tap water contains around 10^{-12} g/g of each isotope. It must be assumed that the D_2O will not arrive with this level of cleanliness, so an extraction plant must be built to remove the impurities. It is also thought that the very high purity of the water may leach these heavy elements out of any surface it is in contact with, so the purification technique must be continuous.

An important point to note is where these gamma decays occur: in the decay of ^{208}Tl in the thorium chain, and in the decay of ^{214}Bi in the uranium chain. If we look back up the chain, we can see that isotopes such as the 2 year ^{228}Th and the 1600 year ^{226}Ra can lead to significant disequilibria where the background production of neutrons can vary over the proposed ten year running cycle. This means that cleaning the detector of only ^{232}Th and ^{238}U, which are at the highest gramme per gramme levels due to their long half lives, will not actually do us any good. We need to extract the isotopes down the chain, especially the last two relatively long lived isotopes (^{228}Th and ^{226}Ra) above the gamma-decaying steps mentioned above.

Perhaps more importantly still, we not only need to extract these isotopes, but we have to have a reliable and calibrated assay system in order to make a background subtraction from the number of neutrons seen in the detector. Again, this assay technique must be sensitive to the isotopes just preceding the gamma producing step in order to give as good an estimate as possible of the flux of these gammas within the heavy water vessel. At present there are two schemes to do this assay using the detector itself, and three assay processes tied in with the extraction techniques.

3.2 Direct assay techniques

Fitting the background wall

The progress made in the Monte Carlo of the detector, SNOMAN, has allowed many hypotheses and schemes to be tried out. From a study of the generated energy spectrum for the detector, an assay technique has recently been proposed to use the shape of the background wall itself to determine the relative concentrations of thorium, uranium and possibly potassium. Of course, the SNO detector was never optimised for such low energy events, but the shape and height of those events below 5 MeV in the electron energy spectrum may give order of magnitude estimates for the rate of high energy beta decays in the detector. This is a good number to know as it gives the activity of the radionuclides at the end of the chain and hence the activity of the gamma producing steps. Also, it gives a direct method of measuring the activity in the vessel, rather than the extraction and assay techniques detailed below.

However, the analysis with the Monte Carlo has shown that such a determination will be difficult at best due to the poor energy resolution and the mis-reconstruction of events that occur outside the acrylic vessel, especially those beta-gamma decays from the glass of the PMT's. Another possible problem may be the presence of potassium 40. Potassium is present in the anhydrous $MgCl_2$ that is added to the detector at a level of approximately 4 ppm, with around 0.01% of this being ^{40}K. The β's of endpoint 1.35 MeV produced by the decay of the isotope will swamp the low energy part of the signal, and energy cuts may be required in order to remove the potassium background. This of course removes a large part of the signal, and will depend on the accuracy of the Monte Carlo.

Beta neutron coincidence

A more reliable method of determination uses a coincidence measurement. The 2.6 MeV gamma transition in ^{208}Pb is preceded by the beta decay of ^{208}Tl, with an endpoint of around 1.8 MeV. The thermalisation and capture time of the background neutron produced by the γ is around 4 ms when the chlorine is present in the heavy water volume. Thus, when the neutron capture event is seen, the PMT data can be scanned over the previous few milliseconds to look for a β event of an MeV or so. Such a signal implies that the neutron was gamma induced, and this method could be used to measure the background in the detector. There is also the prospect of a volume cut being applied around the neutron capture, but the reconstruction of the beta event will have a large error on it due to the small number of PMT hits associated with the low energy beta, and this could smear this fiducial cut.

Although there are some problems, the technique could be very useful. The use of a coincidence measurement eliminates many of the backgrounds associated with the small energy electron signal. The most obvious candidate for the production of spurious low energy betas was at first thought to be the ^{40}K that will be present in the $MgCl_2$ added to the detector. This decay gives electrons with an endpoint of 1.4 MeV, which although smaller than the 1.8 MeV signal expected in β-n coincidence, can contribute due to the poor response of the detector for such low energy events. Studies with the Monte Carlo have shown that the 4 ppm level of potassium in the $MgCl_2$ is too high, but recent results on the extraction of K from $MgCl_2$ using TPB[40] indicate that the reduction of the potassium to the level required should be easily achieved.

Although both direct techniques have problems, they allow the determination of the background within the detector itself, and do not rely on the extraction of the radionuclides. This could be important if the radionuclides are not extractable with the recirculation system. The main advantage of an in situ determination is the demonstrability of the method, in comparison with the slightly esoteric chemistry involved in the extraction and concentration processes.

3.3 Extraction and assay techniques

Of course, if the radioisotopes are to be extracted, it would be useful if they could also be collected in such a fashion that they could then be counted. We can start by examining the characteristic times involved with such a process. The lead isotope, ^{212}Pb, is the relatively longest lived precursor to the major γ ray in the thorium chain. Its concentration level is thus the best indication of the flux of 2.6 MeV gammas in the D_2O. However, its 10 hour half life also means that any counters that are to be used for its assay have to be situated near the mine itself, as any delay in getting the sample to a counter will inevitably lead to a loss in sensitivity. It is for this reason that all extraction and counting techniques which monitor the level of lead have a characteristic time small compared to ten hours. This also reduces the amount of time available for complicated chemical processing.

The next characteristic time is four days. This is the approximate half life of both ^{224}Ra in the thorium chain, and ^{222}Rn in the uranium chain. Immediately preceding these two isotopes are the important long lived nuclides, ^{228}Th and ^{226}Ra. The total extraction volume, ideally the 1 kilotonne of D_2O, will need to be circulated as quickly as possible, as any embedded traces of ^{228}Th and ^{226}Ra will release these four day isotopes over the course of the experiment. To this end, the heavy water in SNO is to be circulated through a water system in the mine drift which will process the whole volume in around 4 days. The turn over time means that such a plant will have to pass 200 liters of water per minute and not introduce activity to such a level as to increase the background.

The extraction will be done by a combination of different systems, all integrated into an overall water system which will allow the different techniques to be used at different times: the exact combination of techniques depends on the detector running mode and the level and type of contamination. The overall water system consists of many parts, with ancillary systems such as heat exchangers, monitors of the isotopic purity of the D_2O and a level control tank, as well as the extractive systems.

The most well known of the extraction systems to be used is Reverse Osmosis. This works by applying an external pressure to counteract the osmotic pressure established across a membrane with very small pores. By the application of this reverse osmotic pressure, there is a concentration of solute molecules on the retentate side, and hence the solvent has impurities removed.

Although the RO system should work to remove the nuclides, it has three problems: firstly, the size of the plant required to filter 200 l/min is so large that its inherent contamination is a major concern. Secondly, the nuclides may stick to the membrane, and not be brought into the reject stream, so making any assay prone to a large systematic error. Thirdly, and most importantly, the RO cannot be used when the $MgCl_2$ is in the detector, as it will tend to remove the salt. What is needed therefore, is an extraction technique which will selectively remove the elements of the Th and U chains, such as radium and thorium, at trace levels, in a 0.2% solution of $MgCl_2$. It is at this point that ion exchangers are needed.

Manganese dioxide, MnO_2, is a fairly well known ion exchanger for the heavier elements. For SNO's purposes, a suspension of particulate MnO_2 is coated onto small acrylic beads; these coated beads are then packed into polypropylene columns across which the water stream is driven. There has been a lot of work on the column geometry and the flow rate across the column, with the current design having very good extraction efficiencies for both ^{226}Ra and ^{224}Ra, but almost no take-up of thorium or lead. The activity in the column is counted by flushing the columns with nitrogen streams, with this gas flow carrying the gaseous radium daughter, radon. The radon is then counted in one of two detectors: the first employs the fact that the radon decay product, polonium, is charged in around 50% of the decays. The radon is contained in a decay chamber held at ground potential, in which is a silicon diode biased at around 1 kV. The charged polonium is thus drifted onto the surface of the diode, which is used to detect the subsequent α-decays. The second counter uses the fact that radon has a freezing point above the boiling point of nitrogen. Two silicon diodes form the sides of a vacuum chamber that is held at 77 K. As the radon passes through the chamber, it freezes onto the surface of the diodes, and again the subsequent α-decays are recorded.

There are a number of problems with the MnO_2 system, but the main difficulty is that it does not extract either thorium or lead: of course, even if it did extract lead, the present counting systems would be of no use, as they can count only radium. It is therefore thought that the MnO_2 and radon counters will prove to be most useful for the purification and monitoring of the ^{226}Ra. Present work is concentrating on the calibration of the extraction and counting systems, both in pure water and in the 0.2% $MgCl_2$ brine solution.

It is obvious that there is a requirement for another ion-exchange system that can remove thorium and lead both in pure water and, more importantly, in the brine solution. After consultation with E. Hooper, a radiochemist at Harwell, the Oxford SNO group began work on hydrous titanium oxide.

3.4 Hydrous titanium oxide.

Hydrous titanium oxide, or HTiO[1], is a mineral ion-exchanger which is used in the nuclear industry for the extraction of heavy ions. The HTiO is formed as a white colloidal suspension by the hydrolysis of acidic titanium salts. This reaction forms particles with a range of sizes[41], with a particular case shown in figure 3.5.

From the figure, it can be seen that the grains show a large variation in size, with the tails reaching 1 micron and 100 microns and the mode at around 10 microns. Not a lot is known about the actual shape of individual particles, and they are usually assumed to be roughly spherical. Although such small particles could cause significant light scattering if they reached the detector, their small size is an advantage in that we have a very large surface area, and hence a large probability of exchange, for a small mass, and hence a small chance of contamination.

The HTiO is used in a technique known as seeded ultra-filtration, or SUF[42]. In this, a certain amount of this suspension is deposited onto the inner surface of a plastic hollow fiber which has radially directed pores, these pores having diameters of the order of nanometers. The plastics used for these hollow fibres are usually cellulose nitrate, or polysulphone. An electron micrograph of a fibre is shown in figure 3.6.

Here, we can see the differentiation of the pore size. The inner wall of the hollow fibre has the very small holes, of order 10nm, and these extend radially for only a

[1]It is important to note that hydrous titanium oxide is a non-stoichiometric compound, and hence the chemical formula HTiO is used only as a representational device.

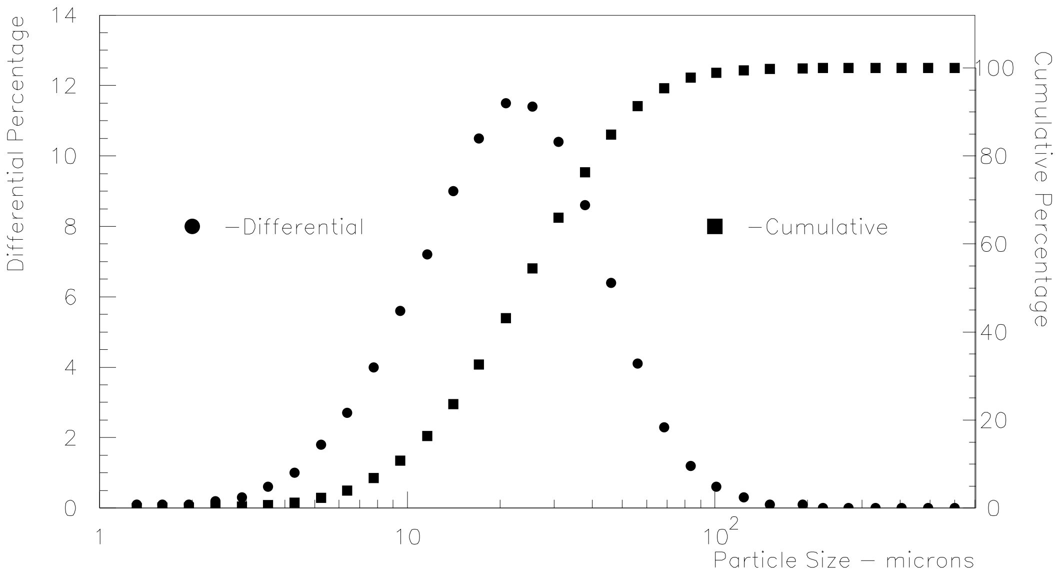

Figure 3.5: A typical particle size distribution for HTiO grains. The measurement was carried out on a sample prepared by the standard technique to be used in SNO.

fraction of the fibre depth. The fibre body is made up of an open, macrocellular structure with diameters of the order of 1μ. Such a structure is obviously able to take a large pressure going from the inside to the outside, while being much more prone to collapse in the opposite direction. Bundles of around 50 of these fibres are sealed in epoxy, with both ends of the fibres open to the retentate water stream. As the pressure on the retentate side is increased, water is forced through the pores to the permeate side. The HTiO, the seed in seeded ultrafiltration (SUF), is put onto the fibre by recirculating a suspension of the material through the fibres. The process volume is then pushed through the pores and the the ions collect on the ion exchanger grains as the water flows through the pores.

There are a number of papers on the use of HTiO[43], although most of these refer to the use of HTiO in the solid form that it produces when the particulate suspension is dried. The solid form does show good extraction efficiencies, but the kinetics of the cation's diffusion into the solid mean that the contact time between the process volume and the ion exchanger is usually of the order of days. Such a process is obviously not applicable to SNO, where a fast recirculation of the process volume is required.

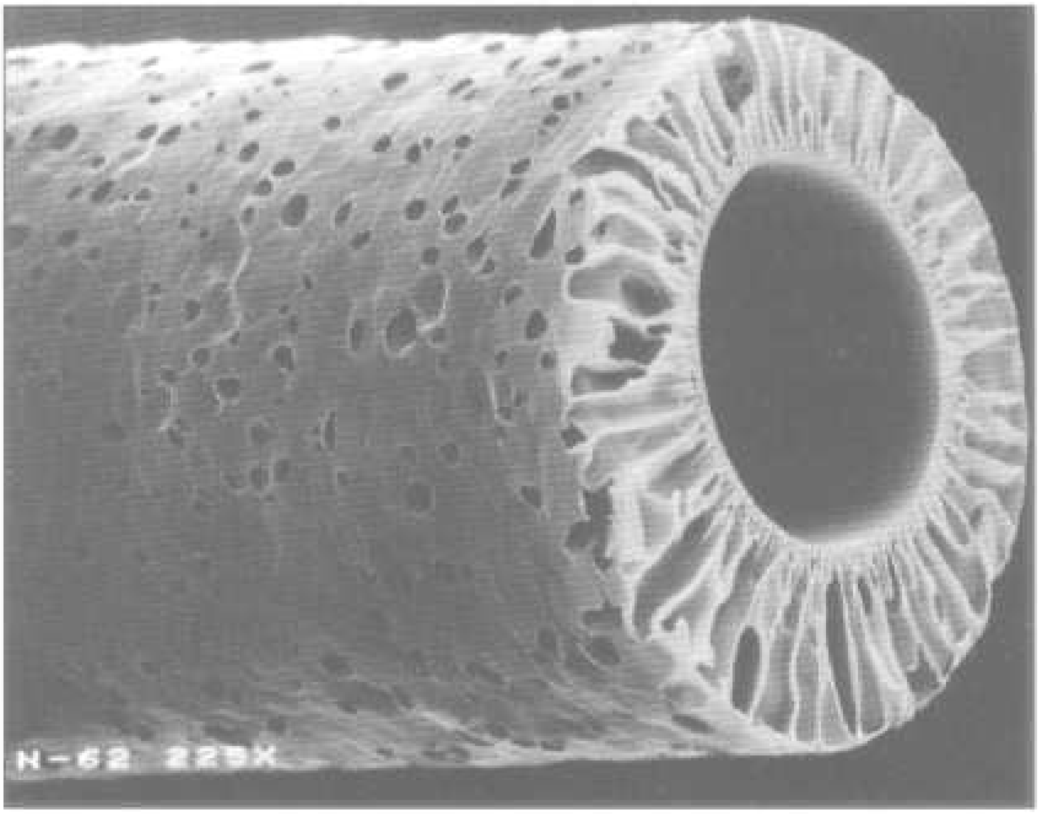

Figure 3.6: An electron micrograph of a typical hollow fibre. The fibre outer diameter is around 2 mm, while the inner diameter is 1 mm. See text for a detailed description.

However, these papers do demonstrate the high affinity for heavy ions that HTiO possess, with one particular paper[44] detailing the extraction of uranium from seawater. Although most of these processes are not at the sort of activity levels required by SNO, there is a degree of encouragement to be taken from the work.

Another advantage of HTiO is the ease with which the extracted nuclides can be washed off the ion exchanger with dilute mineral acids, a process known as elution, with the acid volume being referred to as the eluate. This means that large volume reductions can be achieved by washing the membranes with acid when the membrane is taken off-line: the acid volume is determined only by the volume of the membrane itself, and not by the process volume, and volume reductions of the order of 10^5 are expected for the system envisaged for the detector. The elution usually brings approximately 80 to 90% of the captured activity back into solution. If this acid is neutralized and then passed through a smaller membrane, the volume can be reduced to almost any figure, with a consequent drop in overall efficiency. This elution process also removes some of the HTiO as well, meaning that the membrane cannot be eluted and then immediately reused, but further acid washes can remove the remaining ion exchanger.

There are only two real problems with HTiO: firstly, both the membranes and the HTiO have to be kept wet during production and use. The membranes tend to degrade when dried, while the HTiO forms the solid form when the water is driven off. The chemical processing involved in HTiO is planned to be done in light water, so each time a membrane is prepared for use in SNO, it has to be deuterated by a treatment with heavy water. This could cause serious problems in that if there is too much hydrogen present in the detector, not only does it capture the neutrons produced in the neutral current dissociation, but the heavy water has to be given back to the owners if it falls below a certain isotopic purity level.

The second problem is that this increase in chemical processing brings with it requirements on the radiopurity of all the acids, alkalis and salts used, as well as the activity of the membranes and ancillary plumbing used for the system. This is much more of a problem than the MnO_2, as the processing stages for this are very simple indeed.

Hydrous titanium oxide - assay techniques

The membranes that are proposed for use in the 200 litres per minute plant for the SNO detector have to be eluted with around 20 litres of acid, with two acid washes currently being envisaged. The first will be a comparatively weak nitric acid elution which will remove the radium, followed by a stronger hydrochloric elution which should remove the thorium and lead. It is these acid eluates which will be the focus of any counting procedure.

Previous to the work in the next chapter, the only counting method for these eluates was another stage of concentration, and then the drying of a second eluate onto a steel planchette. This was then put into an alpha counter consisting of a silicon photodiode. This does have the advantage that if the energy resolution is good enough, the various isotopes can be separated. However, the elution process does involve the dissolution of some of the HTiO, with the subsequent appearance of titanium ions in the acid solution. When such an eluate is dried on the planchette, it forms titanium salts which then scatter the α-particles. This affects the measured energy to such an extent as to make alpha counting, and separation of the peaks, very difficult for these eluates. There were proposals to use an electro-plating technique which puts the radioisotopes onto the planchette, which is used as one of the electrodes, but this process is time consuming and prone to contamination.

It was this slightly odd situation which faced the author. Although HTiO offers the best extraction properties for a range of the ions in the thorium and uranium chains, there was no reliable counting method. The highly selective nature of HTiO offers the only method for the cleaning of the $MgCl_2$ to be used as the neutral current additive, and a measurement of the activity of the two tonnes of solid to be used is vital. A more reliable and robust counting method was required. The development of new counters is detailed in the next chapter.

Chapter 4

Low level radionuclide counting in SNO

There are many things we need to consider when developing a counting system which is linked to the extraction processes proposed for SNO. First, it should count as many of the important isotopes as possible; it should be free of background and not liable to contamination; it should have as high an efficiency as possible; and finally, the counting system should be able to count many different samples simultaneously: this imposes limits on the cost and ease of use of the counter. We will go on to detail the production of a new counter developed by the author which meets all these requirements.

4.1 Source disequilibria.

There are two tasks that have to be performed by an extractive systems in the water plant in SNO. The first is the extraction of the long lived isotopes ^{226}Ra in the uranium chain and ^{228}Th in the thorium chain. Such an extraction will gradually clean up the detector over the course of time and thus reduce the rate of gammas of >2.2 MeV, which give the background signal in the neutral current channel. The second task is the measurement of the rate of these gammas, in order to allow a background subtraction to be made on the observed rate of neutron production. In order to make this second measurement, it would be better to count the activities of those isotopes nearer the gamma decaying radionuclides. This becomes important when the long lived isotopes are in a form which renders them inaccessible to the extracting systems, but where they can still produce a steady stream of mobile daughters which can reach the acrylic vessel. It is for this reason that 'chemical' techniques, such as ICP-MS, are of little use to us, as there can be significant disequilibria between the long lived parents, and the gamma-producing daughters. Thus, a radioactive counting technique must be employed.

Considering only the thorium chain to begin with, this means that as well as measuring the ^{228}Th, we have to measure the activity level of the daughters, ^{224}Ra and ^{212}Pb. This can be done in one of two ways: we can either measure the level of all three directly, or count an isotope below all three, and by monitoring the time dependance of the count rate, fit the resulting time spectrum with competing exponentials. In the thorium chain therefore, we can measure the decay of ^{212}Pb, or the decay of any of the isotopes below it. Looking below lead, the most striking feature is the very quick α-decay of ^{212}Po. Such a short half-life immediately suggests that a coincidence between the β-decay of ^{212}Bi and this alpha decay could be very useful indeed. In such low rate counting experiments, it is usually a very good idea to use a coincidence as, for a random, uncorrelated rate of R, the rate of background counts, B, in a single counting experiment is:

$$B = R$$

if no other cuts are applied. For a coincidence with a characteristic time τ:

$$B = R^2\tau$$

and hence if Rτ is less than one, the background is reduced. So, for a given random rate, the smaller τ is, the greater the reduction of the background will be.

4.2 The β-α coincidence

Now, if we are to use the very quick (300ns) decay of ^{212}Po, then we need a counter that is both sensitive to the beta and the alpha. In order to be sensitive to alphas, we need to either deposit the alpha-bearing material as a thin film, or to mix the alphas with some sensitive media. In the thin-film approach, β-α coincidences can be seen in silicon diodes, when the beta's pass through the silicon and deposit a fraction of their energy in the diode[45]. This produces counts which lie above the high energy edge of the alpha particle peak, although the efficiency of such a process is very low, as both the alpha and the beta have to impinge on the diode.

Another technique which actually uses the same β-α coincidence is detailed in [46]. In this experiment, the solid sample for counting is viewed by two photomultipliers: one of which has a disc of plastic scintillator on the front window, while the other has a ZnS screen attached to the front window. The two PMT's are optically separated from each other by a sheet of black plastic between the sample and the plastic scintillator. The β produces light in the plastic scintillator, but not the ZnS, while the α stops in the plastic used for the optical isolation, and produces a flash of scintillation light in the ZnS screen. The two photomultipliers are thus labelled as Beta and Alpha, and a signal from 'Beta' then another signal from 'Alpha' within the gate time of 700 ns indicates a β-α coincidence, and hence the decay of ^{212}Bi. However, the counting efficiency for such a set-up is quite low (with a theoretical maximum of 25% due to the isotropic nature of the β's and α's), and the paper also details problems with muons passing through both the ZnS and plastic scintillator producing prompt (<30 ns) coincidences.

Rather than using this thin film approach, which had already been shown to be a difficult procedure for the samples generated in HTiO eluates, it was thought that the better method was to mix the sample containing the alphas with a sensitive medium. The obvious choice for such a mixing is a liquid scintillator[47]. This has the added advantage that modern liquid scintillators can be formulated to accept aqueous samples, such as those produced by eluting SUF membranes with acid. The water mixes into very small agglomerations from which the alphas can travel out without losing too much energy. The other driving force behind liquid scintillators is their response time when coupled to photomultipliers. The speed of the alpha decay does reduce background, but it also means that the response of the counter has to be quick enough to resolve the two signals.

Preliminary tests

Tests of the method began using a toluene-based liquid scintillator in a glass jar the bottom of which had been ground to an optical flat. Around 100 millilitres of the scintillator were mixed with around 10 Becquerels of ^{228}Th in aqueous solution. The scintillation source was then optically coupled, using a silicone grease, onto a 3 inch flat face PMT and used for a test of the pulse heights and the timing response. The basic electronic circuit for the conditioning of the pulse train is shown in figure 4.1. This involved the use of two discriminators: both had the PMT signal put into them, the signal having been split passively, and one signal delayed with respect to the other. The beta discriminator, which had the delayed signal put into it, was used for three specific functions: to start the timing gate of 1.5μs duration; to start the TDC; and to gate an ADC to record the pulse height of the start pulse. The output of the undelayed alpha discriminator was put into a coincidence unit with this timing gate, and the output of the 'AND' coincidence of the two signals stopped the TDC and gated an ADC which recorded the pulse height of the alpha. The delay introduced between the two discriminators stops them firing on the same pulse and so has to be longer than the output pulses of the discriminators. This meant it was usually of the order of 20 ns, and hence not too many counts were lost. The timing sequence is shown in figure 4.2 also.

This system was used to record the pulse heights and the time spectrum for the ^{228}Th source. This data is shown in figure 4.3, with an exponential fit overlaid on the TDC spectrum.

Some interesting results were seen even with such a primitive setup. The time spectrum was fitted and the half-life was measured at around 290ns, with the discrepancy explained by the calibration of the TDC. Also, the outputs of the ADC's were observed to follow roughly the correct shapes. The beta pulses show an approximate correspondence to the expected first forbidden energy spectrum, while the alpha pulses are peaked, with the relative energy calibration showing that this alpha peak occurred at around 1 MeV equivalent electron energy, this being consistent with the usual 10 to 1 quenching factor for alphas[48]. This alpha peak (at around 600 channels) was measured to lie at around 100 photoelectrons, and hence any problems with the PMT dark current were not important.

Another run, not shown here, looked at the energy spectrum of events due to room gammas scattering off the electrons in the scintillator. The spectrum roughly followed an inverse power law, which meant that a significant background rejection could be achieved by increasing the lower threshold for the alpha discriminator. Such a procedure could obviously not be applied to the beta discriminator however, since the energy spectrum of the beta's extends down to zero.

The only problem with this preliminary set up was cosmic ray muons passing through the scintillator and producing such large signals that the discriminators could refire and produce coincidences with zero time separation. This meant that a cosmic

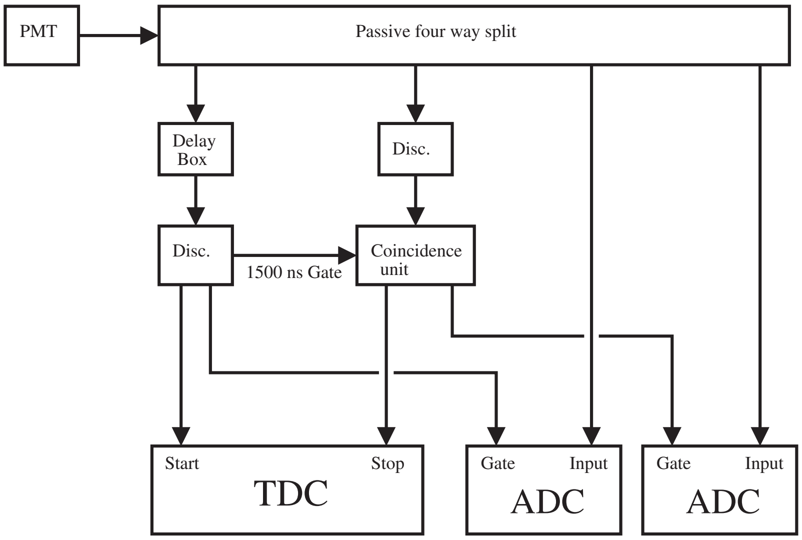

Figure 4.1: A basic electronic setup for the observation of coincidence. The use of two discriminators allows two different thresholds to be set, while the gating of an ADC with the coincident pulse means that a measurement of pulse heights of the beta and alpha pulses is achieved, as well as the time difference between the pulses.

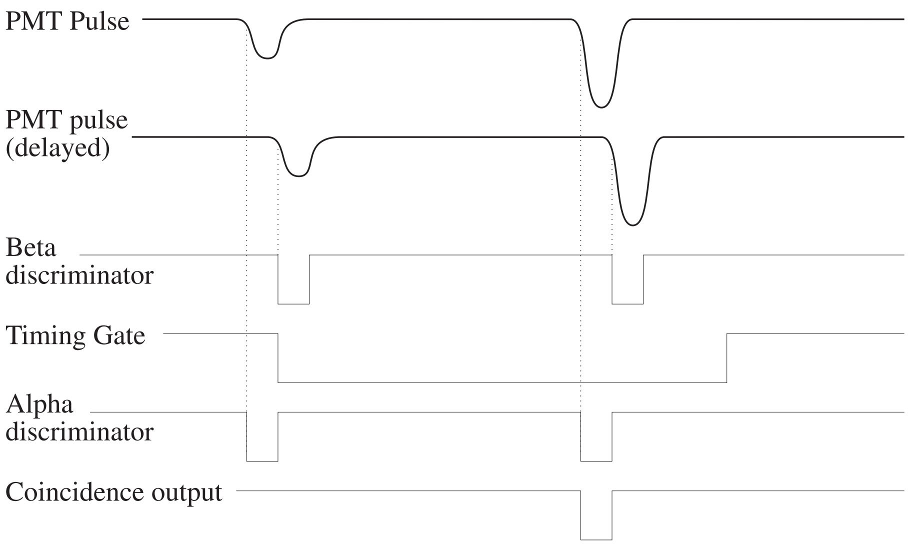

Figure 4.2: The timing diagram for the observation of β-α coincidence. The two discriminators are the ones shown in the electronic diagram above.

veto was introduced, this essentially comprising an upper limit on the beta discriminator, with the system being time-vetoed for 1 μs after such an event.

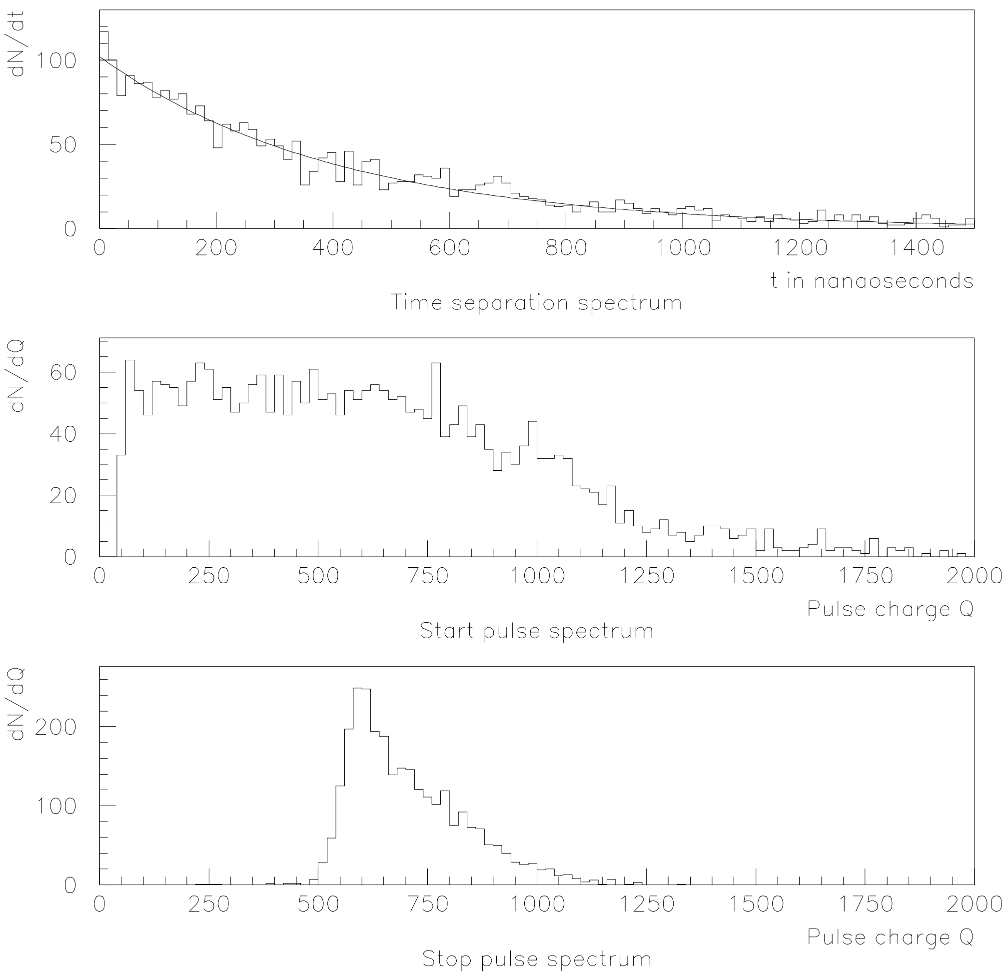

Figure 4.3: The time separation, start pulse and stop pulse spectra for the β-α coincidence. Note the time fit to the time separation, giving a measured half-life of 290 ns, and the low level discriminator cut-off on the start pulse.

Development work

Although this test system was not optimised, it did show the power of the technique, in that the measured efficiency was around 40%, while the background was of the order of 10 counts per day. Numerous developments were obvious, including better shielding and cleaner materials for construction. However, there were some limitations on the progress as well.

The main problem with the method at this point was the use of a glass jar to hold the scintillator. This was necessary due to the powerful solvent characteristics of the toluene, on which the scintillator was based, meaning that plastic jars could not be used. Usually, glass is avoided for low background work due to its inherent radioactive contamination. The use of specifically clean jars, for example clean quartz, was seen as impractical, as their high cost would mean that they would have to be used more than once. This meant that they would have to be cleaned, and such cleaning processes are obviously prone to contamination. Obviously, cheaper plastic jars that could be used once and then discarded would be far better.

In order to move away from glass, a solvent other than toluene was needed, and after some consultation with the manufacturers of scintillation cocktails, trials began on new scintillators which have as their base di-isopropyl napthalene, hereafter referred to as DIN. This has a very limited attack on the weakest plastics, and is perfectly stable in such transparent plastics as polycarbonate (PC) and poly-methyl pentene (PMP).

This was a great step forward for β-α coincidence. The use of plastic jars brought with it two major improvements. Firstly, the radiopurity of the plastic meant that backgrounds from the jars themselves could be all but ignored, and the lower cost of these jars meant that they could be used once and then discarded. This obviated the lengthy cleaning and decontamination process that was required for the glass jars. It also introduced one of the great advantages of the β-α counters, namely source/counter separation. Any contamination of the counter only occurs while the source is being prepared, and the dirtiest sample can be counted, and then immediately replaced with a very clean sample without any lengthy cleaning process being required. With the introduction of these plastic jars, the time had come to design and build a complete system. This is detailed below.

The two inch PMT system

The production of the complete system began with the selection of the PMT and jar combination. We settled on a two inch PMT, onto which could be mounted a standard 60 ml plastic jar. Such a jar allowed 12 ml of an aqueous sample to be counted with 40 ml of the liquid scintillator. The ratio of aqueous sample to scintillator gives clear sources; a greater ratio, i.e. more aqueous, does still produce sources that scintillate to an acceptable level, but it was found that there were large variations in pulse height for apparently similar water/scintillator ratios.

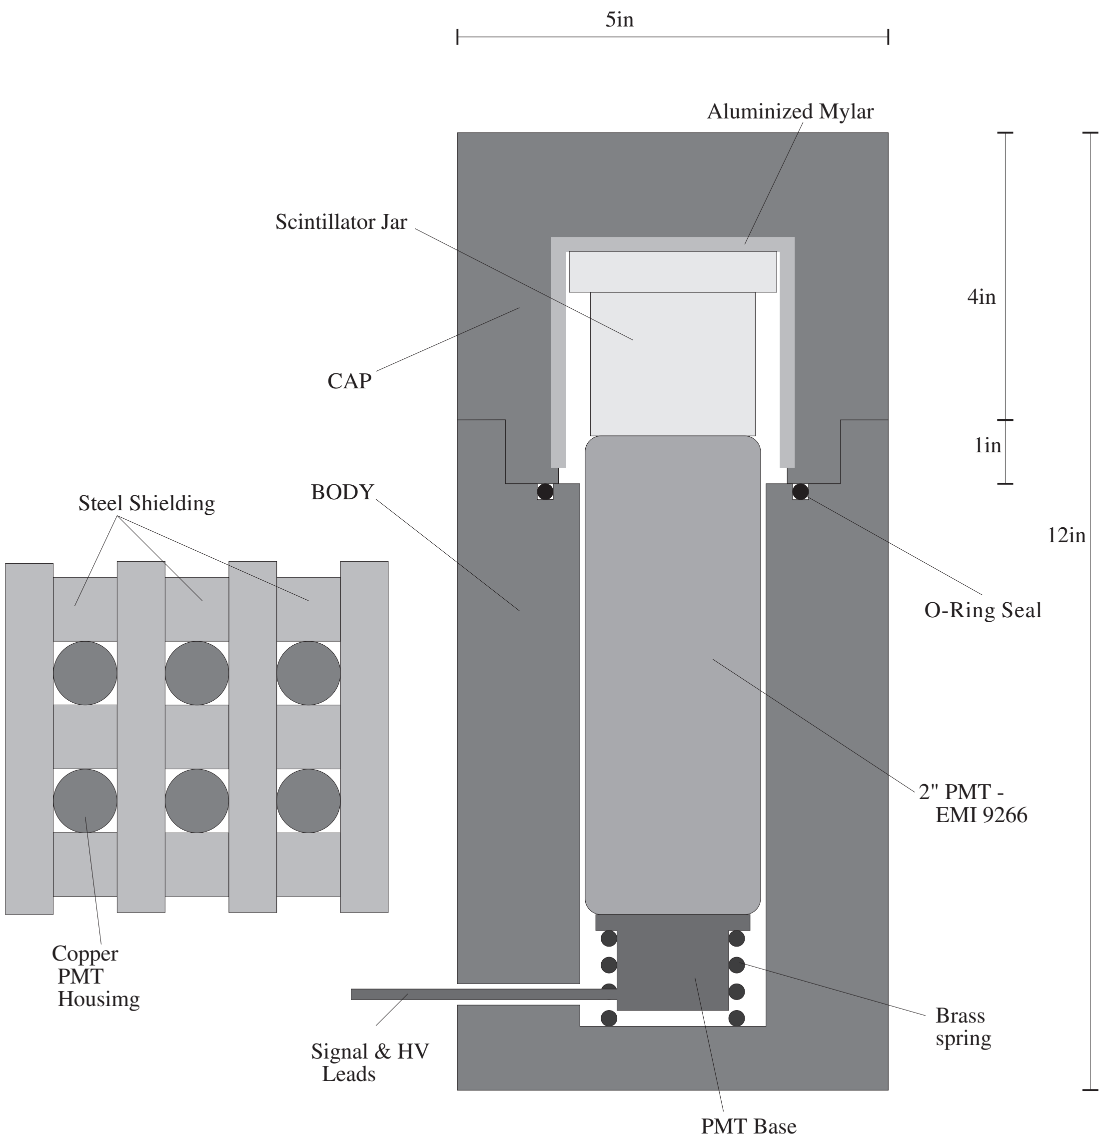

Figure 4.4: A schematic diagram of the two inch β-α counter. The approximate dimensions of the copper housing are shown, as is the plan view of the six counters in the steel shielding.

It was also thought necessary that the counters be shielded against soft room gammas by the use of passive metal shielding. The shielding consists of one inch of oxygen-free, high conductivity (OFHC) copper. This grade of metal is basically a high purity variant. Around this was placed 4 inches of mild steel for further shielding. The arrangement of PMT, jar and shielding is shown in figure 4.4.

As can be seen, the counters were arranged to give a degree of modularity in their design. Also, there is an O-ring under compression between the cap and the body in order to keep the inside of the copper light tight. The two inch PMT is an EMI 9266XB. This is made from special low background glass to reduce the rate of singles in the scintillator. The PMT is spring loaded via a lip on the base so as to keep the jar/PMT join under compression, while the aluminised mylar was used to improve the capture probability for the photons, although it has been subsequently observed that it has almost no effect whatsoever. The system currently in use by the water group at Oxford consists of six of these counters as shown in the figure.

Pulse shape discrimination

It was with this system that the first tests of a discrimination method using the shapes of the pulses were carried out. Pulse shape discrimination (PSD) method is a well known technique in α/liquid scintillator counting [49]. The α particle sets up a larger ratio of triplet to singlet fluorescence states than an electron, and when two of these triplet states react, they can produce an excited singlet and a ground state molecule. This reaction is diffusion limited, and hence occurs more slowly than the normal singlet de-excitation. The scintillator light output therefore consists of a fast component and a slow component. The proportions of these two components is systematically different for different initial ionising radiations, and hence allows a separation between β's and α's.

In terms of the signal from the PMT, the pulse shape differences distort the falling edge of the electronic pulse to an extent that allows pulse by pulse discrimination between betas and alphas. This was viewed with some enthusiasm as the non-correlated single pulses are almost entirely electron events: the short range of alphas means that any generated outside the plastic jar will be ranged out, while the great majority of non-correlated alphas produced within the liquid volume will come from decays within the thorium and uranium chains themselves. If the activity is high enough to produce a large number of these uncorrelated alphas, the rate of real β-α coincidences would swamp any random β-α coincidences.

The first step in quantifying the difference between beta and alpha pulses was to look at the signals themselves. This was done by digitizing the output of the PMT with a Hewlett Packard 54512B digital oscilloscope. This takes 1 giga-samples per second, this being equivalent to a 1 ns bin. The trigger for the digitisation was set to give well separated betas and alphas. Two thousand such pulse shape samples were taken, and

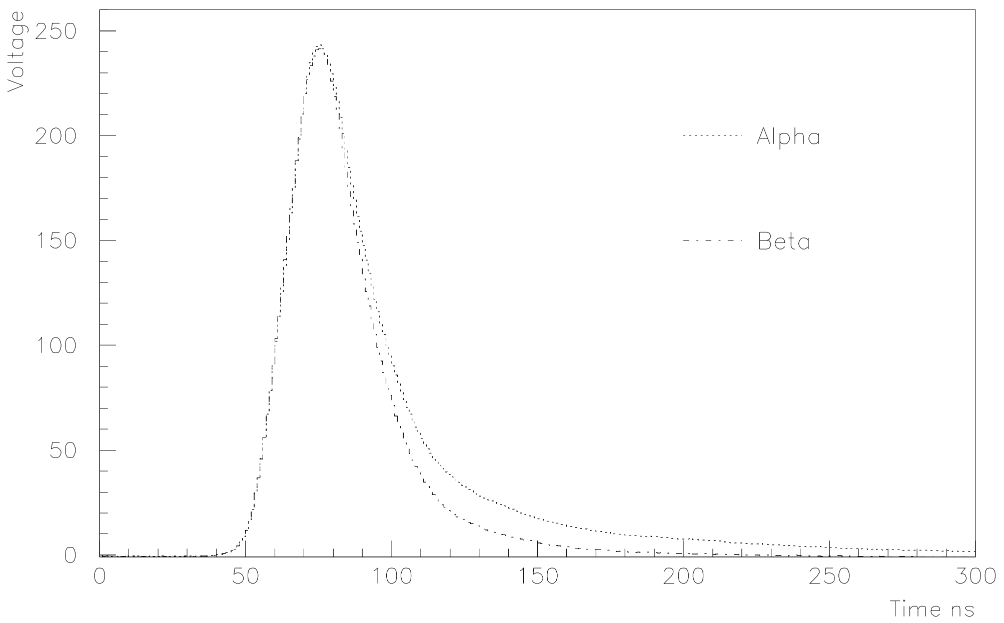

Figure 4.5: The average pulse shapes for beta and alphas. The distributions are the average of two thousand coincidences which were digitized at 1 giga-sample per second. The pulses have been shifted in time so that they cross at the 60 ns point.

the averages of these samples are shown in figure 4.5

These pulse shapes were formed by first separating the beta and alpha pulses for each coincidence. In order to add the samples together properly, it was necessary to cut and paste the data from each pulse in time so that they all crossed a certain point at 10% of their peak height. The forms are fairly similar over the majority of the pulse, but where they really do differ is in their tails. The alpha is much higher over the range between 100ns and 300ns, which was the maximum digitisation time. In order to quantify this difference, many different formulae were suggested, but the most obvious was to measure the charge in the tail of the pulse, between 100ns and 250ns. This however is not a good measure as the betas have a spectrum of heights, while the alphas constitute a relatively narrow peak. Thus, the simple formula of charge in the tail divided by the total charge in the pulse was used. A scatter plot of the two quantities is shown in figures 4.6.

Looking at the figure 4.6 it can be seen that the two decays do seem to follow

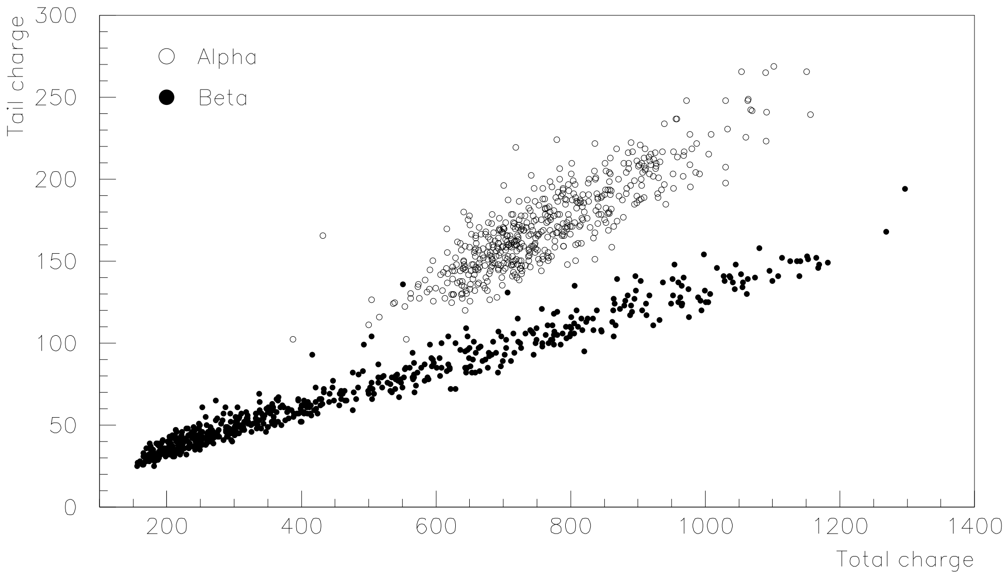

Figure 4.6: A scatter plot of tail against total charge. Note the suppressed zero on the x-axis corresponding to the lower level discriminator on the beta pulse.

different slopes (i.e. different values of the ratio), although the lower energy betas conform slightly less well. This discrepancy can be explained by two effects: either PMT dark current is adding to the tail, and has a bigger effect when the main pulse is so small, or radio noise on the coaxial lines is smearing the tail charge so that the statistical variation pushes some of the tail over charge ratios upwards. Although it should be possible to subtract suitably sampled average noise rates from the pulses, the statistical nature of the noise means that, for any particular small beta pulse, there could be a large enough number of noise counts so as to 'infect' the tail to a level that the pulse is mis-classified.

Of course, such low energy electron type pulses can be discounted in the stop pulse thanks to the total charge lower cut on the alpha pulses. If we select the beta events to lie within the alpha energy window, then when we divide the tail charge by the total charge, we obtain the distributions shown if figure 4.7, where there is a clear separation of the two decays, with only a few percent overlap.

In order to use such a system of separation in the counters, we could follow one of two paths. Either we could digitize all the pulses, or we could measure the peak

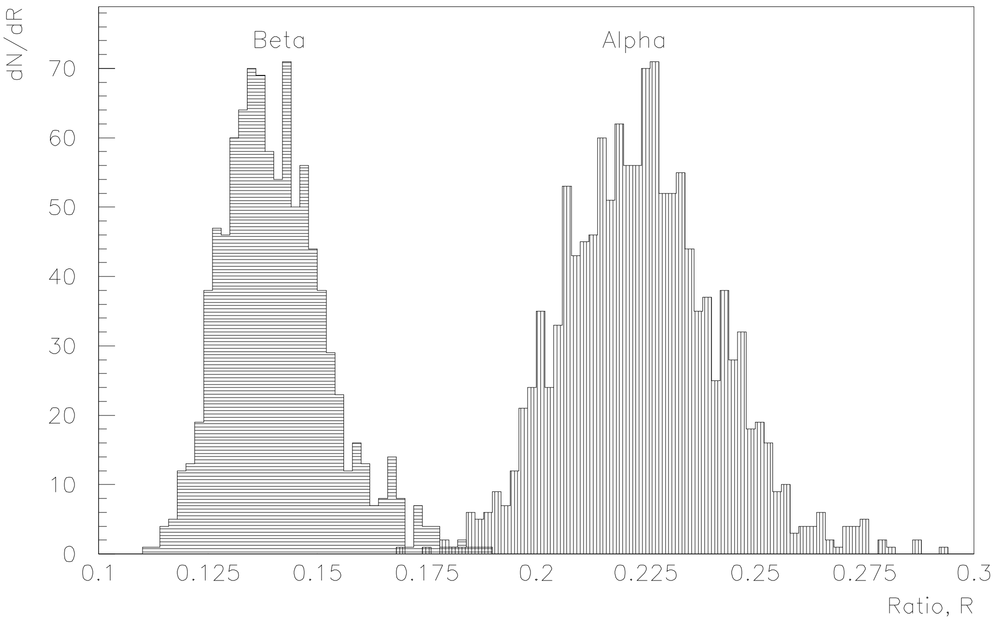

Figure 4.7: The distributions of the ratio of tail charge to total charge for betas and alphas. The beta events have been selected so that their full pulse height lies within the range of the full alpha pulse height. The reasons for such a cut are described in the text.

and tail charge with two ADC's gated on different parts of the pulse train. The first method was seen as fairly expensive and also fairly slow in terms of the transfer of the pulses from the scope to a host computer, and work concentrated on whether the ADC's could be made to be sufficiently discriminating.

The first problem is that it is useless to use a two ADC system to measure the tail charge over total charge ratio for the first pulse of the coincidence, since if we measure the charge for up to 250ns after the main body of the pulse, then any β-α event with a time separation of less than 250ns will include some of the alpha pulse within the beta tail, leading to an increase of the measured tail charge for the beta. However, we do not need PSD on the first pulse, as it is generated by an electron in both the real β-α case and in the random γ-γ case.

The ADC system therefore was set up to measure the PSD ratio only for the second pulse. In order to calibrate such a system, one needs a source which can produce a large number of stop pulses that are electron events. Such a source was made by using around 200 Bq of ^{90}Sr. This high activity means that we induce a high rate of random beta-beta coincidences. The PSD ratio for these stop pulses could then be compared to the PSD ratio for stop pulses from a ^{228}Th source, and hence we can look at the separation for various gate timings.

An important point to note about the ^{90}Sr source is that the larger of the two beta decays, at 2.27 MeV, encompasses the charge range of the alpha decay. This meant that the strontium source could reproduce those random gamma events that could occur within the alpha charge window. In order for this source to fire a coincidence, the coincidence window was increased to 700μs, this gate time having already been set up to look at the uranium chain coincidences. The large rate of strontium decays meant that random coincidences were seen at a rate of around 1 per second.

When these PSD distribution were compared to the digitisation distributions, it was seen that the degree of separation was almost exactly the same, and this two ADC technique of pulse shape discrimination is used in the current counters. Perhaps the only reservation is that the cuts on the pulse shapes are thus hardwired into the acquisition program. Having said this, the stability of the system over its use for around 1 year now is such that no adjustment has been needed to these cuts.

Uranium chain counting

It is at this point that mention must be made of using the counters for activity measurements of the uranium chain. Looking at the uranium chain, it can be seen that a β-α coincidence occurs in the transitions between ^{214}Bi, ^{214}Po and ^{210}Pb. Also, although the branching ratio to the coincidence is higher, the half life is around 500 times as long. This of course leads to an increase in the random background rate of 500 but otherwise there are no inherent differences, and hence no reason why the uranium chain activity cannot be found using the coincidence.

Although tests of the counters were made, and indeed seemed to be quite positive, the backgrounds for the uranium chain counting were much higher than the thorium chain. This came about through a number of reasons: the α energy is lower, and hence the background rate within the alpha energy window is slightly higher. This lower energy also produces a greater overlap between the alpha and beta distributions in the PSD ratio plots, and hence there is a greater chance of mis-classification. Perhaps the most obvious problem was that even very clean samples were seen to have a background which grew in exponentially with a characteristic time of 4 days. This is attributed to the mobile daughter of ^{226}Ra, ^{222}Rn, diffusing into the scintillator from the surrounding materials.

The high background in the system means that the system would not be good enough to measure the level of ^{214}Bi to the required accuracy for SNO. The level of ^{226}Ra in the detector will be measured to a better degree by the MnO_2 system detailed in chapter 3, so that this poor performance should not adversely affect the applicability of the overall HTiO system. At present, the counters do still measure the rate of ^{214}Bi decays for a particular reason. The half-life of 164 μs means that, on average, 1% of the uranium decays will occur within the 1.5 μs gate used to detect thorium chain coincidences. Thus, the bismuth activity is required, although not to a great degree of accuracy, in order to subtract this cross-talk between coincidences.

4.3 Counter Performance

The counting system has been used extensively for the investigation of the parameters involved in seeded ultra filtration, and has been calibrated for this use. The efficiency has been measured with standard radioactive sources and is 50% for thorium activity. With the branching ratio to the β-α sequence in the thorium chain given as 64%[39], the efficiency for detecting a beta followed by an alpha is thus 80%.

This figure of 80% detection efficiency fairly reasonable, with perhaps a 10% loss due to the low level β cut, and 10% in the timing. This first loss is due to an imposition of a low level discrimination on the beta signal at around 100 channels. The main body of the beta spectrum is roughly a flat distribution out to 1200 channels, with a decreasing tail out to 1500 channels. Assuming the beta spectrum is flat from channel 100 down to channel 0, this gives around a 10% loss in counts.

The loss due to timing occurs because the full range of the exponential is not included. If we lose the first 30 ns due to the imposed delay between the two discriminators, and we say that the signals have to be within 1500 ns, or 5 half lives, of each other then the efficiency, E, is:

$$E = \frac{\int_{30}^{1500} e^{-\frac{t}{t_m}} dt}{\int_0^{\infty} e^{-\frac{t}{t_m}} dt}$$

$$\Rightarrow E = \left[e^{-\frac{30}{t_m}} - e^{-\frac{1500}{t_m}}\right]$$

$$t_m = \frac{300ns}{\ln 2}$$
$$\Rightarrow E = 90.2\%$$

As for backgrounds, we have not fully pinned down some of the more troublesome points. In the thorium chain, the background rate has been measured consistently at approximately half a count per day. This is thought to come from contamination of the scintillator and perhaps some radionuclides in the surface layer of the plastic which are leached out. All the indications would suggest that these are true β-α coincidences, and although it would be nice to reduce this rate, perhaps by redistilling the scintillator materials and other such processes, it is seen as both difficult, and to all intents and purposes unnecessary for the required levels: if the extraction and elution system works at a conservative estimate of 30% efficiency, and the smallest volume for the extraction is of the order of 100 tonnes of D_2O, then the decays per day should be around 30, and hence around 15 counts per day in the counters in the limit of 10^{-15} g/g of thorium in the detector.

Another property of the scintillation sources which has been noted is their stability with time. This is important as the only way to measure the ^{228}Th activity is to allow its daughters to grow into equilibrium, which takes of order 2 weeks. Hence the scintillation source should not change in any way over this time. The most obvious difficulty is the possibility that the scintillator 'goes-off' in some sense and gives smaller peak heights. Sources have been seen to count at the true rate over the course of a few months, although the sources themselves do have to be kept in the dark during this time, as exposure to UV light ages the scintillator, and hence reduces the peak heights. This is a problem as, at present, the cuts applied to the charge spectra are hard-wired in the acquisition program. It should be possible to calibrate the detectors externally using a γ-ray source and look for the Compton edge in the singles spectrum, and hence factor-out any significant changes in pulse height with time. Were an alpha calibration needed, it may also be possible to add low levels of ^{241}Am, although it would be important to check any such additive for the levels of thorium and uranium present. Perhaps the most persuasive calibration would be the addition of a few mBq of ^{228}Th to the source at the end of the counting time, this giving a direct check on the efficiency.

4.3.1 Time spectra and fitting.

As mentioned above, if we count at one particular point in the chain, then we can estimate the relative activities of those isotopes above this decay by allowing the chain to grow into equilibrium. In the case of β-α coincidence, we really want to look at the relative proportions of ^{212}Pb, ^{224}Ra and ^{228}Th. This can be done by recording the number of counts per hour over the course of around 2 weeks or more. During this time, the time spectrum will vary due to the three competing exponentials of these three decays. A sample time spectrum, with a fit from the procedure detailed below, is shown in figure 4.8.

The first point to make is about the gaps in the histogram. These are regions where the sources were not counted, and such gaps are not included in the fit. This has no effect on the accuracy of the fit as Monte Carlo calculations have shown that the majority of the information which the fit uses is in the first two or three days. This sets the levels of lead and radium, and counting the source after a few weeks sets the level of thorium. The fit itself is simply three exponentials with the coefficients able to be related to the activities at t=0. The algebra for the fit is shown in appendix B.

The real work in this procedure was done on the use of a log likelihood method for the fitting[50]. A simple Gaussian-distributed bin height fit is inappropriate as the spectrum is divided into bins of an equal time of one hour in order to accurately describe the 15 hour ^{214}Pb. However, with a lower limit on the rate of the order of 1 count per day, there can be many bins with no events in them. Thus, a log likelihood method with Poisson distributed bin heights must be employed. The fitting program has been developed by the author, and seems to be fairly robust, and has been checked with Monte Carlo generated time spectra. The errors on the fitted parameters are found by tracing the contour where the log-likelihood falls by a prescribed amount, and for low activity levels, severe deviations from simple Gaussian approximations have been observed.

The final conclusion is that we now have a counting technique for the eluates from HTiO SUF procedures which has a high efficiency and a low enough background for use in the SNO detector. These counters have also allowed previously excluded regions of the extraction procedure to be studied, and some of this work is shown in the next chapter.

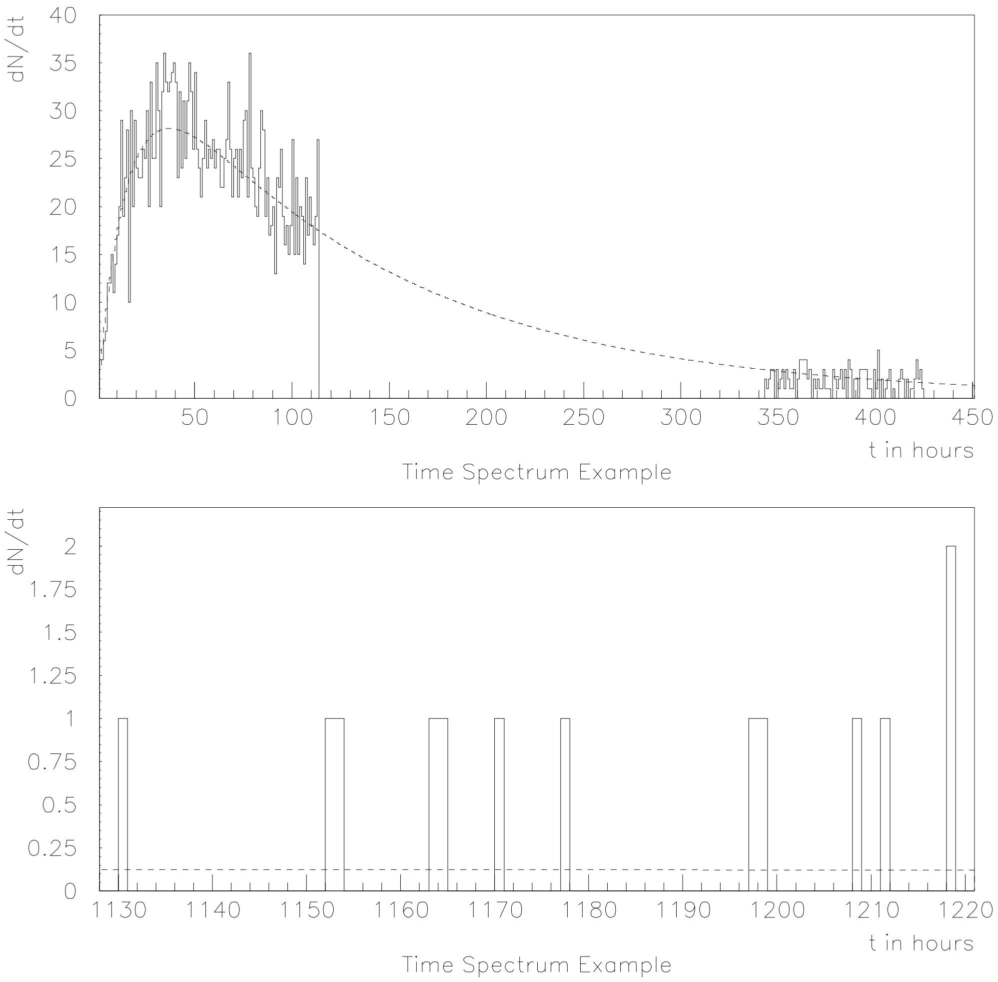

Figure 4.8: A sample time spectrum with a fit superimposed. The time is recorded in the number of hours after source production, and in this case the source has a high level of radium, with little lead and thorium.

Chapter 5

The application of Seeded Ultra Filtration.

5.1 Previous work.

The Oxford SNO group has been using HTiO for some time now and a detailed study of the properties of the ion exchanger was carried out by Katherine Howard[51]. Although other papers[43] have been written on the performance and properties of HTiO, they usually contain information only on the dried form of the material, which is unsuitable for use in SNO due to its slow take-up of cations. The main conclusions of the Howard thesis were on the synthesis of the particulate suspensions of HTiO, with a standard procedure being established, and the changes in the cation exchange capacity (CEC) in different strength salt solutions. These latter tests indicated that the particularly selective nature of HTiO lead to only around a 50% fall in CEC for 5% $MgCl_2$ solutions, with virtually no change in CEC for a 0.2% solution. The CEC was measured using γ-ray counting of a ^{133}Ba source: this demonstrated the selectivity between magnesium and barium, and it was postulated that the separation between magnesium and radium (and lead and thorium) would be the same, if not better.

When the HTiO was first used in basic tests of SUF, it was demonstrated that it could extract the thorium, radium and lead isotopes in the thorium chain from pure water volumes[52]. Further work in [53] and by the author showed that all three nuclides captured on HTiO could be removed from the exchanger using different strengths of mineral acid. The tasks facing the group at this point were to demonstrate that all these pieces of work could be put together into a coherent whole.

The experiments therefore required were scale tests of SUF to demonstrate that the

Membrane name	Membrane Area m^2	Pore Size	Trans-membrane pressure required for 1 l/min flow
H1-MP01-43	0.015	0.1 μm	10 psi
H1-P30-20	0.03	30 kiloDalton	30 psi

Table 5.1: A comparison of the two membranes used in the scale tests. A pore size of 30 kiloDalton is roughly equivalent to 5 nm.

technique would work both with large process volumes and at the high flow rates (and hence small contact times between the radionuclides and the ion-exchanger) required in the detector system. More importantly still, the reproducibility of the process had to be investigated, as large changes in the extraction efficiency for apparently similar process volumes would lead to large systematic errors on the background subtraction for the NC rate.

It is important to note that only extraction efficiency scale tests are required. The calibration of the elution and secondary concentration techniques at 'detector' scales do not require tanks of the 'detector' scale, as the membranes themselves contain all the activity. Thus, investigations of the elution and following stages can be carried out on a full-scale replica of the actual rig to be used in the detector system, with membranes having been spiked with activity in ancillary systems.

The first set of experiments detailed here describe the use of a scale version of the extractive part of the final system, and the appearance of an unforseen effect, radon breakthrough.

5.2 The single one foot membrane tests.

The first tests concentrated mainly on the performance of a new type of membrane, the Amicon H1MP01-43. This replaced the membrane used up until this point, the Amicon H1P30-20. A summary of the properties of both membranes is given in table 5.1.

When the membranes are compared head to head, it can be seen that the MP01 membrane passes more water per unit pressure than the P30. This is only to be expected, given the difference in pore sizes: it should be said that the pore density, i.e. the number of pores per unit area, is approximately the same for the two membranes. In the detector water system, where the flow rate is determined by the requirement that the whole detector volume be passed in four days, this reduces the number of

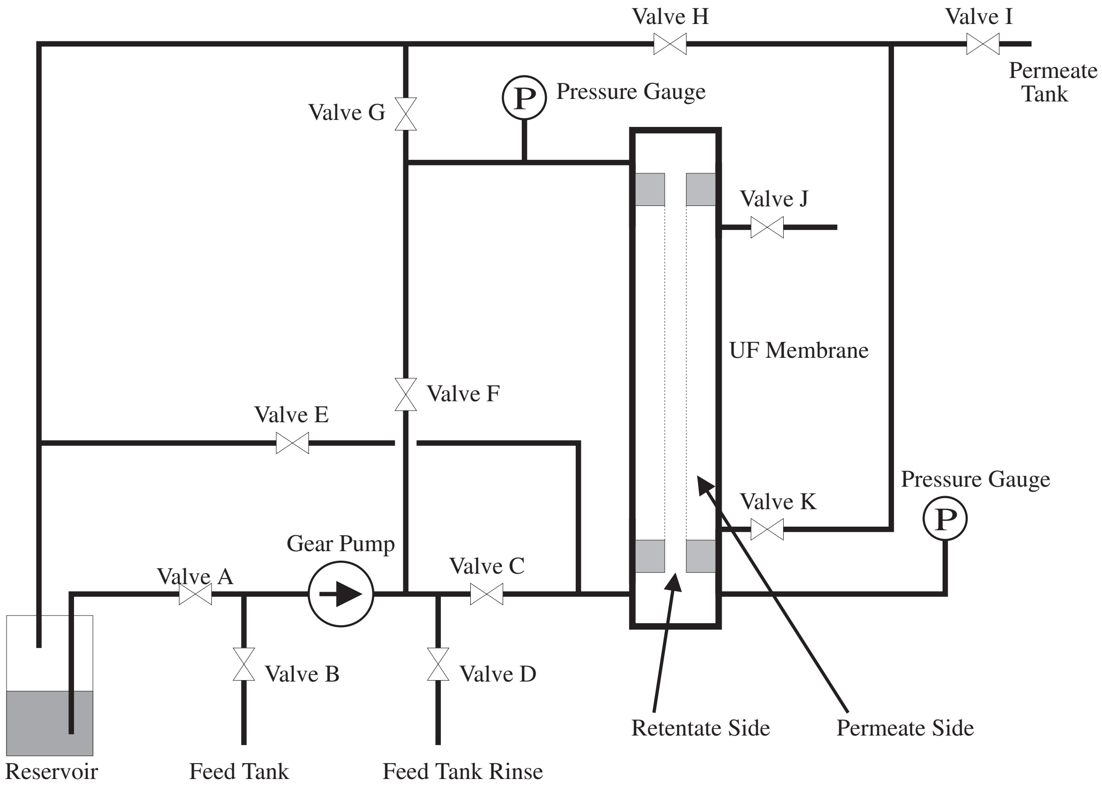

Figure 5.1: An example seeded ultra-filtration rig. The operation of the valves for the various running modes is described in appendix C.

membranes needed from 4 to 2, and hence leads to a reduction by a factor of two such problems as membrane cleanliness and hold-up volumes. However, going to such large pores could have associated problems, most notably situations where the ion-exchanger breaks-up into smaller particles which can then pass through the much larger pores of this new membrane. Such fine particles, or 'fines', have a much higher surface area to volume ratio, and hence have a higher exchange capacity per unit mass. If they are small enough to pass through the membrane, they could carry a significant proportion of the activity with them, and produce errors on the extraction efficiency, as well as scattering the Čerenkov light generated in the detector.

In order to use membranes for SUF, they need to be inserted into systems for the delivery of the water and solutions to the membrane. We will begin with an example SUF rig consisting of a membrane and this ancillary pipework, as is shown in figure 5.1.

Here the UF membrane is shown in a schematic form: the dotted lines represent the walls of a single porous fibre, while the shaded boxes representing the non-porous epoxy into which the fibres are potted. Hence the topmost and bottommost connections are on the retentate side, while the two connections towards the middle are on the permeate side.

All the various running modes are described and shown in appendix C. It suffices to say for the main body of the text that the most important valve is valve G. This controls the amount of water pushed through the membrane as it provides a hydraulic impedance to the water flow in parallel to the impedance of the fibres themselves. Hence if it is fully closed, all the water is forced through the membrane, while if it is fully open, the water flows along the fibres, without no water being forced through to the permeate side.

The scaling factor most appropriate for tests of the SUF is the amount of water passing unit area of membrane per unit time. This has the dimensions of a flow velocity, and for the system proposed for the detector this is 40 litres per minute per metre squared of H_2O at 20 degrees centigrade. Given this scaling, and a 300 cm^2 surface area membrane[1], the required flow rate is 1.2 l/min. The largest convenient volume in Oxford for preliminary tests was 100 litres, which corresponds to 10 tonnes of D_2O if such an experiment were scaled to the detector system.

Previous work had suggested that an amount of HTiO equivalent to 0.5 grammes of titanium per square metre of membrane area (hereafter referred to as g/m^2) led to high extraction efficiencies for the radionuclides[54]. In order to prepare the membrane the HTiO was deposited on the membrane by recirculating the particulate solution. Samples of the solution were taken for spectroscopic titanium analysis in order to quantify the amount of titanium on the membrane. The radioactive spike, around 1 kBq of ^{228}Th in the form of an acidic solution, was taken from a source bottle and counted using a standard GeLi gamma counter. The relatively high activity level of the spike allowed 12 ml samples to be taken from the system directly and counted using the β-α counters. The γ-counting procedure for the spike consisted of the measurement of the rate of the 583 keV γ-line from the decay of the ^{208}Tl. The change in the rate of these gammas with time indicates the relative concentrations of lead, radium and thorium. The feed sample was added to the 100 litre tank and the process volume neutralised to pH 7 using sodium hydroxide. The addition of sodium ions at these

[1] Such a membrane is referred to as a 'one-foot' as this is the length of the membrane cartridge.

levels was not important as the HTiO had been prepared in sodium-rich solutions, and hence would already have sodium on some of the extraction sites.

The process volume was pumped through the membrane at 1.2 l/min, with the trans-membrane pressure adjusted so as to give this flow rate. The membranes were used in 'dead-end' mode, whereby there is no recirculation on the retentate side, and all the water delivered to the membrane is pushed through to the permeate side. During the run the 12 ml samples were taken both from the permeate line and the feed tank. The feed tank was stirred briefly before taking the sample in order to make sure that the radionuclides were well mixed. Samples were taken after 10, 50 and 90 litres of the process volume had passed through the membrane. With a flow rate of 1.2 l/min, the processing of the volume took around 1.5 hours.

These tests were mainly concerned with the extraction efficiency, and not the subsequent elution and concentration stages, for the scaling reasons stated above. The system was not particularly well optimised for elution for this reason, but the membranes were eluted anyway as this constituted a part of the cleaning procedure between the runs. The first elution consisted of 400 ml of 0.5M HCl run for around 5 minutes, while both the second and third eluates were 400 ml of 0.5M HCl run for around 1 hour. The first, quick elution was used to remove the majority of the radio-nuclides into as small a volume as possible, as will be the case in the full scale system. The second and third elutions were carried out in order to remove the HTiO from the membrane. This cleaning with respect to titanium is required as the first acid wash of the membrane, which removes the majority of the radioactivity from the system, always removes a fraction of the HTiO and will obviously distort the particle size distribution. Rather than try to regenerate the HTiO in situ, the membrane is cleaned thoroughly and a new application of HTiO is used for each new experiment. Samples of the eluates were reserved for γ-counting, as described above, while further samples of these eluates were used for spectroscopic titanium analysis so as to determine the fraction of the HTiO that was dissolved by each elution.

The feed and permeate samples were used to make scintillation sources and these began counting immediately after production in order to avoid the loss of the 10.6 hour half-life ^{212}Pb.

Experiment Number	Membrane type	Trans-Membrane pressure required
NS 1	H1P30-20	43 psi
NS 2	H1MP01-43	6 psi
NS 3	H1MP01-43	15 psi
NS 5	H1MP01-43	46 psi

Table 5.2: A summary of the NS experiments. The experiments are described in more detail in the text.

Results

The different experiments are summarised in table 5.2. The membrane type is the manufacturers designation, with the P30 being the 30 kDalton ($\approx$ 5nm) pore membrane, while the MP01 is the 0.1 μm pore membrane. The pressure figure is the pressure that was required in order to pass 1.2 l/min.

The first experiment, NS1[2], did not actually use the new MP01 membrane but the P30 membrane instead. This 30 kDalton membrane was used to check the system could perform as expected and allowed a comparison with previous results. Owing to the availability of thorium sources at the time of the experiment, only around 50 Bq of ^{228}Th were used. The counting of the feed and permeate were consequently more prone to uncertainties. The subsequent three experiments used the same MP01-43 membrane. The increase in the required trans-membrane pressure over the course of the three experiments will be discussed in the conclusions.

The γ-counting and titanium analysis results of the experiments are summarised in table 5.3.

The table is fairly extensive, and some general points must be indicated. The 'feed' lines show the strength of the source that was added to the 100 litre process volume. In experiments NS2 and NS3, only the lead activity was measured, whereas for NS5 and NS1, a small sample of the feed was reserved for further counting of the radium and thorium. The lines labelled 'HTiO' are the mass, in a standard volume, of ion-exchanger deposited on the membranes: the figure shown is the difference between a before-deposition sample and an after-deposition sample. The elution efficiency lines are calculated from the ratio EL1/(EL1+EL2). This should be independent of disequilibria in the feed, and assumes that all the activity is contained in the two eluates, as was

[2]The prefix for these experiments, NS, is a mnemonic for New membrane Seeded ultrafiltration.

Experiment	Sample	Lead c/s	Radium c/s	Thorium c/s	Titanium /ppm
NS1	Feed	0.45	0.45	0.45	–
NS1	HTiO	–	–	–	46
NS1	EL 1	0.34	0.37	0.2	15
NS1	EL 2	0.04	0.05	0.17	29
NS1	EL 3	<0.03	<0.04	<0.04	2
NS1	El.Eff.	89%	88%	54%	–
NS2	Feed	14.6	–	–	–
NS2	HTiO	–	–	–	23
NS2	EL 1	10.8	6.9	3.6	1.2
NS2	EL 2	1.3	2.0	2.1	5.7
NS2	EL 3	<0.03	<0.04	<0.04	0.7
NS2	El.Eff.	89%	78%	63%	–
NS3	Feed	14.5	–	–	–
NS3	HTiO	–	–	–	22
NS3	EL 1	10.3	8.5	4.7	1.5
NS3	EL 2	1.7	2.2	2.7	5
NS3	EL 3	<0.03	<0.04	<0.04	0.9
NS3	El.Eff.	86%	79%	64%	–
NS5	Feed	10.1	6.5	5.5	–
NS5	HTiO	–	–	–	28
NS5	EL 1	7.7	4.9	2.2	1.0
NS5	EL 2	1.2	1.1	1.2	4.3
NS5	EL 3	<0.03	<0.04	<0.04	0.7
NS5	El.Eff.	87%	82%	65%	–

Table 5.3: The results relating to the eluates for the NS experiments. The abbreviation EL refers to the eluate, with EL 1 being the 5 minute elution, and EL 2 and EL 3 being the one hour elutions. The activity columns are in counts per second on the GeLi, which is around 1% efficient. The feed samples are the activity added to the 100 litre process volume, while the HTiO samples indicate the mass of ion-exchanger deposited on the membrane. The El.Eff. lines give the approximate elution efficiency, calculated as the ratio EL1/(EL1+EL2).

confirmed with the observation that the activities in all the EL3's were not significantly above background.

Essentially, this table is included for completeness, and primarily demonstrates that the elution performance of the system was consistent with previous work[53], with 80% of the lead and radium being washed off in the first wash, while 60% of the thorium is also recovered. Both NS1 and NS5 had the feed activity monitored: the source used

Sample name	Litres passed when taken	Thorium DF	Radium DF	Lead DF
NS1-P1	10	105^{+201}_{-61}	79^{+321}_{-42}	66^{+409}_{-30}
NS1-P2	50	77^{+133}_{-49}	>131	17^{+23}_{-13}
NS1-P3	90	79^{+137}_{-49}	>100	12^{+16}_{-9}
NS2-P1	10	12^{+13}_{-11}	>515	36^{+40}_{-33}
NS2-P2	50	14^{+15}_{-13}	>316	26^{+28}_{-24}
NS2-P3	90	19^{+20}_{-18}	40^{+43}_{-37}	12^{+13}_{-11}
NS3-P1	10	171^{+199}_{-147}	>1008	178^{+224}_{-144}
NS3-P2	50	57^{+63}_{-52}	87^{+96}_{-79}	23^{+25}_{-21}
NS3-P3	90	48^{+53}_{-43}	$6.4^{+6.5}_{-6.3}$	13^{+15}_{-12}
NS5-P1	10	192^{+220}_{-168}	>2000	120^{+148}_{-100}
NS5-P2	50	108^{+124}_{-96}	>516	29^{+33}_{-26}
NS5-P3	90	92^{+100}_{-80}	49^{+52}_{-46}	15^{+17}_{-14}

Table 5.4: The DF results for the NS experiments, as obtained from the β-α counting of the feed and permeate samples. The results are shown for the samples extracted after passing 10, 50 and 90 litres in each run. The different isotope levels were obtained from fits of the time spectra of the samples. A lower limit is shown for those samples where the permeate activity was consistent with the background.

for NS2 and NS3 was taken from the same bottle as NS5, and it was demonstrated from the counting of the β-α sources that the disequilibria observed in the NS5 feed activity was also present in NS3 and NS2. The most important conclusion is that the behaviour of the system is consistent over all the experiments.

The β-α counting results are contained in table 5.4, which shows the DF's. The DF (decontamination factor) is defined as the feed activity divided by the permeate activity: hence a DF of ten corresponds to an extraction efficiency of 90%.

Concentrating only on the thorium and radium DF's initially, the DF figures for the experiments are fairly encouraging, with the extraction of radium only dropping below 90% in one determination. A tentative explanation of this fall in radium DF was the absorption of carbon dioxide into the unsealed process volume: it is well known that the extraction efficiency is dependant on pH[51], although no work was done to confirm that this was the case in these experiments. It should be noted that the large variations in DF are somewhat exaggerated, as they correspond to very small changes in the extraction efficiency. The thorium DF results for NS2 are slightly lower (around 5% in extraction efficiency) which was attributed to a poor feed-stock neutralisation: if the figure for the first (i.e. 10 litre) lead DF is also considered, it can be seen that NS2 has a relatively low value as well, suggesting that there was something of a problem in the feed tank.

Summary of NS experiments

There are a number of conclusions to be drawn from the four experiments, and if we consider what the original motivations were, we can make some important statements on the use of HTiO and SUF.

The reproducibility of the experiments is encouraging. As stated already, even though the system was not optimized for elution, with such refinements as the reduction of dead volumes and better control of the acid strength and elution time being introduced in the final system, the variation in the washing-off of the lead, radium and thorium seems good to within a few percent. The final calibration of these values will of course be carried out on a full scale system, with appropriately spiked membranes, but these tests indicate that the errors on the final elution efficiencies should be relatively small.

More importantly, the variation of the extraction efficiency is also quite reasonable, with a random error of a few percent. The fall in radium DF with time is seen in all the experiments, and would be a cause for concern if it indicated that the extraction efficiency were to fall to zero as the extraction continued, but the double membrane tests below show that this is not the case. The thorium DF, which determines the extent to which the D_2O will be purified in the long term, is more than adequate, with very small errors when NS2 is excluded. The lead DF, which directly affects the efficiency of the measurement of the rate of >2.2 MeV gammas, is routinely high at the beginning of each experiment, but then falls in a remarkably reproducible fashion for all the runs. This is discussed in the next section.

The only problem which arose during the tests was the poor recovery of titanium in the eluates of all the MP01 experiments. This was at first thought to be a major sticking point, as it could indicate that the HTiO was being broken down into particles small enough to pass through the pores of the membrane, and hence not available for solution into the eluates at the end of the run. This was thought to be unlikely, mainly as the DF's should have been much lower were this the case. The increase in the MP01 membrane's hydraulic impedance was thought to be related, with the titanium blocking the pores of the membrane, and in the final experiment the hydraulic impedance of the membrane was monitored throughout the entire process. It was noted that the only increase in the impedance occurred during the one hour rinses with HCl, with no change being noted during the passing of the 100 litre process volume. It was thus clear that the one hour acid washes were conditioning the HTiO into a form which blocked the membrane.

It was thought that if the titanium was somehow remaining on the membrane, then by soaking the membrane in acid, the titanium should slowly dissolve off. The MP01 membrane used in the NS experiments was thus immersed in 1M HCl, and left for around 1 month. It was then tested, and the hydraulic impedance of the membrane was seen to be around twice the initial impedance. Such a drop in performance is expected for membranes that have undergone a large amount of use, both from information supplied by the manufacturers, and from previous experience of other membranes at Oxford. Hence, it was surmised that the titanium had been dissolved away to an extent that it was not affecting the membrane's impedance to a large degree. This one month time scale is fairly long, but the membrane was not tested during the time, so it could have recovered after a few days or so.

5.3 Radon breakthrough

As previously noted, the ^{212}Pb DF's in the NS experiments were very similar for all four runs. This systematic behaviour is therefore independent of the pore size of the membrane. This property excludes, to a certain extent, the 'fines' explanation, i.e. the reduction in the size of the HTiO grains and their subsequent passage through the pores in the membrane and into the permeate side. The best explanation was thought to be that between ^{224}Ra and ^{212}Pb, there is the one minute radon, ^{220}Rn. This is known to be very soluble in water, and is a noble gas. Thus it will not exchange on the HTiO, and if the flushing water stream is travelling fast enough, will be carried through the pores of the fibre before decaying to polonium and subsequently to lead. With a flow rate of 1.2 l/min, and a membrane cartridge volume of around 100 ml, it is obvious that the residence time of the water is small enough for such an argument to be valid. This effect is referred to as radon breakthrough.

A theoretical discussion on the effects radon breakthrough is shown in appendix D and if we take the final result for experiments where the total extraction time is small compared to the ^{224}Ra half life, as in the NS experiments above, we obtain:

$$\frac{1}{DF_{Pb}} \approx \lambda_{Pb}\frac{\Delta N_{Pb}}{A_0 \Delta V} = (1-\alpha_{Pb}) + (1-\epsilon)\alpha_{Ra}\lambda_{Pb}t$$

In this equation the α_i's are the probability for the capture of the isotopes from an aqueous solution onto the ion exchanger. The value of ϵ is the probability that the decay of a radium atom captured on the membrane leads to a captured lead atom. Thus near t=0, the ^{212}Pb DF will be dictated by the chemistry of lead take-up on HTiO, whereas for times of the order of the ^{212}Pb mean life, the DF will tend towards a value dictated by the second term in the above equation, depending on the size of ϵ.

Qualitatively, we might expect the value of ϵ to be fairly close to zero, and hence the second term in the above equation becomes very important. If infinite DFs are assumed, then the ^{212}Pb DF is purely determined by this second term (as α_i=1), and the lead activity in the permeate will go from 0 to a value which is directly proportional to the time divided by the ^{212}Pb mean life. For the above experiments this would imply that the lead DF for the three samples would be approximately 100 for the 10l sample, 20 for the 50l sample, and 12 for the 90l sample. As can be seen, this seems to map to the measured values fairly well.

Of course, this is all purely speculative, and depends on the determination of the radon breakthrough parameter, ϵ. It could be that this is actually one, and the lead DF simply falls off with time due to the lead being washed off by water passing over the HTiO. To get a better feel for the value of ϵ, another experiment was needed.

Radon flushing experiments

In order to attempt a measurement of the radon flushing, two filters were used in the rig shown in figure 5.2. These filters, known as Mediakap 10's or MK10's, have a pore size of 0.2 μm, thus matching fairly closely the MP01 filter used in the NS experiments. Both filters were loaded with 0.5 g/m^2 of HTiO, and the flow rate was set at the scale level of 0.6 l/min (the MK10's have a surface area of 150 cm^2). This ensures at least a resemblance to flow patterns in the NS tests above and the 4ft membranes proposed for the detector.

With such a rig, if we load MK-1 with an equilibrium ^{228}Th source, then as the radium decays any radon produced can breakthrough into the 9 litre reservoir. Here it will reside for $9/0.6 = 15$ minutes, during which time the radon decays to polonium and

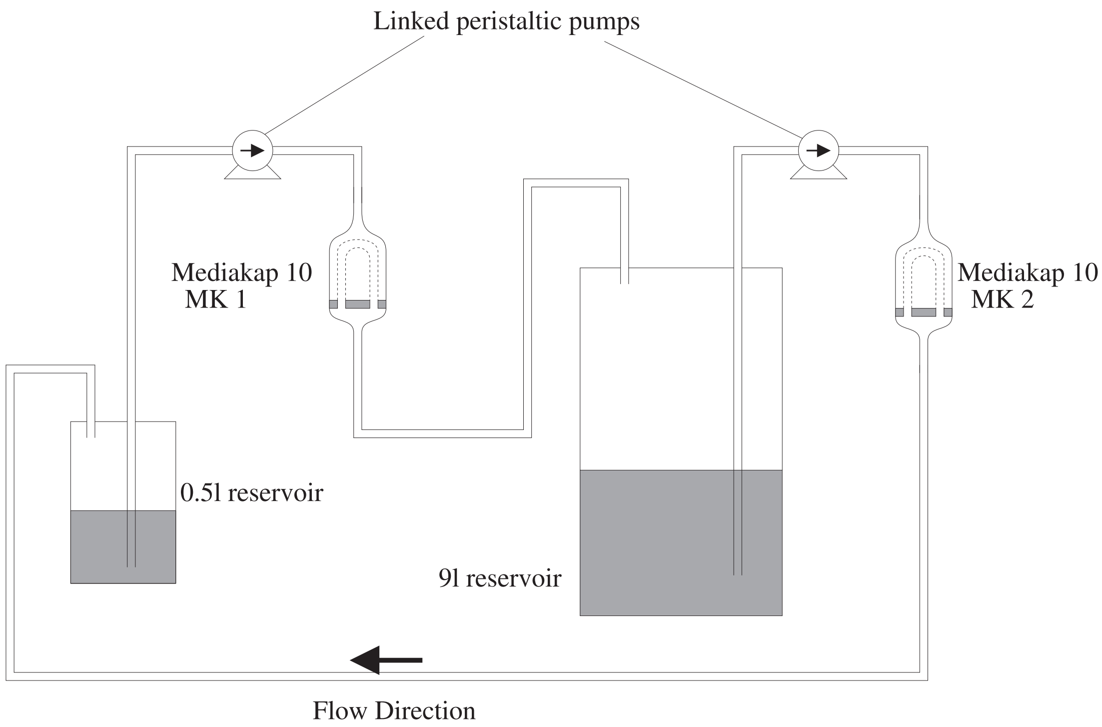

Figure 5.2: The double membrane radon flushing rig.

then to lead, which will then be extracted on MK-2. If, rather than radon breakthrough, we have fines problems, then the fine particles washed off MK-1 and hence we should see a disappearance of all three isotopes, rather than just the lead.

Both membranes were coated with HTiO in a separate rig, while the spike was deposited on MK-1 by recirculating the ^{228}Th, in aqueous solution, in yet another rig so as to avoid contamination. The system shown above was filled with ultra pure water, and a small volume of sodium hydroxide was added to the 9 litre reservoir, bringing the pH of the process volume to 10. This buffered the solution with respect to the absorbtion of carbon dioxide from the air: the absorbtion of CO_2 into the volume had been shown, in previous experiments, to acidify the system to such an extent that the radium was actually eluted off MK-1, and hence radon breakthrough could not be observed (this also went some way to confirming the observation of the falling radium DF's in the NS experiments).

The water was pumped for around 45 hours; this is equivalent to around 5% in ^{212}Pb terms (i.e. if all the radon came off MK-1, there would still be 5% of the lead left on it after 45 hours). The filters were γ counted after being removed from the system

Filter	^{212}Bi activity at t=0, %	^{212}Pb activity at t=0, %	^{224}Ra activity at t=0, %	^{228}Th activity at t=0, %
MK-1	27±4	17±3	98±4	96±2
MK-2	45±6	54±4	4±1	4±1

Table 5.5: The results of the radon flushing experiment, RF5. The activities are shown as fractions of the total activity of each isotope that was taken up on MK-1 at the start of the experiment. The errors are absolute errors on these percentage values, and not percentage errors. The t=0 refers to the start of the counting time, and hence to the end of the pumping.

and the time spectra fitted with four characteristic exponentials - the one hour ^{212}Bi was fitted, as well as the lead, radium and thorium. The results are shown in table 5.5.

The experiment has an obvious signature for radon breakthrough: on MK-1, there is a significant disequilibrium of lead and bismuth, while the radium and thorium do not seem to have been washed off to any significant extent. If the figures are studied, it can be seen that there is also a 30% deficit in the total of lead and bismuth. Around 1% of this deficit will be made up of the radon atoms contained within the reservoir itself, while the rest of the deficit is most probably due to the lead atoms plating-out onto the inner surface of the reservoir, an effect which has been seen in other work.

Were the radon breakthrough hypothesis accepted, the radon loss parameter is around 75%, or to put it more succinctly, 3 out of every 4 radon atoms produced on the retentate side of the membrane find their way through to the permeate side. Another possible explanation is that lead and bismuth are being washed off, whereas radium and thorium are not. This is a little hard to believe, with no obvious mechanism to differentiate between the four isotopes and their interaction with HTiO.

It would seem that two important conclusions can be taken from this experiment: firstly, the radon breakthrough parameter is quite high, at around 75% for these membranes. Of course, in the hollow fibre cartridges, there will be different geometry and flow pattern effects, and the extrapolation of this result to a 4 ft membrane seems a little cavalier.

Secondly, it would appear that there are no obvious fines problems with the HTiO in this experiment. The pore size of these MK10's is 0.2 microns, which is at least twice the size of any of the proposed hollow fibre cartridges, and hence any fines problem would be more pronounced in these runs. A signature for fines would be a disappearance of all the radioactive species on MK-1 as the titanium breaks up and takes the activity with it. As is demonstrated by the 'holding' of the radium and thorium in the above experiment, there appears to be no such distribution of activity.

5.4 The double one foot membrane tests.

We had raised some interesting points with respect to the removal of radionuclides even at the 100 litre level. However, this process volume represents only around 10 tonnes when scaled to the detector system, and it was thought that a natural progression would be to increase the process volume and run the extraction for a more representative time. Also, it was thought necessary that there be a demonstration of radon breakthrough in a scale system. The next step was therefore a 'convolution' of the two existing experiments.

In order to remove the upper limit on the process volume, and hence the running time, a new rig had to be built, and is shown in figure 5.3. The water purification part of the system, consisting of the R/O and IX systems on the left, was designed to give a continuous ultra-pure water stream at the 1.2 l/min required. This clean water stream was spiked continuously at the point shown with a ^{228}Th source; to dismiss the problem of calibrating the dosing procedure, 12 ml samples of the feed and both permeates were taken directly from the system at the points shown. The two H1 membrane systems were essentially reproductions of the system used in the 100 litre tests. The 10 litre tank between the two membranes meant that, when switched-in, the water stream would be held-up for around 10 minutes. This means that 99.9% of any radon appearing on the permeate side of the first membrane should have decayed to lead before reaching the second membrane.

The removal of the upper limit on process volume allowed the investigation of such properties as the reduction in the ^{224}Ra DF observed in the NS experiments, and also allowed the first positive determination of fines from a 0.1μ pore size membrane: the second membrane in the series was a 5 nm pore filter, and hence would catch any HTiO particles smaller than 100nm, but larger than 5nm. An excess of titanium on the second membrane together with a corresponding loss from the first membrane would hence signify that particulate HTiO was being washed off and through the pores. Such a determination could be made either by a direct measurement of the total titanium on the membrane, or an increase of the extraction efficiency of the second membrane, with a corresponding drop in efficiency of the first. Although no fines problems had

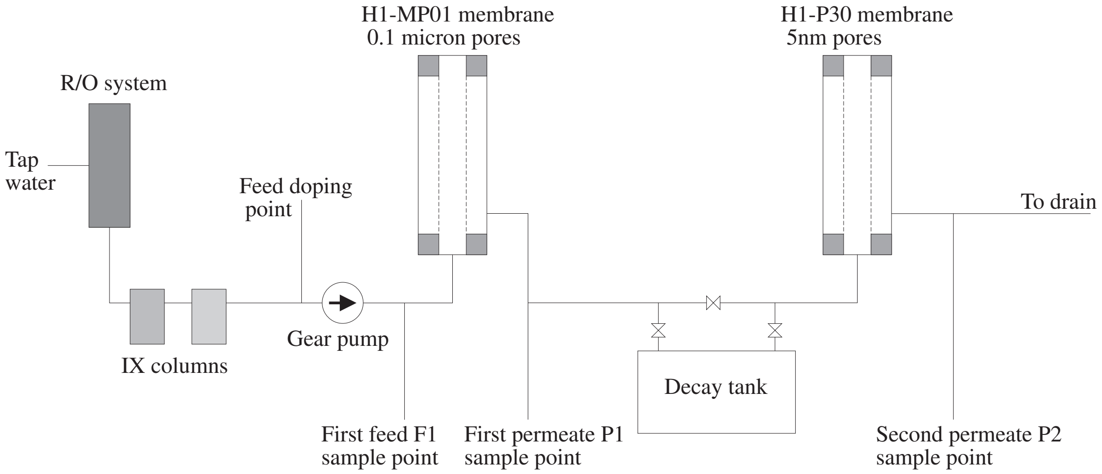

Figure 5.3: The two H1 system. The 'sample line's shown give the points in the system where the samples known as F1, P1 and P2 were taken from.

yet occurred, it was thought that the phenomenon could appear with longer run times.

The system was set up with the MP01 and P30 in series, both of which had 0.5 g/m^2 of HTiO deposited on them, and the flow rate was set to 1.2 l/min. The spike was adjusted to give a constant doping rate of around 1 Bq per litre in the feed stream, and the system was run for around 4 days, or around 5 tonnes total throughput. If this is scaled to the detector system, it is roughly equivalent to a 1 kilo tonne volume. Samples of the feed and both permeates were taken at selected times into the extraction and counted immediately using the β-α system. The time spectra were fitted again, and the DF's calculated.

Two runs were carried out with the system: the first experiment, DM2, had the 10 litre tank switched out, while the second experiment, DM3, had the tank switched-in. The cumulative DF's (i.e. F1/P2 in the diagram above) for the two experiments are shown in table 5.6.

Again, the thorium and radium DF's are approximately at the 95% level, and show no obvious time dependance over the time of extraction. This implies that the fall in radium DF seen in the NS experiments was indeed due to carbon dioxide absorbtion, an effect that could not arise in the DM tests, as the water was not allowed to stand open to the air.

The lead DF's, although obviously out of the ordinary, are easily explained using the radon breakthrough effect. The DM 2-1 sample, taken after 6 minutes, shows a good lead DF. The DM 2-2 sample, after 53 hours, shows significant radon breakthrough,

Sample	Taken at	^{228}Th DF	^{224}Ra DF	^{212}Pb DF
DM 2-1	6 minutes	36^{+45}_{-28}	196^{+419}_{-115}	67^{+182}_{-33}
DM 2-2	53 hours	106^{+362}_{-56}	29^{+52}_{-18}	$0.35^{+0.47}_{-0.26}$
DM 3-1	1 hour	180^{+331}_{-106}	>513	20^{+28}_{-15}
DM 3-2	50 hours	181^{+307}_{-115}	64^{+297}_{-30}	$1.3^{+1.7}_{-0.9}$
DM 3-3	115 hours	736^{+2000}_{-267}	91^{+277}_{-47}	$2.3^{+3.3}_{-1.6}$

Table 5.6: The cumulative DF's for the double membrane experiments. The first experiment, DM2, had the decay tank switched-out, while the second, DM3, had the decay tank switched in. Notice the difference in the later lead DF's for the two cases.

with around twice the amount of lead in the permeate as in the feed stream. This is as expected for radon breakthrough, with the radium on the first membrane acting as a second source of lead activity, and hence producing more lead on the permeate side of this first membrane than is being supplied on the retentate side: the residence time with the decay tank out was much shorter than a radon half life, so that radon produced by the first membrane would not have decayed by the time it reached the second membrane: thus, the excess lead appears in the cumulative DF also. Using the equations from appendix D, a tentative radon breakthrough parameter of 0.6 (ie 60% of the radon is washed off) gives the observed signal. This is significantly different from the MK10 tests above, but could probably be explained by the different flow geometries.

The addition of the holding tank between the membranes in the second experiment gave the radon that was washed-off an opportunity to decay to lead before it reached the second membrane, and hence there was a higher lead cumulative DF for the double membrane system. Particular attention was paid to this cumulative DF, and to its two components. For the DM 3-3 samples, by counting the P1 sample between the two membranes, it was shown that the DF for the first membrane was 0.4 (in agreement with DM 2-2) while the DF for the second membrane was 6. This did not agree very well with the raw chemical DF for lead, thought to be around 100. The main reason for this discrepancy could be due to improper mixing in the delay tank, with the average residence time being much shorter than the simple volume over flow rate calculation: such a behaviour could be altered by adding baffles to the tank.

Qualitatively speaking, it is obvious that, were the system exposed to a constant source of thorium, or even radium, the removal of lead from the water stream over the four day time scale is impossible without a decay tank. This is desperately important for the SNO experiment, as it is the daughters below lead which produce the NC-mimicking gamma rays.

The other important conclusion from the experiments was that there was still no obvious fines problem. This was established by the amount of titanium eluted off both membranes. The total titanium mass recovered for both membranes was around 90%, with the other 10% expected to still be on the membranes, and although a 10% loss of titanium via fines could not be excluded, it would be a surprise were the HTiO to be smaller than the 5 nm pores in the second membrane.

5.5 Conclusions from the scale tests.

The scale tests offer us pointers as to what we need to do and what we can expect from the SUF plant in SNO:

- Both the thorium and radium DF's were routinely high, with the extraction of radium perhaps being slightly lower. From the NS experiments, the consistency of DF's between different experiments indicates that the extraction is reproducible to the +/-5% level, while the DM experiments confirm the DF's stability over the course of a 4 day extraction, which is the longest time scale proposed for the full system.

- There seems to be no obvious problem with HTiO fines, with no significant excess of titanium being observed on the second membrane even after running for four days. A 'loss' of titanium is observed, as well as an increase in the hydraulic impedance of the membrane, when the membranes are eluted with 0.5 M HCl, but this problem can be worked around by long term exposures to concentrated acid.

- The DM experiments show that the introduction of a delay tank leads to a significant improvement in the extraction of ^{212}Pb, in those cases where there is also radium present in the feed stream. The decay tank thus stops the appearance in the detector volume of those lead atoms produced by the decay of the radium on the membrane.

It is important to note the following caveat at this point: the work done on the double membrane and 100 litre systems was instructive and helped in the design of the full scale plant[55]. Also, it demonstrated that the reproducibility was good to the 5% level: in the final water system, it is expected that the root N counting statistics of the samples generated will give 20% errors, and hence a 5% systematic error is not particularly significant. However, there are two variables that were not investigated: the change in extraction efficiency as the activity level falls to 1 decay per day per tonne; and the presence of $MgCl_2$ in the water stream. Both of these could conceivably have a significant effect, but there are also problems associated with the experiments required to investigate them:

The reduction of the activity level to 1 decay per day per tonne is of course impractical for scale tests, as it would require huge process volumes. As well as this, the accurate measurement of the extraction efficiency would depend on the elution and secondary concentration efficiencies as such low levels of activity would not allow the direct counting of the feed and permeate streams.

Although it has been shown[51] that the exchange capacity of the HTiO is not adversely affected by the addition of 0.2% of $MgCl_2$ the competition between the magnesium and the radioisotopes will only be exaggerated as the activity level is reduced. This may be a worry, but it is also true that a representative Mg/Th ratio can only be obtained with a one kilo-tonne process volume.

The best solution to both these concerns would be to spike the detector itself with a short lived radionuclide and calibrate the system with a process volume consisting of the detector itself. Such schemes are already encouraging, with the repeatable production of samples of ^{224}Ra with less than 1% by activity of ^{228}Th already being carried out by the author. Other approaches, for example the carrying of ^{220}Rn by gas streams bubbling through solutions of thorium, are being actively investigated, and may offer the SNO experiment the only possible calibration of the water system. Of course, such schemes will only tell us about the extraction of the short lived isotopes themselves, and not about their long-lived parents, but given that such radionuclides as ^{224}Ra and ^{220}Rn are above the >2.2 MeV gamma-producing steps, we can be fairly confident about the prospects of an accurate determination of the NC background in SNO.

Chapter 6

Eight inch PMT characterisation.

There follow three slightly disparate sections, the first detailing experiments with the Hamamatsu SNO eight inch tube, the latter two involving tubes made by Electron Tubes Ltd, or ETL. This company was formerly known as Thorn EMI Electron Tubes, but recently underwent a management buy-out. The author was sponsored by the company under the CASE award scheme, and much of the work in the next two chapters was carried out in compliance with the terms of the award.

The ETL tubes used were either type 9350, type 9351 or type 9353. These are all eight inch tubes resembling those used in the SNO experiment, but are optimised for slightly different performance characteristics. The 9350 tube is currently in use by the MACRO detector at Grand Sasso, the 9351 tube is to be used in the BOREXINO solar neutrino detector, while the 9353 is in use by the AMANDA cosmic ray observatory.

We will begin with an investigation of some of the properties of the SNO tube, and more particularly the changes in performance when the voltages across the dynodes are changed.

6.1 SNO PMT base optimization.

The photomultipliers used in SNO are Hamamatsu R1408 eight inch tubes, which have a bialkali photocathode and a Venetian blind dynode structure. The R1408 is principally a development of the 20 inch tube used in the original Kamiokande detector. The tubes are used in the measurement of the three important quantities of the relativistic electron: the position, the direction and the energy. Each of these measurements is affected by the three important tube characteristics: transit time spread (TTS), noise rate and relative efficiency.

The solar neutrino induced electrons expected in the detector should have a range of a few centimeters, and will undergo a large number of scatterings over this path. The light output will thus be a smeared Čerenkov cone projected onto the PMT shell. Even the most complicated algorithms for the reconstruction of the position of the event have to use the time of flight of the photons from the event to the PMT. The dynamics of the shower within the dynode stack mean that the time between the photon striking the cathode, and the pulse's appearance at the anode, also referred to as the transit time, has a gaussian spread, known unsurprisingly as the transit time spread, or TTS. Thus, the fitting of the event is limited not only by the efficiency of the PMT/concentrator combination, which increases the number of struck PMT's for a given energy, but by this inherent spread in read out time.

Noise rates also contribute to bad reconstruction in that if an uncorrelated noise hit occurs within the set time window around a hit pattern, it can pull the fit to one side or another. This effect can be reduced by constraining the fit to successively smaller time windows, but as only around 10 PMT hits occur per MeV electron energy, this discrimination becomes more difficult at the lower energies of around 5 MeV. For example, a 100 nanosecond time window for all 10,000 tubes, as has been proposed for the SNO, leads, on average, to one noise hit if all the tubes had a noise rate of 1 kHz. Thus an increase to 10 kHz for the average noise rate scales directly to 10 noise hits within the window. It has been found from early Monte Carlo calculations that an increase in TTS of 0.1 ns for all the tubes is equivalent, in reconstruction accuracy terms, to an increase in noise from 2 to 4 kHz for all the tubes.

The direction of the electron is found from the addition of the vectors joining the reconstructed event and the phototube. As the event position is subject to all the properties mentioned above, the error in the direction is also subject to the efficiency, the transit time spread and the noise of the tubes. Again, cuts in both time and position within the detector can be applied to throw away noise hits, but such cuts are again restrained by the small number of real PMT hits. Having said all this, the main error on the direction will be the inherent uncertainty due to the multiple scattering of the electron.

The tubes are also used in a calorimetry sense to measure the energy of the electron. By counting the number of tubes hit within the time window, the energy of the electron can be found, as the number of Čerenkov photons produced is proportional to the track length, and hence related to the energy of the electron. The most important concern in this case is the overall efficiency of the tube/concentrator combination and its dependance on the angle at which the photon strikes. Inclusion of noise hits leads to overestimation of energies: smaller noise rates are unimportant in that they will be included in the calibration of the detector, but large noise rates have large random fluctuations, and thus are to be avoided.

There was thus a drive towards trying to improve the performance of the PMT's, and also the try to look at the balance between all the performance characteristics. In order to change the properties of the tubes, the only freely available variables are the voltages across the dynodes, especially those voltages across the first few dynodes. This voltage variation was carried out on six of the SNO tubes, selected at random from the PMT stock in Canada, using a coincidence arrangement in a piece of kit known as the mortar.

The Mortar.

Each tube was inserted into the Mortar, shown in figure 6.1.

The Mortar is composed of a fast Hamamatsu 50mm photomultiplier mounted at one end of the Mortar tube (approx. diameter 0.5m) at the other end of which the SNO tube was mounted at a distance of 1.5m. A Čerenkov source was used which consisted of 300 kBq of ^{90}Sr in 20 ml of water. The water is held in an acrylic sphere, this acrylic being of the same type as that used in the SNO acrylic vessel. The coincidence arrangement allowed the transit time spread to be measured by using the smaller tube to provide a reference time. The charge of such coincident events was also recorded, and their rate was used as a rough indication of the efficiency. Such a correlation comes about through the stability of the source over the testing period and the small coincidence window meaning that random coincidences constituted only a small fraction of the rate. Finally the dark current was recorded by measuring the rate of pulses when the source blocker was closed.

Each tube's total voltage was set so as to give a gain of 10^7, and with the variation of the voltages across the first few dynodes, the correct total voltage had to be found for each run. The software discriminator on the charge spectrum was set at 1/4 of a single photoelectron, and this cut was applied to the events in the time spectrum and also to the coincidence rate so as to simulate the circumstances in the detector, where

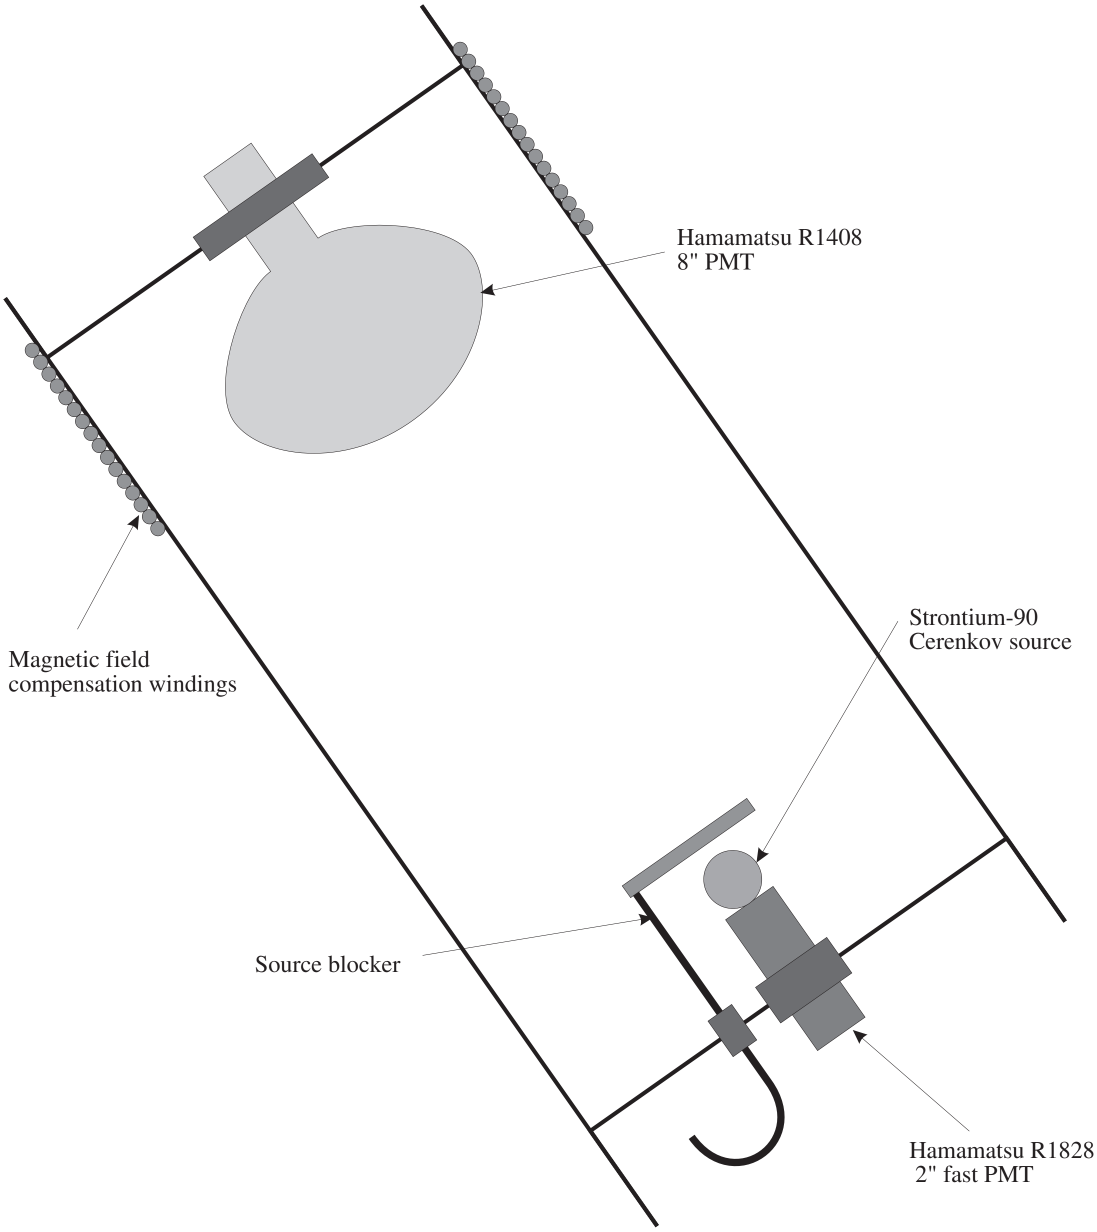

Figure 6.1: A diagram of the Mortar. It is essentially a coincidence arrangement, with the light source fixed to the small tube and the large tube placed 1.5 metres away. Great care was taken to make the arrangement light tight, and to stop reflections from the side of the Mortar itself.

the electronics will also have this discriminator setting.

Noise

The noise rate in kHz is plotted against the voltage FD1 for the three values of

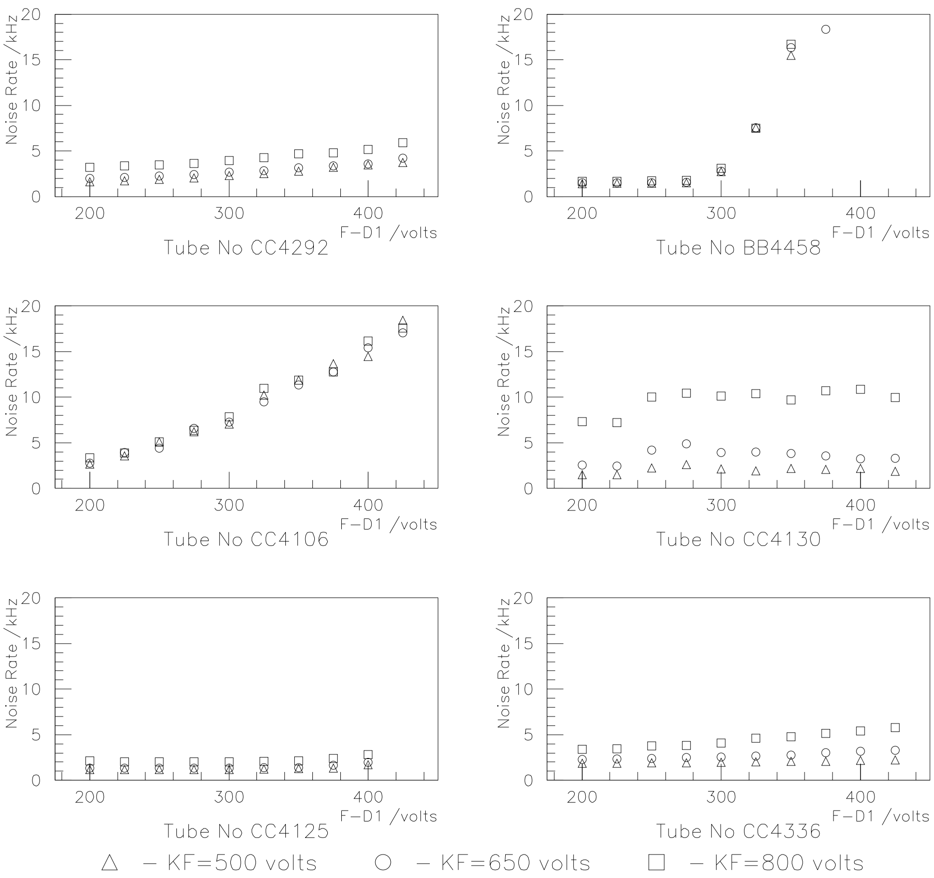

Figure 6.2: The noise behaviour for all six tubes. The tubes are grouped into two rough categories, with the sharp increases of BB4458 and CC4106 being separated from the other four tubes. There is also an increase with KF for some of the tubes, although not in such dramatic fashion as BB4458.

KF in figure 6.2: there will be many references to the mnemonics. The K refers to the cathode, F to the focus electrode, D1 to the first dynode and A to the anode. Hence KF is the relative voltage between the cathode and the focus electrode, D1A between the first dynode and anode, and so on. All six tubes are shown in the figure, with the different symbols representing different values of KF. The standard SNO base has the voltages KF and FD1 in the ratio $\frac{KF}{FD1}=\frac{5}{3}$.

It is readily apparent that there is a huge variation between one tube and the next. The catastrophic breakdown of BB4458 is the first point to note. This had been seen before, as this tube was tested in the KD1 variation regime, but here we see that the breakdown is due to the value of FD1, not KF. The sharpness of this increase in noise rate would tend to suggest some sort of critical phenomenon, such as electrostatic breakdown or corona discharge. However an important point to note in our search for optimization is that for FD1 below 300V the noise is acceptable (we will see later what this implies for other properties). The tube BB4458 is only approached in behaviour by CC4106. Although it shows no obvious step at any FD1 voltage there is clearly a dependance on FD1 which is only seen elsewhere in BB4458.

The other four tubes are loosely grouped together in that they do not show the dependance on FD1 that these first two tubes do. However there are also some similarities in these four tubes that can be used to classify them as having different noise phenomena to BB4458 and CC4106.

The dependance on KF in these four tubes is evident, but another point to note is that three of them have noise rates which stay well under 10 kHz even at high FD1. However the rise due to KF is evident and is interesting in that the noise rates at KF=500 and 650 volts are similar, but the noise appears to step up on approaching 800 volts in much the same fashion as BB4458 does with increasing FD1, though perhaps not quite as dramatically. The similarity between tubes is enhanced by all the tubes having a noise of around 3 kHz when KF=650 volts and FD1=390 volts, as is the general case for the normal SNO base.

Transit Time Spread

Again we must first note some points about the plots themselves, which are shown in figure 6.3. The three shown are the complete runs for BB4458, CC4125 and CC4130. The axes are the same for all tubes, and the higher FD1 measurements for BB4458 were ignored as the noise became too large to contemplate running the tube in SNO.

The main point to make here is the very good straight lines obtained throughout all six tubes. Although only three are shown here, the three not present all showed remarkably similar behaviour over the ranges of FD1 and KF. The higher KF points consistently led to lower transit times, as did the the higher FD1 voltages. Also, we know that, to first order, the TTS of the small tube adds in quadrature to the measured TTS, and hence we would expect the slight tailing off of the TTS at very small values. The difference in end points for the tubes could be assigned to slight manufacturing differences in both the dynode stack and the glass bulb. The slopes of the lines against

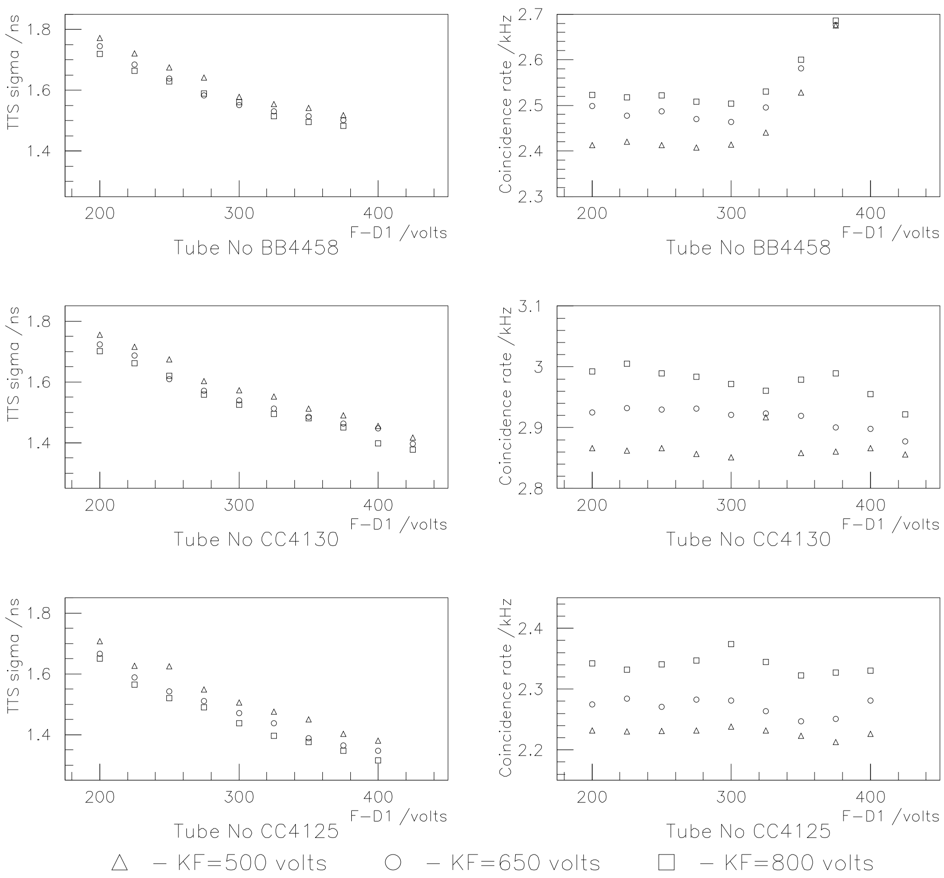

Figure 6.3: The transit time spread and coincidence rate for three of the PMT's. The coincidence rate is a fair indication of the efficiency of the tube as the source itself was stable over the testing time for all six tubes.

FD1 are shown in table 6.1, along with the slope against KF and the TTS at KF=800V, KD1=425V (the endpoint) which represents the smallest possible TTS value.

A good deal of agreement can thus be seen from tube to tube, except at the end-points of the TTS against FD1 graphs. This small variation is simply due to the variation in manufacture from tube to tube.

Now we can see that to bring about an decrease of, say, 0.1 ns in the transit time spread we would have to increase FD1 by around 75 volts, whereas to bring about the same change by increasing KF, we should have to increase it by 500 volts. This has interesting repercussions. If, for example, we wanted to reduce the operating point of FD1 by 50 volts in order to bring it below the noise catastrophe in BB4458, then we

Tube No	TTS Slope FD1 ps/V	TTS Slope KF ps/V	TTS endpoint ns
BB4458	1.1	0.18	1.48
CC4125	1.4	0.20	1.32
CC4130	1.4	0.18	1.38
CC4336	1.5	0.17	1.36
CC4292	1.3	0.20	1.34
CC4106	1.4	0.22	1.30

Table 6.1: The measured values of the TTS parameters. The change in TTS per volt for the focus-first dynode voltage is of the order of ten times the the similar change for the cathode-focus voltage. The large variation of the endpoints means that each tube has to be tested before insertion into the detector, as indeed is being done.

would have to increase KF by around 300 volts in order to compensate the rise in TTS due to this change of FD1. This has serious side-effects, not least of which is the noise behaviour of the second group of tubes. There is every reason to expect that the noise rate of BB4458 would increase in much the same way as these tubes, and so we would find ourselves with a noisy tube once more (although the rise due to this increase in KF would probably be much less dramatic than the sharp increase against FD1).

Relative Efficiency

The coincidence rate plots for three of the six tubes are shown on figure 6.3 also. The y-axis scales are different tube to tube, and only three tubes are shown as the other three showed a good degree of similarity in nature to the plots here. As stated before, due to the Čerenkov source's stability with time, these coincidence rates can be said to be directly proportional to the relative efficiency from tube to tube, even though no overall calibration is given.

The efficiency plots all showed the same flat behaviour against FD1. The really interesting feature in these plots, is not the variation against FD1 but that against KF. The magnitude of the change in count rate is somewhat surprising, but could be explained by the increase in the probability that the isotropically-generated photo-electron reaches the dynode stack as the voltage rises. This is also referred to as the collection efficiency. A point to notice is that of the high FD1 behaviour of BB4458. The tendency for the count rates to rise at large FD1 was also seen to a lesser extent with tube CC4106, the other high noise-FD1 tube, and is due to the high noise of the tube, with noise hits in the R1408 occurring in the coincidence window. For normal noise rates of around 4kHz this effect is negligible, corresponding to less than 1% of the counts, whereas for noise rates of 40 kHz, as BB4458 reached at the high FD1 end, this random coincidence rate can be seen to affect the plot considerably.

Conclusions

Although this work did not bring about any change in the SNO base, it was useful as an introduction to the properties of the eight inch tube. It also highlighted what will be discussed in the next chapter - the breakdown noise phenomenon. Otherwise, the characterisation of the PMT properties seemed fairly straightforward: increasing the cathode to focus voltage did increase the collection efficiency, although it would also increase this efficiency for the collection of thermally generated noise-electrons from the cathode as well.

The decrease of TTS with KD1 was a well known effect in these tubes as the tracks of the electrons through the glass bulb straighten with increasing electric field, but the dependance on the two components of KD1, namely KF and FD1, was somewhat unexpected. This does offer some scope for improving the TTS, but the behaviour of the noisier tubes worked against any such change.

6.2 Absolute efficiencies.

The measurement of efficiency for a PMT is not an easy process either to define or carry out systematically. Obviously, the best result would be to quote a single value for any given tube, but the number of variables precludes this. Such parameters as discriminator setting, KD1 voltage and photocathode area lead to complications which mean that efficiencies simply cannot be attributed on a tube-by-tube basis. With this in mind, Boardman defined the single photon counting efficiency as the probability that a photoelectron (produced by a Cerenkov photon) is detected in the electronics after discrimination.

In order to measure such a parameter, we need a standard light source. If we consider a Cerenkov source, such as a radioactive β emitter in aqueous solution then we can find the expected count rate using the assumption that averaged over a large number of decays, the light produced will be isotropic. This is not true for any single decay, but due to there being around one million decays per second for this particular source, we can assume isotropy with a good degree of confidence. The formula for the rate is then:

$$R = 2\pi\eta c\, A\Omega \int_{300\ nm}^{700\ nm} Q(\lambda)\, T(\lambda) \left(\frac{dN}{d\omega}\right) \frac{d\lambda}{\lambda^2}\,, \tag{6.1}$$

Here R is the measured rate, A the activity of the radioactive source in decays per second and $dN/d\omega$ the number of Cerenkov photons produced, per decay, in a unit frequency interval. This should be independent of frequency and must be found by considering the mechanics of Cerenkov production within the source. The number given here refers to the isotropic case, so we have to multiply by the solid angle subtended by the exposed area of photocathode Ω in order to compare with measurement. It is at this point that a second assumption is made: that in the experimental set-up used, the Čerenkov source can be considered to be a point source.

The functions Q(λ) and T(λ) refer to the quantum efficiency and the transmission of the SNO acrylic respectively, both of which are dependant on wavelength. The acrylic was introduced between the source and the tube so as to define a low wavelength limit on the Cerenkov light. The values for these functions are defined to range between zero and one and were found independently of this experiment. The limits on the integral are such that all the useful information was contained, as the acrylic used cuts off at around 370nm, while the quantum efficiency of the bialkali photocathode falls rapidly above 500nm. The parameter η is the single photon counting efficiency, which is the measure of the probability that a photoelectron produced at the photocathode is detected after electronic discrimination of the PMT output pulse. By comparing predicted and measured rates, the value of η can be found, and by then convoluting this number and the quantum efficiency, a value for the absolute efficiency can be found.

Method

The Čerenkov source used consisted of a solution of SrCl in water enclosed in a glass sphere of radius two inches. The glass is transparent to UV radiation down to around 200 nm, while the acrylic cut-off is at around 370 nm, so the first assumption was that the glass/water system in which the Cerenkov light is produced is transparent in the wavelength region we are interested in. This system had been modelled extensively using the EGS4 code for electron-gamma showers, and the results from this have been accepted here without prejudice.

The activity, A, was found from the quoted activity as given by the suppliers[1] on a particular date and adjusted to account for the decay in the specific activity up to the time of this experiment. Again, the original number from the suppliers was taken as accurate, and any systematic errors are considered later.

The quantum efficiency graphs were produced by the standard practical setup at E.T.L. consisting of calibrated photodiodes, a tungsten filament lamp and a deuterium lamp. Due to the acrylic cut-off at 370 nm the data from the tungsten lamp was used more extensively, thus obviating any possible problems about cross calibration between the deuterium and tungsten lamps. The Q.E. for all six tubes was measured between 200 and 750 nm. A typical response is given by the eight inch PMT KFL 9353 S/No 254, shown in figure 6.4.

The transmission of the perspex was also measured, this time using a Lambda-9 spectrophotometer, for light wavelengths between 300nm and 750nm and the spectrum is shown in figure 6.5. As can be seen there was a very sharp cut-off in the region of 370 nm, while above this value the transmission was almost constant at 92%. This value was below one because of Fresnel scattering and reflections at the top and bottom surfaces of the acrylic. This will not be such a problem in the SNO experiment as here we are considering an acrylic/air interface, and not the acrylic/water interface that will occur in the detector.

The convolution of these two gave the shape shown in figure 6.6, which also includes a multiplication by the integral conversion factor of $1/\lambda^2$. This combination of all three was integrated over the appropriate region so as to give a value applicable to rate equation. This was done numerically for each tube using its individual Q.E. measurement and the values are shown in table 6.2:

As can be seen there was a large variation tube to tube. This can be put down to variations in the thickness and chemical composition of the photocathode. The Corning blue value is also shown: this is the current from the tube when it is exposed to a tungsten lamp at 2720K, with a standard Corning filter applied to the light. This filter is a band pass device, with a peak at around 420nm, and a FWHM value of around 20 nm, and thus gives a fair estimation of the response of the tube to blue light. This is obviously a far easier value to measure than the complete quantum efficiency response

[1]Amersham International PLC

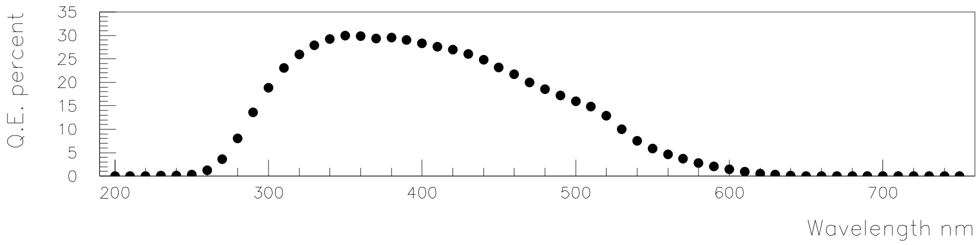

Figure 6.4: A measurement of the quantum efficiency for one of the eight inch tubes used for the subsequent work. The distribution itself follows a familiar pattern for bialkali photocathodes, with the peak at around 400 nm.

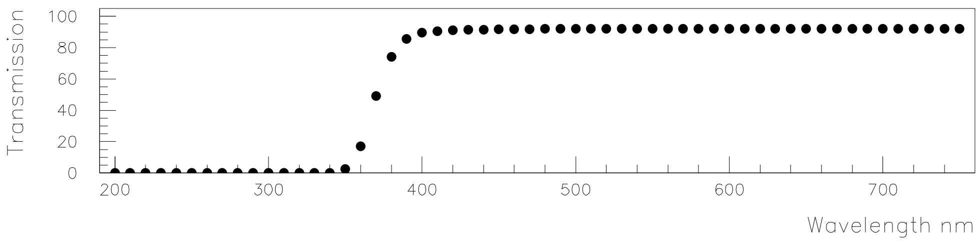

Figure 6.5: The transmission of the perspex used for the experiment. This plastic was a sample of the material that is to be used to make the acrylic vessel in the SNO detector, and shows a fairly low cut off in transmission at around 370 nm.

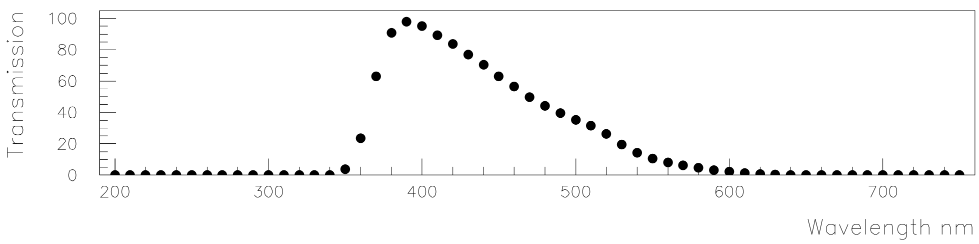

Figure 6.6: The convolution of the Q.E., acrylic transmission and a $\frac{1}{\lambda^2}$ factor. This is the function that was numerically integrated in order to calibrate the response of the PMT, and represents the spectral response of the PMT to a light source of the sort which will be seen in the SNO detector.

Tube Type	Serial Number	Value of Integral $/m^{-1}$	Corning Blue μA/lm
9353	254	178763.9	12.2
9353	257	172434.9	11.5
9353	258	185017.5	12.5
9351	1277	195125.8	13.0
9351	1283	179406.6	12.5
9351	1286	185355.2	12.0

Table 6.2: A table of values for the integral involved in the rate equation given above. The Corning blue value, found using a standard filter is also shown for comparison.

of the PMT, and correlates fairly well with the value of the integral.

Rate measurements.

The experimental rate was measured using a standard ADC, with the dead time being factored out by using an in built live time measurement. The ADC was calibrated using a pulser and a digital oscilloscope, and each tube had its total voltage set so as to give a gain of 10^7. Each tube was exposed to the source for around 20 seconds, and then the source was covered, and a dark current run was taken. The pulses from the PMT were recorded and appropriate regions of interest were set so as to simulate a change in the lower level discriminator. The rates were calculated by summing the number of events over the ROI and dividing by the live time. The two rates, source-on and source-off, were subtracted to give a true rate, and the whole process was repeated two more times. The three values were then combined to give a mean and standard deviation of the rate for each tube.

Although only six tubes were tested, a fair indication of the variation from tube to tube was possible, as the tubes studied were selected randomly from the production batch. The data in table 6.3 is for KFL9353 S/No 254, and represents a typical data set.

When these mean rate results were plotted against the lower level discriminator, the graph in figure 6.7 was obtained. It should be noted that the zero charge channel is 52, and the s.p.e. peak is at 152 channels. The error bars on the graph represent five times the standard error on the mean and are typical of the six tubes tested. The distribution of the rate is much as expected with a steady, almost linear, part to begin with, but then dropping more quickly as the lower level discriminator starts removing

Lower Level Channel No	Lower Level S.P.E.	Mean Rate /Hz	Standard Error on Mean /Hz
53	0.01	12877	156
57	0.05	12563	142
62	0.1	12438	150
67	0.15	12195	144
77	0.25	11703	136
82	0.3	11440	135
102	0.5	10279	123
127	0.75	8440	105
152	1.0	6192	104

Table 6.3: The results for a typical discriminator setting run. The errors were calculated from the three runs taken for each tube. Note the zero offset of the ADC at channel 52, with the single photo-electron peak at channel 100 above this.

some of the events in the single photo-electron peak of the PMT charge spectrum.

Calculation of counting efficiency.

In order to calculate the parameter η, certain preliminaries are needed. Firstly, the activity of the source is given by:

$$A = A_0 e^{-t/28.6}$$

where t is in years, and the activity A_0 was the activity on January 8, 1992 and was given as 1129.67 kBq. For these experiments, the value of t was 1.175 and so A = 1084.2 kBq. The value of Ω, the fractional solid angle, was calculated from the exposed photocathode area, and the yield number was given as 1.317×10^{-14} from the Monte Carlo of the source. So:

$$\eta = 9.838 \frac{R}{\int_{300nm}^{700nm} Q(\lambda) T(\lambda) \frac{d\lambda}{\lambda^2}}$$

It is also necessary to correct for the afterpulsing differences from tube to tube. A special piece of kit known as an autocorrelator was used to give a dead time-less measurement of the probability of an afterpulse occurring. Such afterpulsing could lead to an increase of the measured rate, and if we say:

$$\alpha = 1 + P$$

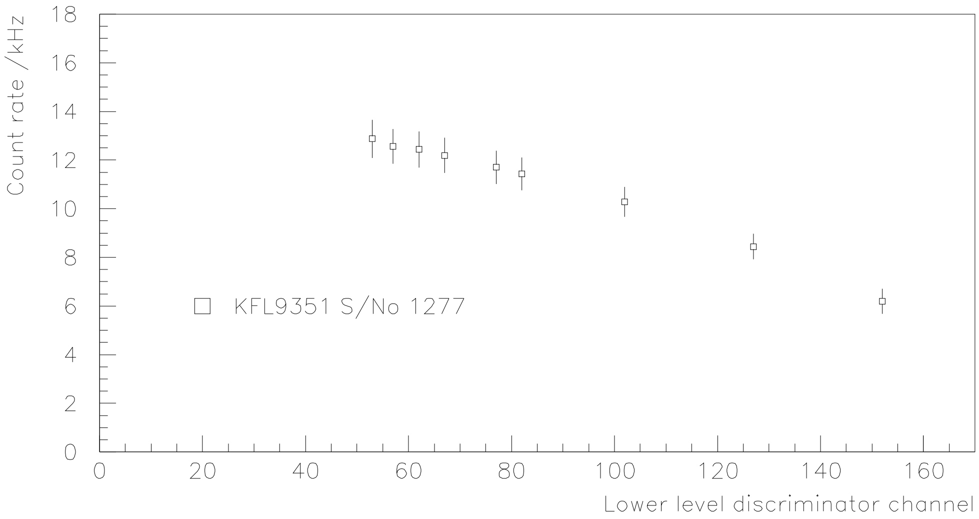

Figure 6.7: A graph of the change in counting rate against lower level discriminator. The fall above around 100 channels corresponds to the removal of events contained within the single photon peak of a charge spectrum. The error bars are multiplied by five for clarity.

where P is the probability of an afterpulse being produced by a primary pulse, then we have to divide our expression for η by α so as to obtain a true value. This argument was extended by including the probability that the afterpulse also produces an afterpulse itself, that this secondary afterpulse produces another and so on. Thus:

$$\alpha = \sum_{i=0}^{\infty} P^i$$

And as P is less than one:

$$\alpha = \frac{1}{1-P}$$

This factor is included as the general variation from tube to tube in afterpulse rates was well known to range between 1% and 5%, and any tube to tube effects that could be eliminated would give a better indication of any possible systematic variations between 9351 and 9353. The true value for the counting efficiency is thus:

$$\eta = 9.838 \times (1-P) \frac{R}{\int_{300nm}^{700nm} Q(\lambda) T(\lambda) \frac{d\lambda}{\lambda^2}}$$

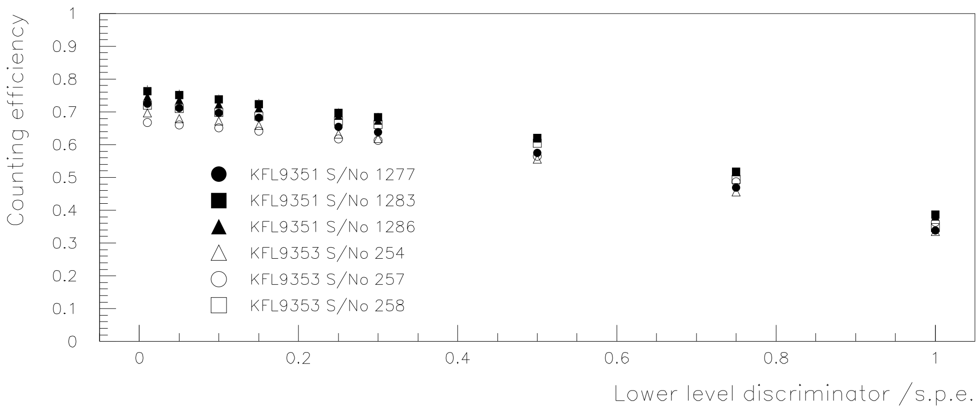

Figure 6.8: The distributions of single photon counting efficiency with lower level discriminator for the six tubes under test. The filled circles are the 9351 BOREXINO tubes, while the open circles are the 9353 AMANDA tubes.

The values for each of the six tubes are plotted in figure 6.8. The horizontal axis shows lower level discriminator in terms of the single photo electron peak, while the vertical axis shows the counting efficiency, here defined to be between zero and one. The filled symbols represent the three 9351's while the open symbols are the 9353's.

The obvious point to note is that while there is a definite tube to tube variation, they all seem to follow the general pattern as described for figure 6.7. The counting efficiencies at values of < 0.4 s.p.e. all seem to be roughly linear, and the gradients of these lines could be said to be approximately the same, although the higher efficiency tubes are slightly steeper.

The actual values of around 70% were in good agreement with previous experiments on tubes such as EMI 9350 [56], RCA 31000 [57] and EMI 9558 [58]. The difference between the 9351 tubes and the 9353 tubes was rather more difficult to confirm, but there did seem to be a slight increase for the 9351. This assertion is countered by the fact that the sample size is only three tubes of each type, and the spread within the tubes of the same type is of roughly similar proportions to the 9351/9353 difference.

If all these results are put together to obtain a mean counting efficiency, we obtain a value for the counting efficiency of around 67% for a lower level discriminator of one quarter of a single photo electron. Thus for every three photoelectrons produced, only two are detected, and if a peak Q.E. of 30% is included, the peak absolute efficiency is near 20%.

Errors - Random and Systematic

As stated before, the errors plotted on figure 6.7 are purely due to measuring the rate three times and then working out the mean and standard deviation for each of these three runs. As they were all taken within a few minutes of each other, effects such as the drifting of the supply voltage could be ignored, whereas the gradual change in noise rate often seen in large tubes like the 9353/1 series could be discounted for two reasons: firstly, the tubes were not used without being allowed to settle inside the gear for at least five minutes; secondly, all tubes had been previously dark-conditioned by being kept in black boxes at low tungsten light levels for at least 24 hours. This means that the average of the runs is fairly good, and standard errors on the means of less than 1% were obtained.

The magnitude of the errors on each reading were seen to be fairly constant for each tube. Thus one tube would show a 2% error at all 9 lower level discriminator points, whereas another might have 0.5% error. This deviation was also seen to be related to the noise rate: high noise meant high errors, whereas low noise meant low errors. This goes against a naive root N expectation, where larger noise rates would lead to higher statistics. However, the stochastic nature of noise events means that $\sqrt{N}$ errors are not applicable. Having said all this, the percentage random error is only 2% at its very largest, and although this could have been improved by taking more runs, the largest single systematic was around twice as large as this, and further runs were not deemed necessary.

Other sources of error on the counting efficiency were considered, including the positioning of the single photo-electron peak on the ADC, the measurement of the solid angle subtended by the exposed area of the PMT and the calibrations of the quantum efficiency and acrylic transmission measurements. One of the largest errors came from the uncertainties involved in the Monte Carlo modelling of the source itself, and when all these were combined, the total systematic error was:

$$(\frac{\sigma(\eta)}{\eta}) = 4.3\%$$

With random errors ranging from 0.5% to 2%, total errors on each data point could be given, but an approximation on the error of each point may be said to be 5%.

The single photon counting efficiency.

The difficulty of defining a true value for any efficiency number in PMT's has already been emphasized, and to conclude with a solitary number is ill-advised. The change in counting efficiency due to discriminator setting is both interesting and useful in the optimization of detectors involving large numbers of these tubes, and the overall performance seems to be consistent with previous data, both on the EMI9350, the close relative of the 9351/3 tube, and other large area PMT's, such as the Hamamatsu R1408.

The most obvious result is that the method gives reproducible results with a good degree of accuracy, and although the testing process takes a certain amount of time per tube, not lending itself to a production test, the possibilities of tailoring such efficiencies by adjusting lower level discriminator are clear.

6.3 Late Pulsing

The appearance of late pulses at the anode has been observed in many different phototubes, and with the use of tubes in such timing experiments as BOREXINO and SNO, the phenomenon is of particular interest. Normal transit time jitter is of the order of 1.5 nanoseconds (σ), but late pulses of the order of twenty or thirty nanoseconds, as observed below, offer particular problems. In order to identify and explain such behaviour it is important to check the variation of the late pulse's position in time and the fraction of a pulses which are late with the parameters available, in this case the voltages across the dynode stack.

One good candidate for the late pulse process could be the elastic backscattering of the primary photoelectron from the first dynode, without the production of secondary electrons. The scattered electron then travels back up to the photocathode before being re-accelerated down to the first dynode again, where it can then cause the multiple scattering which gives an anode pulse. Such a late pulse's position in time with respect to the 'primary' or main peak would then depend strongly on the voltage between the cathode and first dynode, V_{KD1}. There would also be the possibility that such a process could occur twice to the same photoelectron, thus leading to a further late peak at twice this characteristic time. The probability of such a double backscatter should correspond to the square of the initial chance of a single backscatter occurring. Such a characteristic is looked for below, although other explanations are not excluded. It must be said however that for the bulk of the text below, the elastic backscattering argument is assumed to be correct.

Transit Time Measurement

The measurement of transit times is a well known technique involving a pulsed light source and a coincidence or time difference measurement between the tube under investigation and a monitor tube of well known properties. Although coincidence measurements lead to a reduction in the random background across the time window, the same effect can be obtained by starting the time measurement on the discriminated output of the large[2] tube and stopping on the pulse from the monitor tube. The random noise from the large tube thus generates either a flat background, or an overflow condition in the time window. The signal to noise ratio from the large tube under test was usually of the order of ten, and random coincidences were to a large extent obviated. Also the dead time associated with the overflow was largely ignored due to the speed of the pulsed source which meant large amounts of data could be collected in fairly short time intervals. This speed reduces such effects as gain drifting and temperature effects. Dead times were not measured as the late peaks are normalised by the main peak, and not by a rate measurement.

The pulsed source used was a hydrogen discharge lamp with a 420 nanometre filter and a neutral density filter inserted in the light path from the source to the hemispherical tube, which was illuminated uniformly over the whole face. The reduction in the lamp flux per pulse with the neutral density filter was primarily to ensure that only single photon events triggered the large tube. The importance of such a discrimination is obvious in that if, for example, two photoelectrons were generated within a few nanoseconds of each other and only one was backscattered, then the non-backscattered electron would provide a start signal and the timing measurement would occur, while the backscattered electron would be ignored due to the dead time of the discriminator. This would bias the amount of counts in the main peak compared to the late peak, and thus give an incorrect measurement of the backscattering probability.

The timing measurements were carried out using a LeCroy QVT which was then

[2]The term large in this context applies to the hemispherical tube under test.

connected to a PC via a CAMAC crate. The data sets thus obtained were saved and analysed at Oxford. The calibration of the timing gear was not measured by the author, but was said to be of the order of 1% from previous work. All the tubes were run with mu-metal shielding to reduce magnetic field effects on the timing. The discriminator on the monitor tube was set at a prescribed level, while the large tube had a low-level discrimination of $\frac{1}{4}$ of a single photoelectron but no upper level. Both discriminators were constant fraction devices to preclude the time walk associated with edge-level triggering.

Results

The analysis of the time spectra is one of the most crucial aspects of the determination both of the time differences and the probabilities of backscattering. All time spectra show the main 'true' peak on the left hand side, with increasing time going towards the right. As an example a data set is shown in figure 6.9.

As can be seen, this plot shows the time channels of the QVT on a log scale. The actual spectrum was obtained from around half an hour's data-taking with one of the EMI 9353 tubes. The data in the primary peak is actually larger than the QVT maximum data size (2^{16}), and the overflowing of such channels is shown as a flattened top. It should be noted that the very first channel on this plot is the overflow bin, and should be ignored, and that the time calibration for this run is 200 ns full scale, corresponding to 0.78 ns per channel.

The time spectrum shows some obvious structure. The first peak is due to normal pulses. The second peak, at around 90 channels, is very asymmetric in time, with a far slower leading edge. The third feature is the plateau between 110 and 160. This is significantly higher than the random background (which can be seen on the extreme right) and has a fairly sharp cut-off. On a quick inspection it can be seen that this cut-off, the first late peak and the main peak are separated by the same time, namely 60 channels, or 45 ns. On a grossly simplified model of the electron optics between photocathode and first dynode we can say:

$$t \approx \sqrt{\frac{2md^2}{eV_{KD1}}}$$

where t is the cathode to first dynode transit time, d is the cathode to first dynode distance and m is the electron mass. This assumes a linear E field between cathode

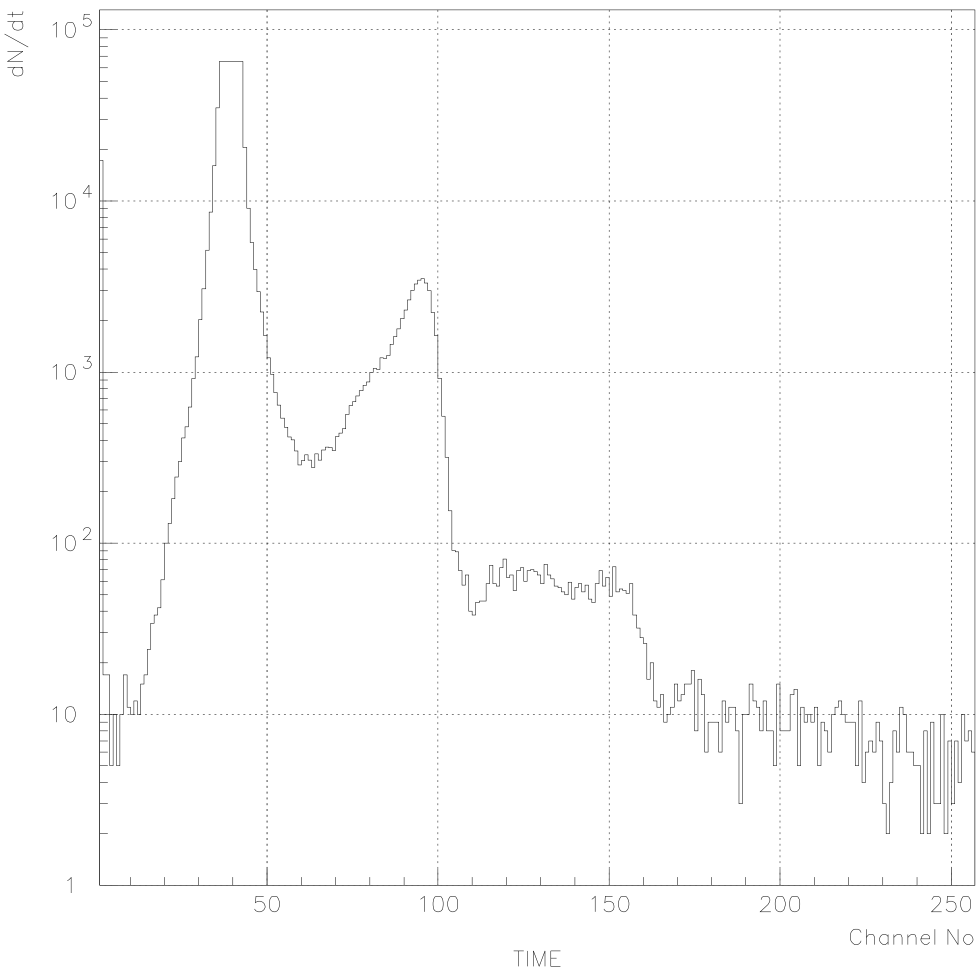

Figure 6.9: An example time spectrum. Notice the (artificially flattened) primary peak at the channel 40, followed by a smeared peak at around channel 90. There is also a third feature at around channel 160. The flat background over the whole range is due to random coincidences.

and dynode and zero velocity for emitted photoelectrons. In the above case this gives a cathode-first dynode transit time of around 17ns. Thus the first late peak seems to correspond to around two transit times, while the second peak is around four KD1 transit times, as expected for the elastic backscattering explanation. Better models including the full photocathode and non-zero initial electron velocities increase the predicted transit time, but the proportionality to $\frac{1}{\sqrt{V_{KD1}}}$ is still good.

To measure the late peak probability is slightly more difficult, in that limits have

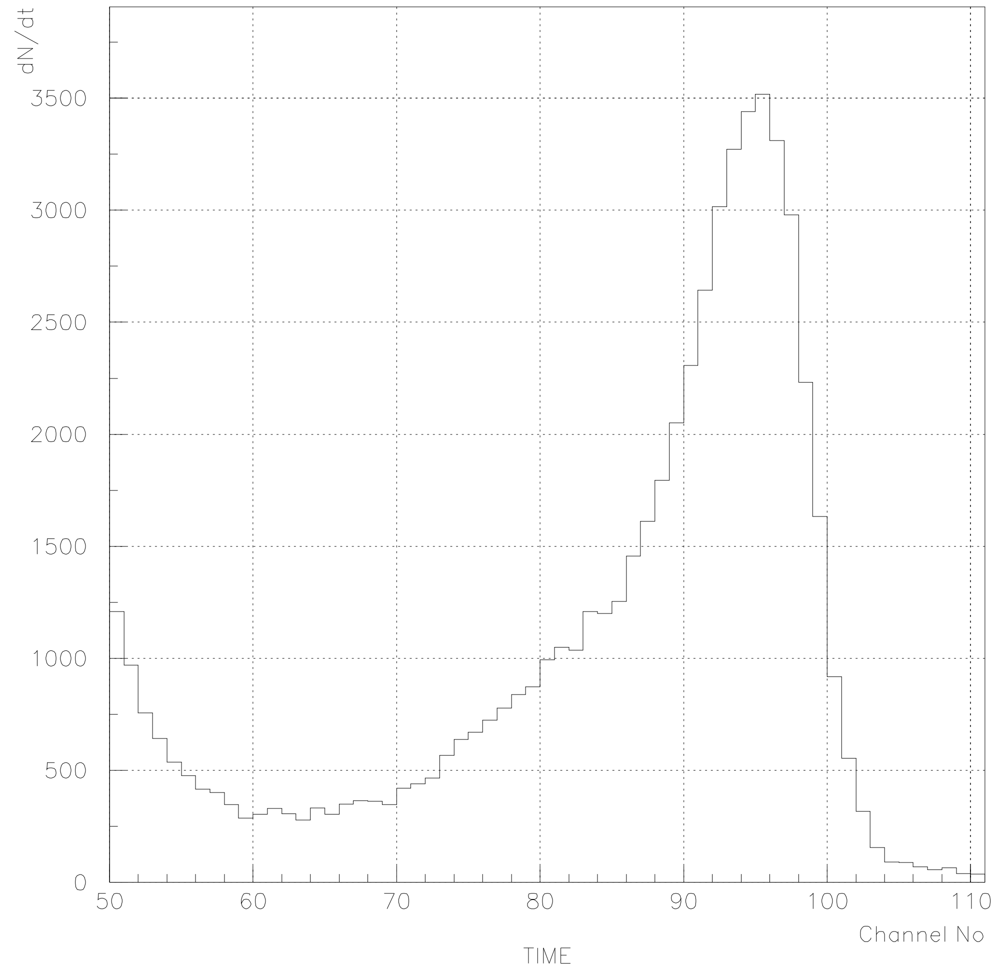

Figure 6.10: The first late peak. The measurement of the number of events in this peak is the integration of the graph from the valley at channel 60 to the point on the other side of the peak which corresponds to the same height as this valle, in this case at around channel 100.

to be set on the windows which correspond to the first and second late pulses. In the remainder of the report a systematic application of the areas shown in figures 6.10 and 6.11 is applied.

For the first late peak the number of counts is integrated from the valley, here at around channel 60, through to the point on the other side of the peak which corresponds to the same height as the valley, here at around channel 102. This gives a fairly good estimation of the total number of late events, and also is independent of the peak shape

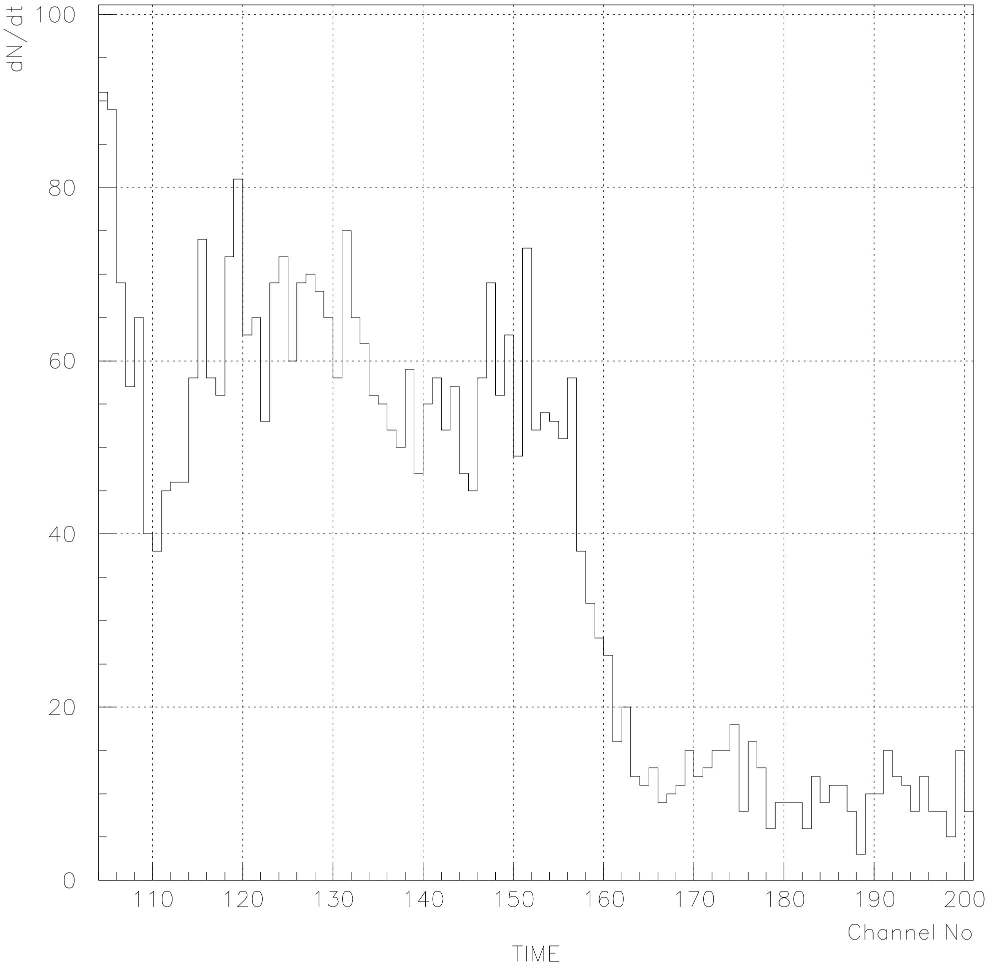

Figure 6.11: The second late peak. The distribution is fitted with a simple step function, and the area of the feature calculated from the rectangular area.

itself. This ensures a fair degree of consistency between different voltages and different tubes. The time of the first late peak is taken as the time of the highest channel in this range.

The second late peak is not really a peak at all, and a simple step function fit is used to give the number of events. The position of the edge gives the time position of the second late peak, and the random background is found from the height to the right of this cut-off. This background is subtracted from the plateau and from the first late peak. Possibly better descriptions could be made of this plateau, but the application

Tube Number	KD1 Volts	Delta t /ns	Late Peak 1 Area	Late Peak 2 Area	Main Peak Area	Ratio of peaks
571	900	43,44	39440	1501	–	3.8
"	450	57,59	96783	1809	–	1.9
"	300	72,75	52127	951	–	1.8
"	100	118	36284	–	–	–
"	900	42	451	–	190770	0.2
"	450	59	2469	–	253760	1.0
"	300	73	3844	–	342230	1.1
"	100	120	12554	–	787900	1.6
565	900	41,45	58970	2633	–	4.5
"	450	60,58	85887	2222	–	2.6
"	900	43	3407	–	216116	1.6
"	450	57	4490	–	275784	1.6
"	300	72	3112	–	331870	0.9
"	100	119	4238	–	245873	1.7
5inch	900	21.5	581	–	119645	0.5
"	750	21.9	774	–	102273	0.8
"	600	23.0	1551	–	198897	0.8
"	450	28.9	1866	–	177548	1.1
"	300	34.4	2989	–	221309	1.4
"	150	45.7	3210	–	321259	1.0
"	100	55.1	4385	–	386831	1.1
R1408	1049	40,37	110125	2590	–	2.4

Table 6.4: Analysed data for all the runs. The ratio of peak heights is in percent, whereas the size of the peaks themselves are in numbers of events. The delta t measurements are between the main peak and the first late peak and between the first late peak and the edge of the second feature.

of a simple fit again should reduce any systematic changes in the shape and allow a good comparison of one run to the next.

Table 6.4 lists the results for all data taken. The tubes investigated were two 8" EMI 9353 tubes (serial numbers 571 and 565), the new 5" EMI hemispherical tube (serial number 100), and a SNO type Hamamatsu R1408 (serial number CB0274). The time differences listed are given in nanoseconds, with requisite calibrations having been used.

Although there is rather a large amount of information in the table, we can look at pieces individually. As can be seen, there are two individual types of run associated with each tube/voltage combination. In the one type the main peak is contained within the range of the QVT, and a ratio of the first late peak to the main peak is obtained. In this case, the main peak area is found from a gaussian fit.

In the other type, the main peak is allowed to overflow the QVT and a ratio of the second late peak to the first late peak is measured. For each run a time difference is obtained, and sometimes for both main to first late peak and first to second late peak. It should be said that the appearance of a second peak was seen on all the tubes tested, although the analysis is not shown due to the inaccurate nature of some of the data.

The ratios of peaks do show some interesting variation. If we concentrate on tube 571, we can see that the first four runs seem to show the ratio increasing with KD1, whereas the second four show the ratio decreasing with increasing KD1. Appropriate, flat backgrounds have been subtracted from all the numbers, and so the changes are real.

In order to explain this away, we can say that as KD1 rises, the counting efficiency for photoelectrons rises also, and so the probability of multi-photoelectron events also increases. This would tend to increase the number of counts in the main peak, while not affecting the late peaks. Thus, it would seem that the rise in backscattering with rising KD1, as displayed by the ratio of the two late peaks, is the more significant of the two.

As a counter argument, it can be said that these backscattered electrons would also be subject to some collection efficiency, and this might be an even stronger function of KD1. Without a simultaneous measurement of the charge of the anode pulse it is difficult to separate these two possibilities.

Of course, if we use the Rutherford scattering cross-section:

$$N(\phi)d\phi = N_0 n t \pi \left(\frac{zZe^2}{4\pi\epsilon_0 m v_0^2}\right)^2 \frac{\cos(\phi/2)}{\sin^3(\phi/2)} d\phi$$

Then, with electrons accelerated through around 1KV, we can assume they are non relativistic, and so:

$$2E_e = mv_0^2 = 2eV_{KD1}$$

So that:

$$N(\phi)d\phi \propto \frac{Z^2}{V_{KD1}^2} \frac{\cos(\phi/2)}{\sin^3(\phi/2)} d\phi$$

And:

$$N'(\phi_0) \propto \frac{Z^2}{V_{KD1}^2} \cot^2(\phi_0/2)$$

where N' is the number of electrons scattered into an angle greater than ϕ_0. This shows a strong dependance with V_{KD1} which seems to be absent in the data, and no strong conclusions about the backscattering probability are drawn.

Having said this, the strongest evidence of a backscattering phenomenon is the times of the late peaks. The results are in figures 6.12 and 6.13 for the two eight inch tubes. The fit is:

$$t = \frac{P1}{\sqrt{V_{KD1}}}$$

where t is in nanoseconds and V_{KD1} is in volts. Thus:

$$P1 \approx 10^9 d \sqrt{\frac{2m}{e}}$$

where d is the cathode to first dynode distance, and m and e are the mass and charge of the electron respectively.

This one parameter fit is applied to the data, errors of 3 ns having been assigned to all values due to calibration and measurement errors. The value of P1 given in the top right of each plot is the fitted value. The chi-squared value is also shown, although this is not a very useful measure of statistical significance when the errors are assigned in such an arbitrary way. However, the two values of P1 agree closely, and this would suggest that the time differences are due only to a tube type's geometry and to the voltage between K and D1. Also, as the stack voltages were different, the time is not dependant on voltages after D1.

The next fit is shown in figure 6.14 for the five inch tube. Again a good fit is seen to the above function, this time with errors of 1.5 ns, as the measurements were taken on a smaller full scale time resolution. If we compare the parameter P1 with those of the eight inch tubes we should be able to relate this directly to the ratios of the distance from the cathode to the first dynode. The ratio of the P1's is 0.47, while the distances are in the approximate ratio 0.5. Thus it can be seen that the same effect appears in both tubes, and appears to be the same effect for both types.

It should also be said that although the second peak was not measured on the five inch tube, it was seen in the longer runs. However, due to the large amount of noise on

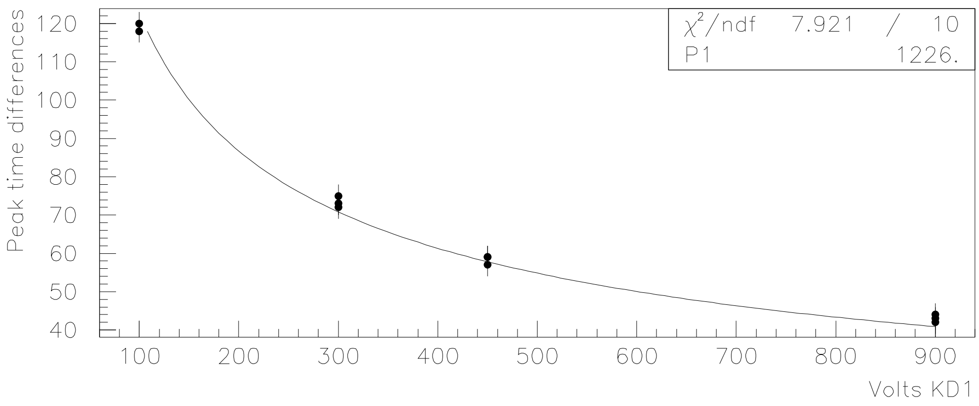

Figure 6.12: A fit of the time differences for Tube 571. The fit function is a simple one over the square root of the KD1 voltage, as is expected from naive determinations of the transit time.

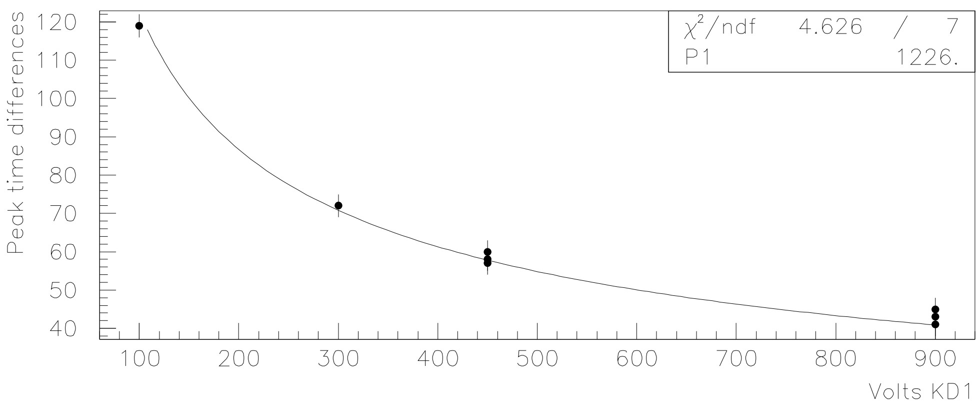

Figure 6.13: A fit of the time differences for Tube 565. Again, the fit is to the inverse of the square root of the voltage.

the tube the random background was such that measurements of the peak area were impossible. The R1408 run was somewhat ruined by a reflection due to the termination on the base. This caused ringing problems and led to a series of peaks overlaying the true spectrum. Still, both peaks were seen, and the time differences corresponded to two transit times.

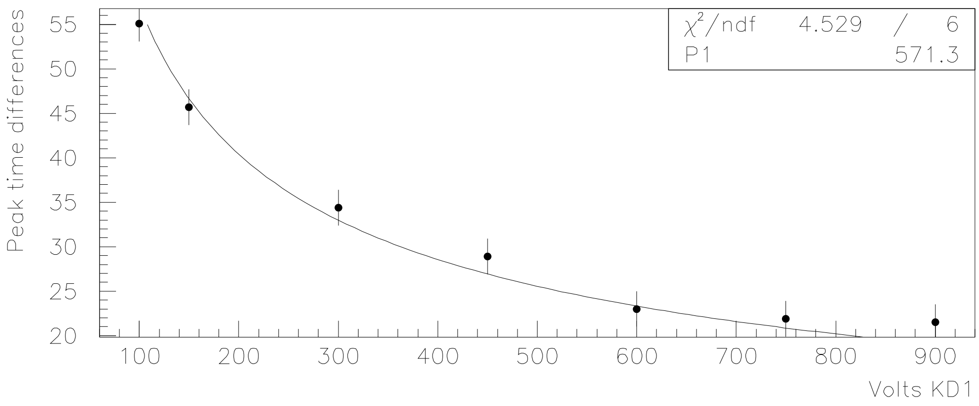

Figure 6.14: A fit of the time differences for the 5 inch tube. The value of the parameter should be proportional to the cathode to first dynode distance, and the comparison of this parameter with the eight inch fits is done in the text.

Conclusions

To conclude, we can say that the backscattering explanation for the late peaks seems to agree with the data in as much as the time differences between peaks are of the correct order. It should be said that these characteristic times are much shorter than ionic afterpulse times, and hence afterpulsing in this sense is excluded.

The evidence from the probability of late peaks occurring is much less certain, both due to the difficulty in measurement and in the theoretical behaviour of the ratio with respect to KD1 voltage. More work would have to be done on both these areas if a true comparison were to be drawn. At the moment it can be said that there appears to be a systematic increase, for single photoelectrons, with increasing KD1 voltage, although this seems to go against a naive treatment of the phenomenon with the Rutherford formula.

Chapter 7

Breakdown Noise in Photomultipliers

In this chapter we consider a problem that has arisen with the widespread use of large area photomultipliers: breakdown noise. Of course, it is well known that PMT's produce noise hits from many sources, some being time dependent and others temperature dependent. Such background effects as radioactivity in the glass and cosmic rays producing Čerenkov light as they travel through the glass can also produce uncorrelated signals[59].

More recently, the extensive use of large area photocathodes has revealed yet another effect, that of breakdown noise. This phenomenon seems to occur at some critical voltage, and produces huge increases in the single photoelectron noise rate. Such noise has now been seen in both the SNO tube, more specifically BB4458 mentioned in the previous chapter, and in three of the ETL eight inch tubes: the 9350, 9351 and 9353. The apparent similarity between these widely differing phototubes would seem to suggest a common cause that is due to some property of these large area tubes.

We will begin by considering the performance of the SNO breakdown tube more carefully.

7.1 Breakdown noise in the SNO tube

During the voltage variation experiments of the previous chapter, one tube stood out for particular attention with respect to the others, namely BB4458. Although at low voltages its performance was roughly equivalent with the other Hamamatsu tubes under test, as the cathode to first dynode voltage increased, the noise rate suddenly jumped by an order of magnitude. The critical voltage was further pinned down to the focus to first point, with the cathode to focus voltage apparently having no effect.

The increased noise was decidedly unexpected, and although it was only seen on one of the six tubes under test, reports of similar behaviour both by other SNO tubes under test in Canada, and the appearance of similar effects in ETL tubes meant that it was worthy of further investigation. The charge spectrum of these noise events was recorded, and apparently showed that the noise was due to single photoelectrons, rather than bursts of multi-photoelectron events, as would be expected for residual gas effects. This also suggested that the noise was originating from somewhere near the photocathode, as the noise pulses were being fully amplified down the dynode stack. This was slightly at odds with the evidence of the voltage variation tests, but at this point the other tubes were undergoing temperature tests, and BB4458 was also subjected to these experiments.

Temperature tests.

The study of noise rate versus temperature has been the subject of much previous published work, and the temperature tests on the SNO tubes were viewed as more of a qualitative confirmation of the results of such work, rather than a quantitative examination of the parameters.

Each tube was placed in a normal domestic fridge which both produced the temperature variation and also worked as a fairly good dark box. The temperature was measured using a standard thermocouple connected to a voltage to frequency conversion circuit. This allowed the measurement of the temperature via a CAMAC crate by connecting the frequency output to a scalar. The temperature probe/conversion circuit was calibrated over the range of temperatures required, and was shown to be linear to a fairly high degree of precision. The noise rate was measured using a standard discriminator/scalar combination, and the charge spectrum of the events was also recorded using a charge sensitive ADC gated on the discriminator.

The temperature variation was achieved by turning the fridge on at the mains and allowing it to fall to around 2 or 3 degrees. The tubes were held at this temperature for some time, usually overnight, and the fridge was then turned off and allowed to warm to room temperature. This obviated any noise pickup by the electrical switching within the fridge as it warmed. Of course, the temperature rose in an exponential

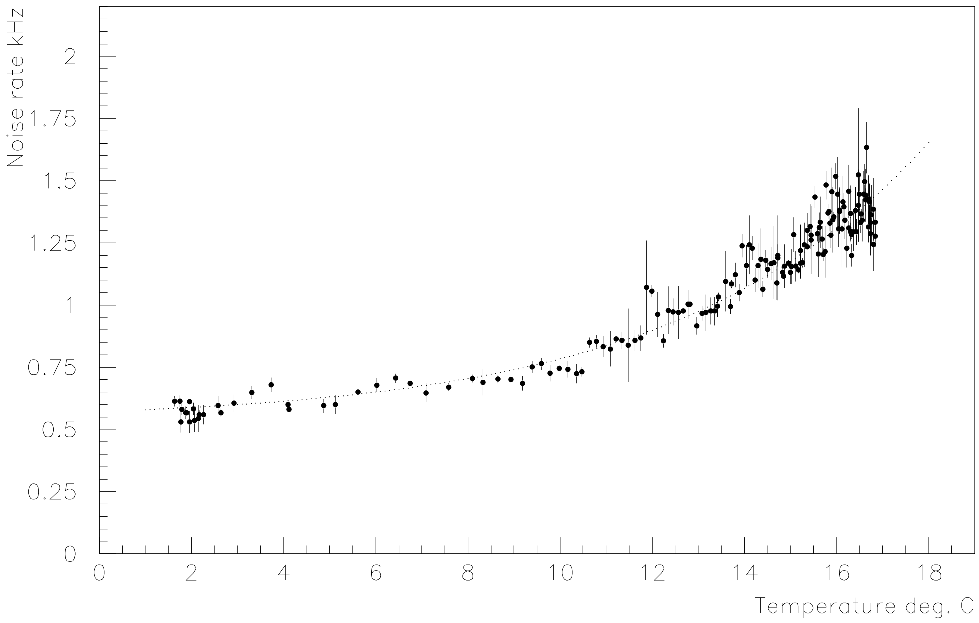

Figure 7.1: A typical noise rate against temperature plot for a SNO tube. As can be seen, the rate is almost flat at the outset, but then rises above 10 degrees, with an increase by a factor of around 2 at 20 degrees. Details of the fit are given in the text.

fashion rather than linearly, but the low thermal conductivity of the fridge meant that the entire warming cycle was of the order of hours, and hence many measurements of the rate could be achieved.

The temperature, noise rate, and fraction of hits above the software threshold of one quarter of a single photoelectron were thus taken at intervals of 5 minutes, with 10 rate measurements carried out and the mean and standard deviation calculated from this series of measurements. As an example of a 'good' tube, the noise rate against temperature plot for CC4292 is shown in figure 7.1.

The superimposed fit has the form:

$$\text{Noise Rate} = P_1 + P_2 T^n e^{-\phi/kT}$$

where T is the absolute temperature, and ϕ is the work function of the photocathode. The first coefficient P_1 is the statistical noise of the PMT, which is temperature independent. The second term comes about from the Richardson equation for

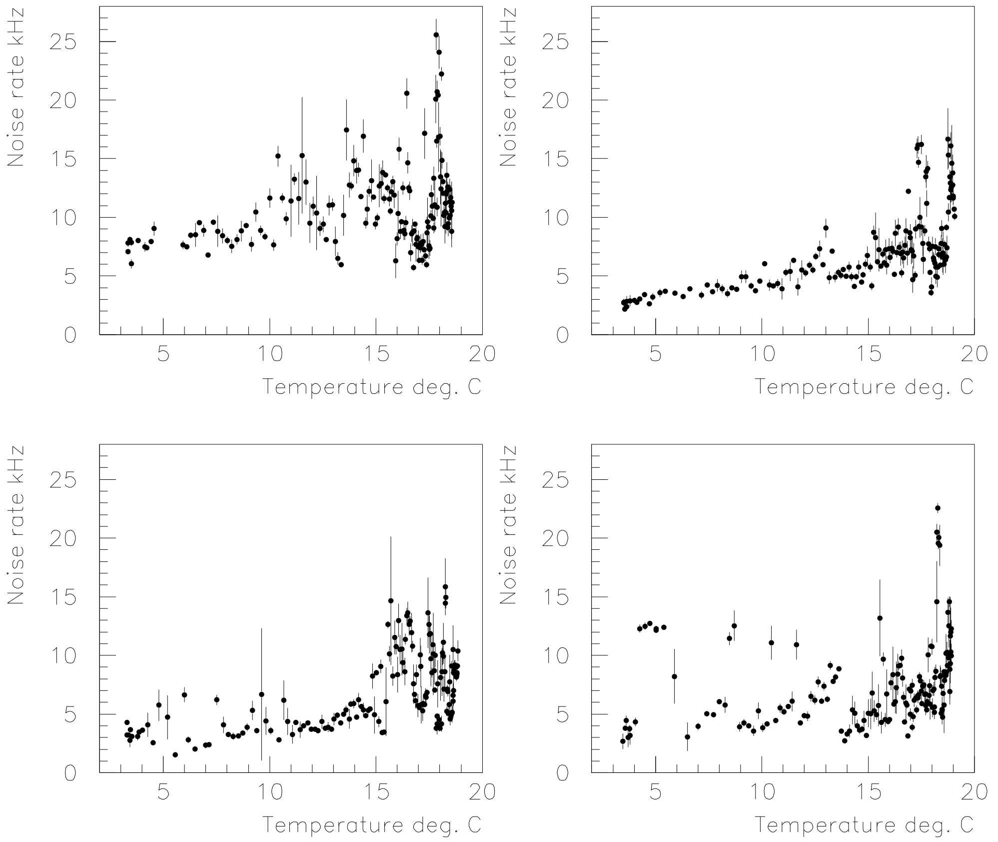

Figure 7.2: Four temperature variation runs for the breakdown tube, BB4458. As can be seen, the variation in noise rate is enormous, and apparently not a strong function of temperature.

thermionic emission from a conducting surface. The value of the index n is due to the type of surface: for metals, a value of 2 is used, whereas for semiconductors like the photocathode, a value of 5/4 is usually assumed[60]. The fit to the above data has three parameters, the two coefficients, and the work function. The fit gives a value for this of 1.35 eV, which is much as expected for a bialkali photocathode.

This form of rise with temperature is very well established, and when the breakdown tube was inserted into the fridge, it was mainly to see if the high noise rate was dependent on temperature in the same way as was expected. Four temperature runs are shown below in figure 7.2.

These figures again show noise rate against temperature, but the difference between these and the previous tube is obvious. Although there does seem to be a variation of the correct order with temperature of the underlying noise rate, the most interesting point is that the noise rate jumps up and down with time, and that these jumps seem to occur independent of temperature. It should be noted that as the fridge warmed exponentially, the difference in temperature between two successive measurements of the noise rate changed over the run. Essentially this means that there was more data taken at higher temperatures than at low temperatures. This means that we would expect to see more of these random jumps at higher temperatures than at the lower end of the scale, as is indeed the case.

When these random increases in the noise rate were seen, it became obvious that they had to be studied more closely. In order to do this, the tube was left in the fridge over a two week period. There was no enforced temperature variation during this time, with the fridge serving as nothing more than a dark box. The noise rate was measured at 10 minute intervals over this fortnight, and the time spectrum is shown in figure 7.3.

The variation in rate is again apparent, with huge increases in noise at apparently random intervals. It was apparent that the high noise rate of a breakdown tube was actually due to these large variations pushing the time averaged noise rate upwards.

Such noise rate increases can arise in phototubes through a number of differing causes. Perhaps the most well known is afterpulsing[61]. This occurs when the electron pulse ionizes a residual gas molecule within the phototube itself. This positive ion then travels back to the photocathode in times of the order of microseconds, and the ions subsequent impact can cause up to five electrons to be generated. These electrons are then accelerated back down to the anode in exactly the same fashion as photon-generated electrons, and can cause further afterpulses, so that the process can self propagate, leading to very high noise levels. This problem has been observed in tubes which are known to contain a large amount of residual gas, but such afterpulsing does not usually have the voltage dependence seen in breakdown tubes.

One of the first explanations of breakdown noise involved point defects at various places within the phototube. These defects can cause corona discharge, which would thus produce large amounts of electrons to be generated. Also, as the charge spectrum of the breakdown noise was observed to have a fairly well resolved single electron peak, it was thought that any such defect must be either on or near the photocathode itself, as electrons produced in the middle of the stack would not be fully amplified. If it

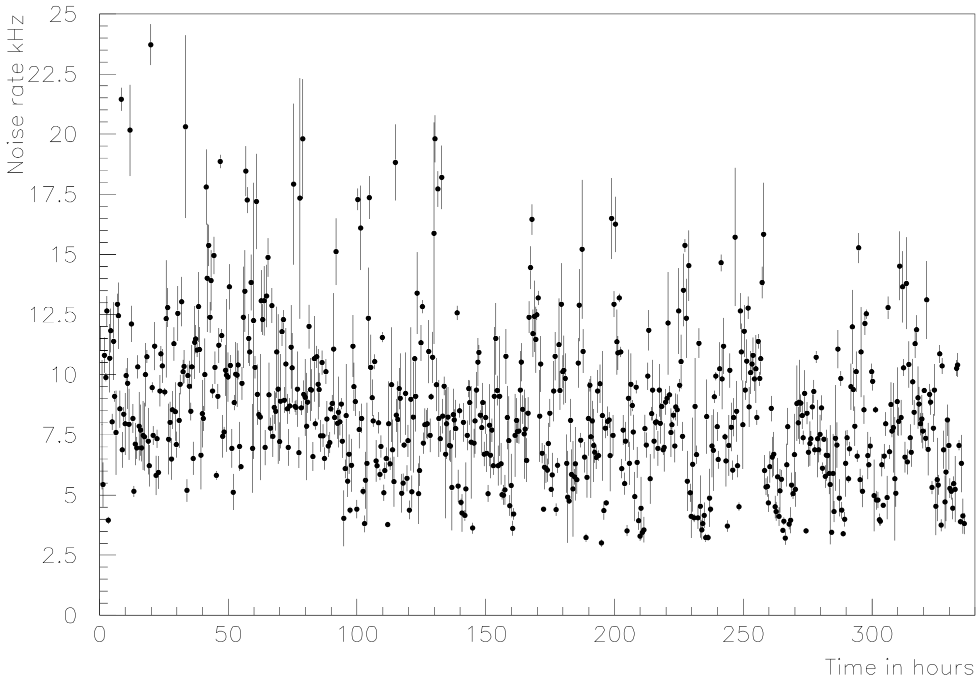

Figure 7.3: A time spectrum of the noise rate for the breakdown tube. The noise rate was measured at ten minute intervals over a fortnight. The variation in measured rate is apparent, with order of magnitude increases and decreases apparently at random.

were defects in the photocathode, it might also explain the appearance of breakdown noise in the eight inch tubes, where the photocathode area, and hence the probability of seeing such a defect, is that much higher due to the increase in photocathode area.

Although such explanations as defects in the photocathode do seem to do fairly well, and indeed would depend on the KD1 voltage, the fact that the BB4458 tube's breakdown did not depend on KF, but on FD1. It was at this point that it was thought basing a large amount of work on a single tube was fairly precipitative, and that other breakdown tubes were needed for investigation. Thus, work was done at ETL, which is detailed in the next section.

7.2 Breakdown noise in ETL 9350 tubes.

In order to use a representative sample, forty type 9350[1] tubes were selected from two main batches; the first batch had been previously tested and rejected by the production specification. They were classified as either BD or DC tubes, meaning that they displayed either breakdown behaviour or dark current readings above the maximum specification. These tubes had been variously altered (some had been hot-aged) and retested, but had still failed. The second batch had passed the specification, and were used as a control set.

Each tube was inserted into a standard test apparatus at ETL known as a flying spot gear. The light source was a normal tungsten lamp with a Corning blue filter tube in the light path. Each tube was ramped up to around 1 kV and the voltage was then fine tuned to give an anode sensitivity of 10 amps per lumen. In all these tests each tube was assumed to have a cathode sensitivity of 10 μA/lm, and so this first point corresponded to a single electron gain of 10^6. The light source was then switched out and the dark current measured using a picoammeter within the flying spot gear.

The light source was switched back on and the gain was then adjusted to 2×10^6, ie an anode sensitivity of 20 A/lm. The voltage and dark current were again recorded and this process continued up to at least 200 A/lm, having passed through the points at 50 and 100 A/lm. For most of the tubes further measurements were taken at anode sensitivities of 500 and 1000 A/lm, but on some of the breakdown tubes the dark current was high enough so as to lead to a noise in signal ratio of more than one, and hence not allow an accurate determination of the anode sensitivity.

The breakdown phenomenon is very variable and in order to consider the tubes behaviour in a systematic fashion, a figure of merit was defined as the dark current at 5×10^7 divided by the dark current at 5×10^6. This ratio should give a value of around 10 for a tube which has a constant background noise rate, whereas breakdown tubes will have a value in excess of 10. Those ratio values significantly different from ten are shown in table 7.1.

As can be seen, the anomalous tubes fall into two groups, those with dark current ratios well below 10 and those with ratios well above ten. The first group are leakage current tubes. In these, there is a conductive path between the cathode and the anode

[1] Also known as the MACRO tube due to its use in the MACRO cosmic ray experiment at Grand Sasso.

Tube Serial No	Ratio of dark currents
KFL8655	2
KFL8729	1
KFL8657	1
KFL8790	4
KFL8740	10000
KFL8762	20
KFL8384	72
KFL8784	110
KFL8791	260
KFL8706	267
KFL8707	2500
KFL8708	123

Table 7.1: A table of the dark current at 5×10^7 divided by the dark current at 5×10^6. Only those tubes which differed significantly from the expected value of ten are shown.

which provides a current path. This current is unaffected by the change in voltage between the two gain settings, and hence the ratio is close to one. This is a well known problem in manufacture, and is not addressed here.

Of those tubes which have ratios well above ten, we can see that there is a fair variation in the actual figure. The fact that the ratio is above ten indicates that the onset of breakdown noise lies somewhere within the two voltages corresponding to the two gains, but that the exact point of breakdown differed among all eight tubes.

Further Testing

More rigorous testing of the breakdown phenomenon proceeded by selecting six tubes for further investigation. Two non-breakdown tubes (8753 and 8774) were used as a control set for the tests, and of the tubes listed above, four were taken. These were divided into two groups of two tubes each: the mediocre tubes (8784 and 8384) which had ratios of the order of 100; and the bad tubes (8707 and 8740) where the ratios were larger than 1000. This distinction was made primarily in order to investigate whether the breakdown effect showed any similarity among these tubes.

The first tests involved taking more data for noise rate against gain curves for the six tubes. These distributions are shown in figure 7.4.

As can be seen, all the tubes start at reasonably low noise rates, with KFL8784

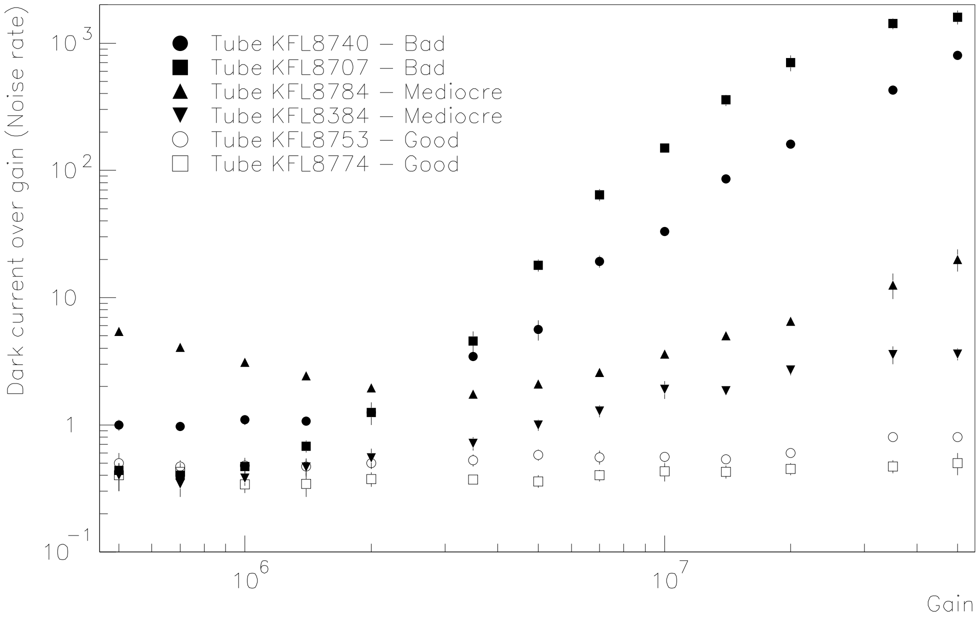

Figure 7.4: A plot of I_d over gain (which is equivalent to a noise rate) against gain for the six tubes. The difference between the three groups of tubes is evident, as is the difference between the tubes themselves. The y scale is in arbitrary units, with the value of one equating to around 4 kHz noise rate.

showing some evidence of leakage in the first part of its curve. However at around 5×10^6 the tubes begin to diverge and we can see the difference in behaviour between good, mediocre and bad tubes. The noise rates of the three groups at 5×10^7 differ by three orders of magnitude, showing the large difference not only between non-breakdown and breakdown tubes, but also among the breakdown tubes themselves. This difference is not just due to the point at which breakdown occurs, but seems to suggest that the strength of the breakdown is somehow different. Thus we can surmise that breakdown noise is not simply an all-or-nothing effect, but does have some fundamental parameters which determine its behaviour.

The error bars shown are by no means an indication of one sigma significance, but simply give indications as to how much the picoammeter fluctuated during the measurement of the noise rate. This fast (<1 second) fluctuation, as was seen earlier in the SNO PMT, was seen most clearly during breakdown, and the error bars show

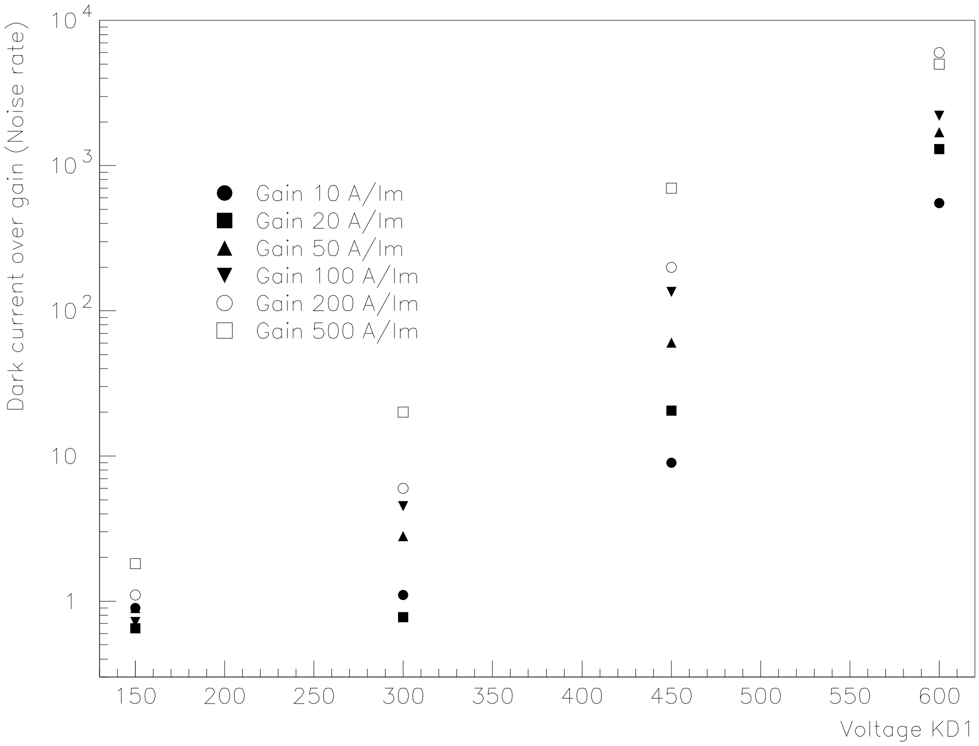

Figure 7.5: A plot of I_d over gain against the voltage KD1 for six different gains. The obvious effect of the KD1 voltage is seen, as well as an effect on the D1A voltage.

such this dependence.

The second set of tests were carried out in order to determine which was the critical voltage that caused breakdown. The KD1 voltage was varied in 150 volt steps using Zener diodes and measurements of the dark current were taken at six different anode sensitivities. The results are shown in figure 7.5.

Only the bad breakdown tube is shown: the good tube showed almost no change in noise rate over the same testing scheme, while the mediocre tube had much the same form as the above results, only with a slightly less drastic effect.

In order to fully appreciate the result, it must be noted that for a given gain, although the total voltage does increase with increasing KD1, the stack voltage, ie the voltage between the first dynode and anode, does not change by more than a few tens of volts. Thus similar symbols in figure 7.5 represent similar stack voltages. As can be seen, at the lower gains, breakdown occurs within the range of variation of KD1, while for the higher gains, the noise rate seems to rise almost linearly on the log plot, and hence exponentially, with KD1. The slope of this increase seems almost independent of the gain.

The change in noise rate for a given KD1 indicates that there must also be a component of the breakdown noise coming from the dynode stack. The total change in stack voltage on going from the lowest to highest gain was around 500 volts, or 50% of the stack voltage itself. Such an increase leads to an order of magnitude increase in noise rate, at least at the higher KD1 values. Perhaps the most interesting point to note from this is that the 9350 type tube has the focus dynode connected to the third dynode to improve transit time spread. Thus making the anode more positive with respect to the first dynode, ie increasing the stack voltage, increases the voltage between F and D1. To draw a conclusion that we could see the same FD1 dependence as in the Hamamatsu tube does seem slightly spurious, due to the very different construction of the dynode stacks in the two tubes.

Incidental experiments.

It was at this point that a series of slightly more speculative experiments were carried out. The first was to determine whether an apparently well behaved tube could be made to breakdown. In order to achieve this, tube 8753 had to be varied way away from its design criterion, with a KD1 value of 750 volts. The anode sensitivity of the tube was set at 5^7, but when the dark current was measured, an increase of only around a factor of 3 was seen in the noise rate, and this could easily be attributed to increased collection efficiency for thermionic electrons and such effects. It was clear that the increase was not breakdown noise in the same sense as the tubes tested above.

The next investigation was to see whether, on coming out of breakdown, there was a significant time lag between the voltage being reduced and the noise rate falling. One of the breakdown tubes was adjusted so that it could be switched from a breakdown regime into a non-breakdown regime within a second or so. When the noise rate was monitored during these changes, it was observed that the noise rate fell to the prescribed, non-breakdown level within one or two seconds. This time constant was attributed to the power supply itself, and hence any time constant associated with the onset of breakdown noise would have to be smaller than this value.

Finally, there was some investigation of hysteresis effects of the breakdown phenomenon. The voltage across one of the bad tubes was cycled up and down across the breakdown point, and the dark current noted on the upward-going and downward-going cycles. Also, the time that the tube was left in the breakdown state before coming back down in voltage was varied. All these tests showed no significant hysteresis effects, with the tubes remaining pretty much constant their breakdown behaviour.

Conclusions on the ETL 9350 tubes.

The main finding of these tests was the identification of the dependence on KD1. The slight contribution from the stack exacerbated the problem if the tube was already breaking-down, but the main point was that in order to stop a tube breaking down, one must lower KD1. This was also, in part, the remedy for the SNO Hamamatsu tube as well, and we have already seen the effects of such a change on the transit time spread and relative efficiency. Although the tubes are significantly different in construction, it is a fair to say that the effects of lowering KD1 would be detrimental on the ETL tubes as well.

Incidental findings are that there seems to be no evident decay time on going from a breakdown to a non–breakdown situation, and no hysteresis in the noise rate was seen on cycling the voltage between breakdown and non–breakdown voltage values.

7.3 Breakdown noise in the ETL 9353 tube.

The final part of this chapter deals with another of the range of eight inch tubes manufactured by ETL. The tube differs from the 9350 by having a slightly different glass bulb shape. This difference comes about due to the 9353's use by the AMANDA collaboration, a cosmic ray experiment at the south pole which involves burying the tubes in 1.2 kilometers of ice. The bulb shape for the 9353 is consequently designed to be slightly stronger in order to resist the pressures at such depths. The 9353 also has fewer dynodes than the 9350, but the structure of the first few dynodes, where the breakdown phenomenon is suspected to occur, is exactly the same.

The first tests on these tubes were again conducted in a flying spot gear as for the 9350 tubes, but the signals from the tubes were then fed into an ADC and the charge spectra studied. Due to the amplifier within the ADC, only the higher gain conditions could be observed with any certainty, which meant that the tubes were usually in

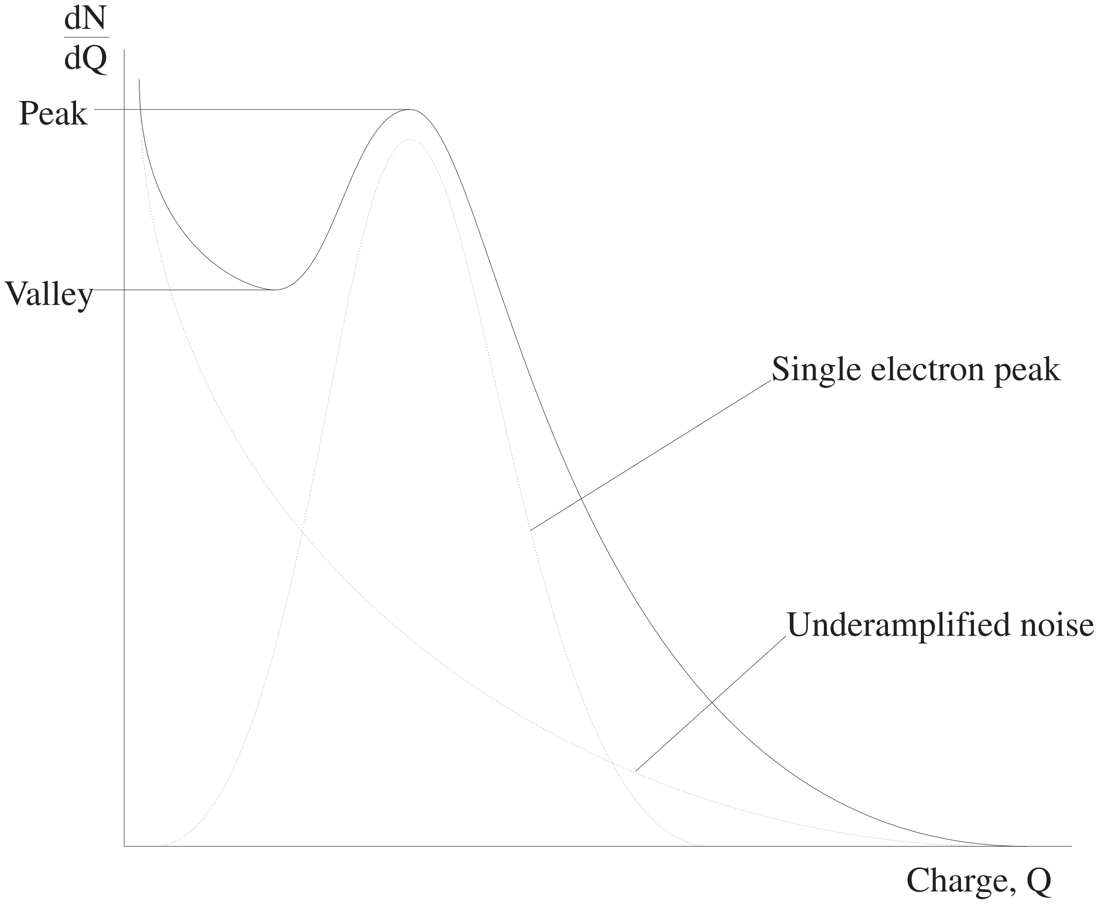

Figure 7.6: A schematic example of the charge spectrum. The main peak is due to the fully amplified cascades, while the exponential comes about due to statistical under-amplification within the first few dynodes. The peak and valley are shown on the y-axis, these two heights being used to form the single electron response, SER, or Peak-to-Valley ratio.

breakdown before a charge spectrum could be taken. Also, the dead time associated with the conversion within the ADC meant that noise rates up to around 60 kHz could be studied with certainty.

However, even with these limitations, some general observations were possible. The most important of these were that of the four tubes tested, two showed good single electron peaks during breakdown, while the other two showed a large increase in underamplified noise. This is important as such a distinction could not be seen by simply measuring the dark current, as in that case we are measuring the integration of the charge spectra.

The two components of the charge spectrum come from different effects within the regions of the dynode stack. If we see an increase in both the exponential and the gaussian, it indicates that the noise is being generated in the region of the photocathode, and the noise events are amplified in the same way as light-induced electrons. Such events will also be underamplified due to the Poisson distribution of the secondary electrons from the first dynode, so that we will see the exponential background as well as the full peak.

However, if an increase in the exponential background is seen without a corresponding increase in the gaussian, this would indicate that the extra noise events are being generated within the dynode stack itself, and are not being properly multiplied by the dynode stack. In order to form an impression of the relative magnitudes of the two distributions without complicated fitting of the full spectrum, we can form the single electron response (SER), also known as the peak-to-valley ratio, from the heights of the distribution at the two turning points. Hence a high SER means that the peak is well resolved, while a low SER implies either excessive under-multiplied noise or poor first dynode gain.

As was said above, two of these tubes had noise SER's of the order of 2, which was perhaps slightly lower than normal, but not wildly different, while the other two tubes had SER's that were almost unmeasurable, due to the poor resolution of the peak in the charge spectrum. This would seem to show that the former two had a source of breakdown noise near the photo-cathode, while the latter two had sources of noise somewhere within the dynode stack itself. This is interesting when compared with the voltage variation experiments on the admittedly different 9350 tubes above. There, it was seen that there were two voltage components for the breakdown dark current, which is of course independent of the spectral shape of the individual events, one KD1 effect and one D1A effect. This would seem to tally with two distinct sources of breakdown noise, one at the cathode, and another within the stack.

Further tests on the 9353 tubes consisted of temperature variation tests in much the same fashion as the SNO tubes detailed earlier in this chapter. Although a larger variation in temperature was available in these tests, the standard gear at ETL was not really set up to investigate the problem to quite the level required. Although the results were similar to those obtained with the SNO tube, with a decrease in underlying noise rate with decreasing temperature, and the breakdown behaviour not obviously temperature dependent.

7.4 Conclusions

Although the tests detailed are extensive, it is also obvious that no one candidate for the breakdown phenomenon emerges. Indeed, the very nature of the breakdown noise phenomenon makes it difficult to pin down whether we can accurately relate one breakdown tube to another. We can summarise what seem to be the major points that apply to breakdown tubes:

1. Breakdown noise seems to arise mainly in eight inch PMT's, and is characterized by a sharp increase in noise rate at some critical voltage. This breakdown voltage is specific to the tube itself, and seems to be related mainly to the inter-dynode potentials between the cathode and the first two dynodes.

2. The breakdown phenomenon consists of large, random increases in noise which last for periods of the order of milliseconds. Between these bursts of events, the tubes seem to behave in roughly the correct manner with respect to their noise behaviour.

3. The rate of breakdown seems unaffected by the temperature at which the PMT is operated. This excludes some of the more obvious effects commonly known to produce noise, such as thermionic emission from the photocathode, and excessive afterpulsing. This second statement comes about as there is a temperature component to the amount of residual gas within the tube which should have been apparent in the temperature regime studied.

4. The charge spectrum of the noise itself seems to vary, with some tubes showing good SER's, indicating cathode defects, while others have very low SER's, pointing to the dynode stack as the source of the noise.

5. Breakdown noise has no obvious hysteresis effects on the tube, with the noise rate falling and rising with the voltage, with no obvious time lag between the two.

6. The breakdown noise effect seems only to affect large area PMT's, with the problem only coming to light as these tubes have started to be produced in large quantities. Certainly, no such effect has been seen in one or two inch tubes, although this could be a function of these smaller tubes being usually operated at much smaller gains.

7. Finally, the similarity in the behaviour of the SNO Hamamatsu R1408 tube, which has a Venetian blind dynode structure, and the ETL 935– range, which has a linear focused dynode stack, would indicate that there must be some new, specific property of these large area tubes which is producing the breakdown effect.

Although this chapter is rather inconclusive, it is obvious that in order to pin down the breakdown effect, yet more characterization is needed, with the obvious avenues being the study of the good SER, bad SER difference; and the continued study of the characteristic times involved in the appearance of breakdown and the breakdown pulses themselves. The need for further work is obvious, but it is also clear that the path to take is not to well sign-posted.

Appendix A

Matter Enhanced Oscillations.

Here, neutrino oscillation is considered in a medium of constant electron density, N_e. As the electrons can scatter via W boson exchange, an interaction Hamiltonian is added to the Schrodinger equation describing the propagation states. Thus, again taking the case for just two neutrino flavours, ν_e and ν_x:

$$\begin{bmatrix} \nu_e \\ \nu_x \end{bmatrix} = \begin{bmatrix} \cos\theta_v & \sin\theta_v \\ -\sin\theta_v & \cos\theta_v \end{bmatrix} \begin{bmatrix} \nu_1 \\ \nu_2 \end{bmatrix} = R(\theta_v) \begin{bmatrix} \nu_1 \\ \nu_2 \end{bmatrix}$$

where R is the normal rotation matrix, and the e and x states are the flavour eigenstates, and 1 and 2 are the vacuum eigenstates. These vacuum eigenstates are determined by the Schrodinger equation:

$$i\frac{d}{dt}\begin{bmatrix} \nu_1 \\ \nu_2 \end{bmatrix} = \begin{bmatrix} E_1 & 0 \\ 0 & E_2 \end{bmatrix} \begin{bmatrix} \nu_1 \\ \nu_2 \end{bmatrix} = H_{vac} \begin{bmatrix} \nu_1 \\ \nu_2 \end{bmatrix}$$

just as before. Here H_{vac} is the vacuum Hamiltonian. The interaction term in the presence of electrons, H_i is given by[25]:

$$H_{int} = \begin{bmatrix} \sqrt{2}G_F N_e & 0 \\ 0 & 0 \end{bmatrix}$$

This matrix is expressed in the flavour basis, with G_F the Fermi coupling constant and N_e the electron density. Thus in the flavour basis, the Schrodinger equation is:

$$i\frac{d}{dt}\begin{bmatrix} \nu_e \\ \nu_x \end{bmatrix} = R(\theta_v) H_{vac} R^\dagger(\theta_v) \begin{bmatrix} \nu_e \\ \nu_x \end{bmatrix} + H_{int} \begin{bmatrix} \nu_e \\ \nu_x \end{bmatrix}$$

Now, we have a non diagonal Hamiltonian, but we can rotate the states to diagonalize by introducing a new parameter, θ_m, and two new states, ν_{1m} and ν_{2m}, where the m stands for matter. So we then have:

$$\begin{bmatrix} \nu_e \\ \nu_x \end{bmatrix} = \begin{bmatrix} \cos\theta_m & \sin\theta_m \\ -\sin\theta_m & \cos\theta_m \end{bmatrix} \begin{bmatrix} \nu_{1m} \\ \nu_{2m} \end{bmatrix} = R(\theta_m) \begin{bmatrix} \nu_{1m} \\ \nu_{2m} \end{bmatrix}$$

and hence:

$$H_m = R^\dagger(\theta_m)(R(\theta_v H_{vac} R^\dagger(\theta_v) + H_{int})R(\theta_m)$$

As we have defined the matter states to be where the Hamiltonian is diagonal, we end up with the requirement that:

$$(E_2 - E_1 \sin(\theta_v - \theta_m)\cos(\theta_v - \theta_m) = \sqrt{2} G_F N_e \sin\theta_m \cos\theta_m$$

So that:

$$\sin^2 2\theta_m = \frac{\sin^2 2\theta_v}{\sin^2 2\theta_v + (\alpha - \cos 2\theta_m)^2}$$

where $\alpha = \sqrt{2} G_F N_e/(E_2 - E_1)$. If we again say that the neutrino's mass is very much smaller than the energy then:

$$\alpha = \frac{\sqrt{2} G_F N_e E}{\Delta m^2}$$

Now, for the resonance condition, we have:

$$\alpha = \cos 2\theta_v$$

$$\Delta m^2 = \sqrt{2} G_F N_e E$$

$$N_e(res.) \approx 66 N_A \cos 2\theta_v \frac{\Delta m^2}{10^{-4} eV^2} \frac{10 MeV}{E}$$

Appendix B

Calculation of activities from fitting time spectra

If we start with the three linked radioactive decays:

$$\frac{dN_{Th}}{dt} = -\lambda_{Th} N_{Th}$$

$$\frac{dN_{Ra}}{dt} = \lambda_{Th} N_{Th} - \lambda_{Ra} N_{Ra}$$

$$\frac{dN_{Pb}}{dt} = \lambda_{Ra} N_{Ra} - \lambda_{Pb} N_{Pb}$$

then we can say immediately:

$$\lambda_{Th} N_{Th}(t) = P_1 e^{-\lambda_{Th} t}$$

where the constant P_1 is the Th activity at t=0. Introducing this formula into the Ra equation:

$$\frac{dN_{Ra}}{dt} = P_1 e^{-\lambda_{Th} t} - \lambda_{Ra} N_{Ra}$$

$$\Rightarrow \lambda_{Ra} N_{Ra}(t) = P_2 e^{-\lambda_{Ra} t} + \frac{\lambda_{Ra} P_1}{\lambda_{Ra} - \lambda_{Th}} e^{-\lambda_{Th} t}$$

where P_2 is a constant. For the Pb equation, the introduction of the above formulae gives:

$$\frac{dN_{Pb}}{dt} = P_2 e^{-\lambda_{Ra} t} + \frac{\lambda_{Ra} P_1}{(\lambda_{Ra} - \lambda_{Th})} e^{-\lambda_{Th} t} - \lambda_{Pb} N_{Pb}$$

leading to:

$$\lambda_{Pb} N_{Pb}(t) = P_3 e^{-\lambda_{Pb} t} + \frac{\lambda_{Pb} P_2}{\lambda_{Pb} - \lambda_{Ra}} e^{-\lambda_{Ra} t} + \frac{\lambda_{Ra} \lambda_{Pb} P_1}{(\lambda_{Pb} - \lambda_{Th})(\lambda_{Ra} - \lambda_{Th})} e^{-\lambda_{Th} t}$$

Now, the counter actually measures $\varepsilon\lambda_{Pb}N_{Pb}$, as opposed to $\frac{dN_{Pb}}{dt}$, where ε is an efficiency constant. If we make the approximations:

$$\lambda_{Pb}, \lambda_{Ra} \gg \lambda_{Th}$$

$$e^{-\lambda_{Th}t} \approx 1$$

then the rate is:

$$Rate = \varepsilon(P_3 e^{-\lambda_{Pb}t} + \frac{\lambda_{Pb}P_2}{\lambda_{Pb} - \lambda_{Ra}} e^{-\lambda_{Ra}t} + P_1)$$

As can be seen, if the source is in equilibrium P_3=P_2=0, and a flat spectrum with respect to time is seen. So the fit has the form of three competing exponentials, and at t=0:

$$\lambda_{Th}N_{Th}(0) = P_1$$

$$\lambda_{Ra}N_{Ra}(0) = P_2 + P_1$$

$$\lambda_{Pb}N_{Pb}(0) = P_3 + \frac{\lambda_{Pb}P_2}{\lambda_{Pb} - \lambda_{Ra}} + P_1$$

Also, the constraints

$$P_1 \geq 0$$

$$P_1 + P_2 \geq 0$$

$$P_1 + \frac{\lambda_{Pb}P_2}{\lambda_{Pb} - \lambda_{Ra}} + P_3 \geq 0$$

can be applied in the fitting process, as negative activities are unphysical. In order to ease the calculation of the activities at t=0, we can transform the coefficients as below:

$$Rate \propto P_3 e^{-\lambda_{Pb}t} + \frac{\lambda_{Pb}P_2}{\lambda_{Pb} - \lambda_{Ra}} e^{-\lambda_{Ra}t} + P_1$$

Now:

$$\lambda_{Th}N_{Th}(0) = A_{Th}$$

$$\lambda_{Ra}N_{Ra}(0) = A_{Ra}$$

$$\lambda_{Pb}N_{Pb}(0) = A_{Pb}$$

$$\frac{\lambda_{Pb}}{\lambda_{Pb} - \lambda_{Ra}} = \alpha$$

$$P_1 = A_{Th}$$

$$P_2 = A_{Ra} - P_1 = A_{Ra} - A_{Th}$$

$$P_3 = A_{Pb} - P_1 - \alpha P_2 = A_{Pb} - A_{Th} - \alpha A_{Ra} + \alpha A_{Th}$$

So that:

$$Rate \propto \left[A_{Pb} e^{-\lambda_{Pb} t} + A_{Ra} \alpha (e^{-\lambda_{Ra} t} - e^{-\lambda_{Pb} t}) + A_{Th} (1 - \alpha e^{-\lambda_{Ra} t} + (\alpha - 1) e^{-\lambda_{Pb} t})\right]$$

Appendix C

A summary of the running modes for SUF membranes.

The running modes for a Seeded Ultra-Filtration system are detailed below:

- Recircultaion Mode. This is shown in figure C.1. Here, the back pressure valve, valve G, is set so as to equalize the rates of water flowing along the membrane and through the membrane. Both these streams are fed back to the reservoir. This running mode is used for many of the procedures carried out on the membrane, including the deposition of the HTiO seed, the elution of the membrane, and, in the full scale system, for the deuteration of the membranes. This last function is anticipated for the full scale system so as to avoid the need for deuterated acids and alkalis. It becomes necessary due to the dead volume associated with the porous structure of the membrane. Hence, the light water contained within the pores at the end of chemical processing is flushed out using repeated washes of heavy water. This requirement comes about due to the presence of hydrogen in the detector volume severly reducing the neutron capture probability on the chlorine.

- Fast flush cleaning. In order to restore the permeate rate of a membrane, the first act is to remove any blockages stuck on the retentate side of the fibres. By far the best way to do this is to drive water in the opposite direction form that normally used, ie from permeate to retentate side. However, due to the construction of the fibres, they have far less strength against collapse than explosion, and hence such a back flushing procedure has to be very carefully controlled. One way to get around this problem is to close off the permeate side when full of liquid, while establishing a pressure difference along the length of the fibres themselves by increasing the retentate flow rate.

When such a situation arises water is forced out of the fibres at the high pressure end, the inlet, and drawn back into the low pressure end of the fibres, at the outlet. Thus a flow of water in the opposite direction is achieved over half the membrane. In order to complete this cleaning process the direction of the retentate flow is then reversed, so that the high pressure end of the fibre becomes the low pressure end and vice versa. These two modes are shown in figures C.2 and C.3. This slightly elaborate system is required as the pumps on both the one foot and four foot membrane levels operate in a single direction only, and hence the flow cannot be reversed easily.

- Dead end mode. Finally, we come to dead end mode, in which the membrane passes all the water from the retentate side to the permeate side. The diagram in figure C.4 shows the retentate side supply coming from the feed tank, and the permeate stream going to the permeate tank. The passage of the 100 litres in the single H1 tests was achieved using this process, and is to be used in the final detector system. in the lab scale system, there were problems with air bubbles in the feed stream collecting in the fibres themselves, as they cannot pass through the pores due to the hydrophilic nature of the fibre material itself. It was therefore necessary to occasionally flush the fibres by opening valve G. Such problems are not expected in the full scale system as the heavy water will be thoroughly degassed during processing.

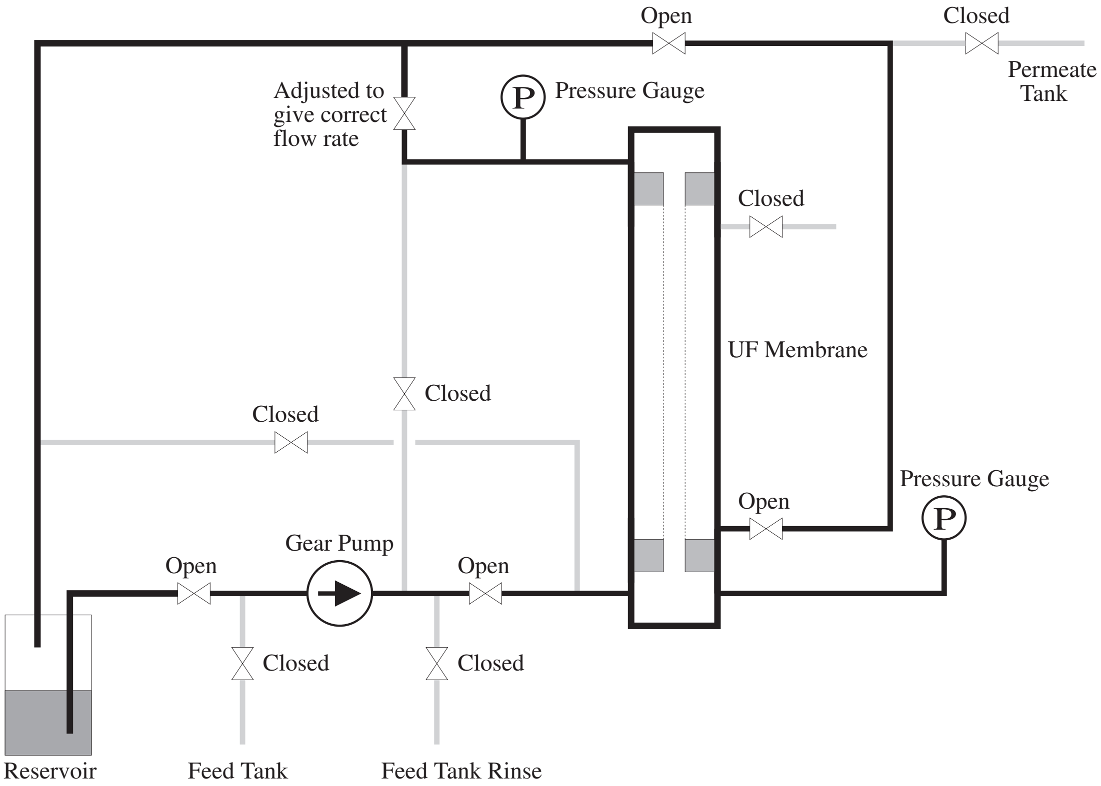

Figure C.1: The flow diagram for recirculation mode. The darker lines show the water path, with the appropriate valves opened or closed as necessary.

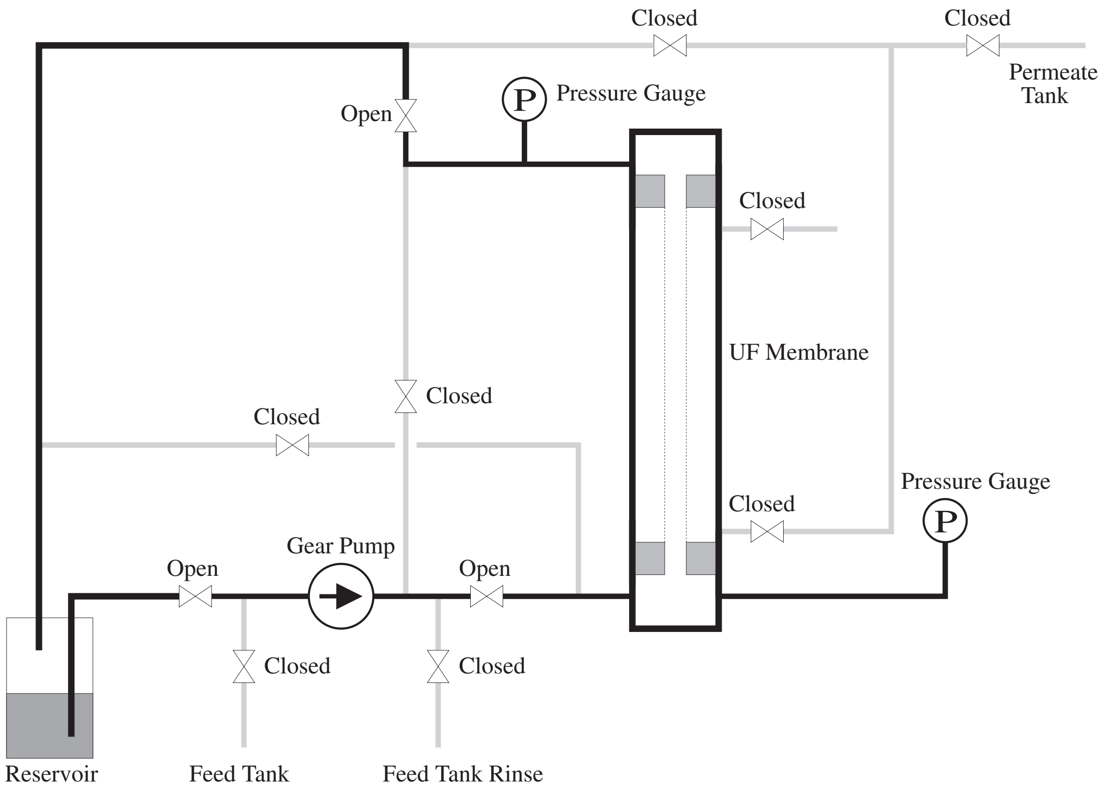

Figure C.2: The flow diagram for fast forward flush up mode. Again, the darker lines show the water path. This mode allows the cleaning of the membrane by creating a pressure difference along the length of the fibres.

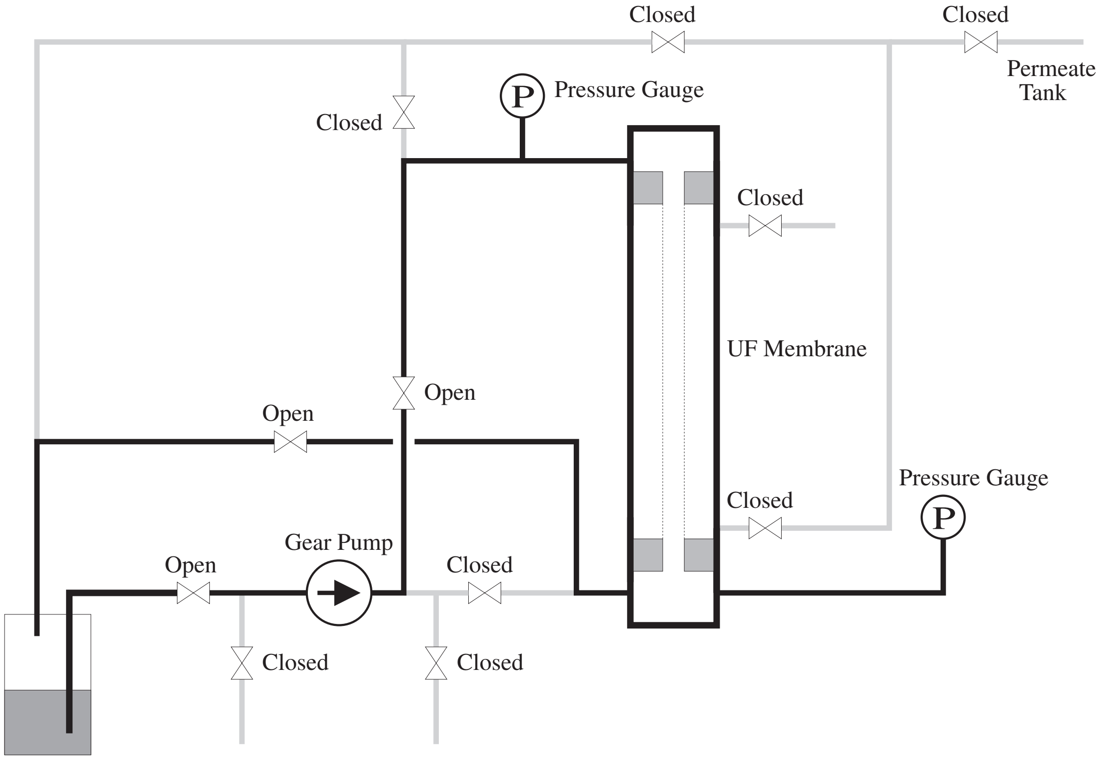

Figure C.3: The flow diagram for fast forward flush down mode. This mode completes the cleaning process by reversing the direction of the flow on the retentate side, and hence establishing the pressure difference in the opposite direction to the previous figure.

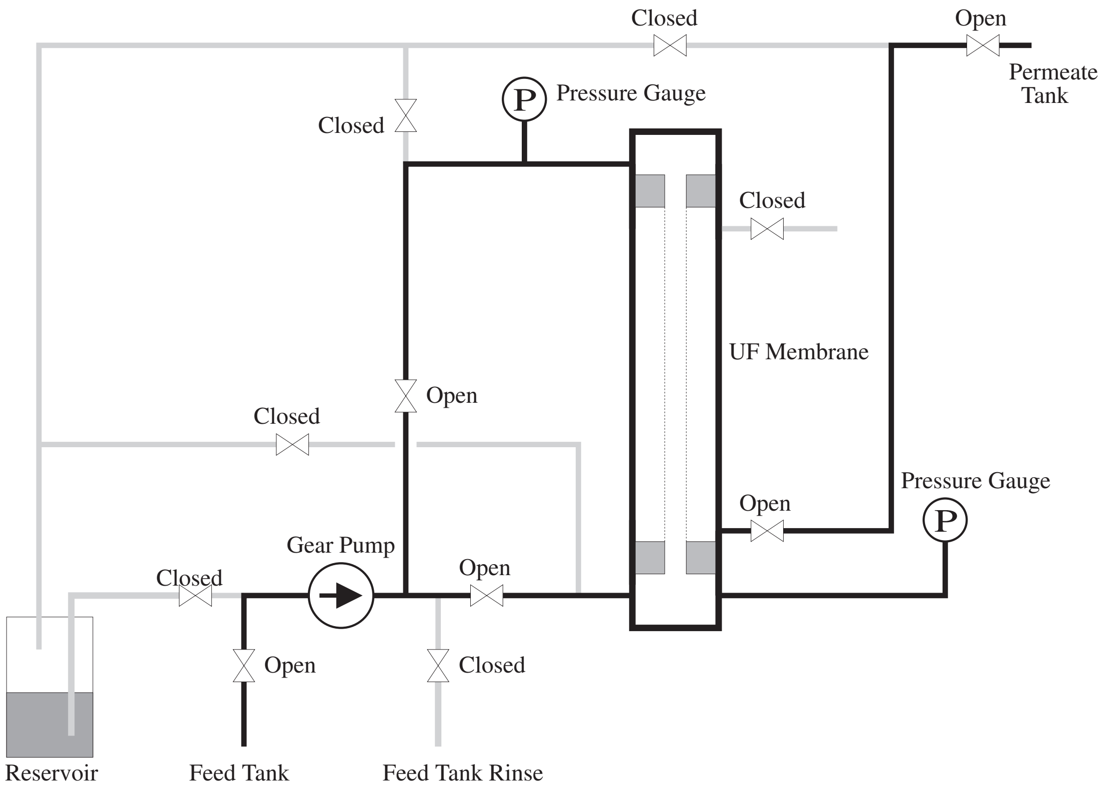

Figure C.4: The flow diagram for dead end mode. In this configuration all the water presented to the membrane is pushed through the pores. In the diagram, the feed and permeate lines are connected to the appropriate tanks used in the 100 litre scale tests.

Appendix D

Radon breakthrough

In this section we calculate the theoretical DF's for an extraction. If we consider the number of atoms, N^F, passing a point just before the extracting section, where the superscript F denotes the feed line, the flow rate is:

$$\frac{dN_i^F}{dt} = \left.\frac{\partial N_i^F}{\partial V}\right|_t \frac{dV}{dt} + \left.\frac{\partial N_i^F}{\partial t}\right|_V$$

where i=Th, Ra, Pb. The second term in this equation is set to zero as we only consider small elements of the feed during the run. Thus,

$$\frac{dN_i^F}{dt} = \left.\frac{\partial N_i^F}{\partial V}\right|_t \frac{dV}{dt} = N_i'$$

This is the amount hitting the extracting section per unit time, and if a fraction α_i sticks, then in the differential equations for the activity extracted, denoted by the superscript M:

$$\frac{dN_{Th}^M}{dt} = \alpha_{Th} N_{Th}' - \lambda_{Th} N_{Th}^M$$

$$\frac{dN_{Ra}^M}{dt} = \alpha_{Ra} N_{Ra}' - \lambda_{Ra} N_{Ra}^M + \delta\lambda_{Th} N_{Th}^M$$

$$\frac{dN_{Pb}^M}{dt} = \alpha_{Pb} N_{Pb}' - \lambda_{Pb} N_{Pb}^M + \epsilon\lambda_{Ra} N_{Ra}^M$$

The parameters δ and ϵ are the respective probabilities that when a captured atom decays, its daughter is re-captured on the retentate side. Any daughters which are not captured are assumed to reach the permeate side, as shown in the equations below.

These are, as yet, parameters about which nothing is known, but they are included for completeness.

For the flow in the permeate stream we thus have:

$$\frac{dN_{Th}^P}{dt} = (1 - \alpha_{Th})N'_{Th}$$

$$\frac{dN_{Ra}^P}{dt} = (1 - \alpha_{Ra})N'_{Ra} + (1 - \delta)\lambda_{Th}N_{Th}^M$$

$$\frac{dN_{Pb}^P}{dt} = (1 - \alpha_{Pb})N'_{Pb} + (1 - \epsilon)\lambda_{Ra}N_{Ra}^M$$

Solving the equations for the captured activity and assuming that the extracting section is clean of activity at t=0 gives:

$$\lambda_{Th}N_{Th}^M = \alpha_{Th}N'_{Th}(1 - e^{-\lambda_{Th}t})$$

So,

$$\frac{dN_{Ra}^M}{dt} = \alpha_{Ra}N'_{Ra} - \lambda_{Ra}N_{Ra}^M + \delta\alpha_{Th}N'_{Th}(1 - e^{-\lambda_{Th}t})$$

The exact solution is thus

$$\lambda_{Ra}N_{Ra}^M = (\alpha_{Ra}N'_{Ra} + \delta\alpha_{Th}N'_{Th})(1 - e^{-\lambda_{Ra}t}) + \frac{\lambda_{Ra}\delta\alpha_{Th}N'_{Th}}{\lambda_{Ra} - \lambda_{Th}}(e^{-\lambda_{Ra}t} - e^{-\lambda_{Th}t})$$

If we now consider the permeate side, we have:

$$\frac{dN_{Th}^P}{dt} = (1 - \alpha_{Th})N'_{Th}$$

$$\frac{dN_{Ra}^P}{dt} = (1 - \alpha_{Ra})N'_{Ra} + (1 - \delta)\alpha_{Th}N'_{Th}(1 - e^{-\lambda_{Th}t})$$

$$\lambda_{Pb}\frac{dN_{Pb}^P}{dt} = (Z_1 + Z_2 + Z_3)$$

where:

$$Z_1 = (1 - \alpha_{Pb})\lambda_{Pb}N'_{Pb}$$

$$Z_2 = (1 - \epsilon)\lambda_{Pb}(\alpha_{Ra}N'_{Ra} + \delta\alpha_{Th}N'_{Th})(1 - e^{-\lambda_{Ra}t})$$

$$Z_3 = (1 - \epsilon)\frac{\lambda_{Ra}\delta\alpha_{Th}N'_{Th}\lambda_{Pb}}{\lambda_{Ra} - \lambda_{Th}}(e^{-\lambda_{Ra}t} - e^{-\lambda_{Th}t})$$

Again, the decays of the atoms on the permeate side are not considered as we are only calculating the activity of small elements of the permeate stream, and not the integrated permeate activity. Now, if we consider the number of atoms, ΔN, which pass a point in the permeate stream between t and t+Δt, then :

$$\lambda_{Pb}\Delta N_{Pb}^{P} = \Delta t(Z_1 + Z_2 + Z_3)$$

Now, if we also define:

$$A_i = \lambda_i \left.\frac{\partial N_i^F}{\partial V}\right|_t \frac{dV}{dt}$$

Then:

$$Z_1 = (1-\alpha_{Pb})A_{Pb}$$

$$Z_2 = (1-\epsilon)(\alpha_{Ra}\frac{\lambda_{Pb}}{\lambda_{Ra}}A_{Ra} + \delta\alpha_{Th}\frac{\lambda_{Pb}}{\lambda_{Th}}A_{Th})(1-e^{-\lambda_{Ra}t})$$

$$Z_3 = (1-\epsilon)\frac{\lambda_{Pb}}{\lambda_{Th}}A_{Th}\frac{\lambda_{Ra}\delta\alpha_{Th}}{\lambda_{Ra}-\lambda_{Th}}(e^{-\lambda_{Ra}t} - e^{-\lambda_{Th}t})$$

Now, we can simplify this expression by saying that for all times over which an extraction takes place:

$$e^{-\lambda_{Th}t} \approx 1 - \lambda_{Th}t + \frac{\lambda_{Th}t^2}{2}$$

Thus:

$$Z_3 = A_{Th}(1-\epsilon)\frac{\lambda_{Pb}}{\lambda_{Th}}\frac{\lambda_{Ra}\delta\alpha_{Th}}{\lambda_{Ra}-\lambda_{Th}}(e^{-\lambda_{Ra}t} - 1 + \lambda_{Th}t + \frac{\lambda_{Th}t^2}{2})$$

If we also assume that the feed source is in equilibrium so that:

$$A_{Th} = A_{Ra} = A_{Pb} = A_0$$

Then we have:

$$\lambda_{Pb}\frac{\Delta N_{Pb}}{A_0\Delta V} = (1-\alpha_{Pb}) + (1-\epsilon)\frac{\lambda_{Pb}}{\lambda_{Ra}}\left[\alpha_{Ra}(1-e^{-\lambda_{Ra}t}) + \delta\alpha_{Th}(e^{-\lambda_{Ra}t} - (1-\lambda_{Ra}t))\right]$$

where we have said:

$$\frac{\lambda_{Ra}t\lambda_{Th}t}{2} \approx 0$$

$$\lambda_{Ra} - \lambda_{Th} \approx \lambda_{Ra}$$

If we finally consider the regime where:

$$e^{-\lambda_{Ra} t} \approx 1 - \lambda_{Ra} t$$

Then:

$$\lambda_{Pb} \frac{\Delta N_{Pb}}{A_0 \Delta V} = (1 - \alpha_{Pb}) + (1 - \epsilon) \alpha_{Ra} \lambda_{Pb} t$$

This is then approximately the inverse of the DF.

Bibliography

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

[2] J.N. Bahcall and R.K. Ulrich, Rev. Mod. Phys. **60** 297 (1988).

[3] J.N. Bahcall and M.H. Pinsonneault, Rev. Mod. Phys. **64** 885 (1992).

[4] J.N. Bahcall and M.H. Pinsonneault, *Solar Models with Helium and Heavy Element Diffusion* , To be published

[5] S. Turck-Chieze, S. Cahen, M. Casse and C.Doom, Astrophys. J. **335** 415 (1988).

[6] S. Turck-Chieze and I. Lopes, Ap. J. **408** 347 (1993).

[7] I.J. Sackmann, A.I. Boothroyd, A.W. Fowler, Astrophys. J., **386** 727 (1990).

[8] F.J. Rogers and C.A. Iglesias, Science **263** 50 (1994).

[9] G. Berthomieu et al., Astr. Astrophys. **268** 775 (1993).

[10] V. Castellani et al., Phys. Rev. D, **50**, number 8 4749 (1994).

[11] J.N. Bahcall, Phys. Rev. D, **44**, number 6 1644 (1991).

[12] B.T. Cleveland et al., Neutrino 94 Conference Proceedings Nuc. Phys. B, **38** 47 (1994).

[13] K.S. Hirata et al., Phys. Rev. D, **44**, number 8 2241 (1991).

[14] Y. Suzuki, Neutrino 94 Conference Proceedings Nuc. Phys. B, **38** 54 (1994).

[15] J.N. Abdurashitov *et al*, Phys. Lett. B **B328** 234 (1994).

[16] J.N. Abdurashitov et al., Neutrino 94 Conference Proceedings Nuc. Phys. B, **38** 60 (1994).

[17] P. Anselmann et al., Phys. Lett. B, **327** 377 (1994).

[18] GALLEX collaboration, Neutrino 94 Conference Proceedings Nuc. Phys. B, **38** 68 (1994).

[19] To be published

[20] To be published

[21] M. Gai, Neutrino 94 Conference Proceedings Nuc. Phys. B, **38** 77 (1994).

[22] S. Parke, Phys. Rev. Lett. **74**, number 6 839 (1995).

[23] B. Pontecorvo, Sov. JETP **26** 984 (1968).

[24] D.H. Perkins, P.F. Harrison and W.G. Scott, *Threefold Maximal Lepton Mixing and the Solar and Atmospheric Neutrino Deficits.*, RAL-94-125

[25] L. Wolfenstein, Phys. Rev. D. **20** 2634 (1979).

[26] S.P Mikheyev and A. Yu. Smirnov, Sov. J. Nucl. Phys., **42** 1441 (1985); Nuovo Cimento **9C** 17 (1986).

[27] E. Bellotti, Neutrino 94 Conference Proceedings Nuc. Phys. B, **38** 90 (1994).

[28] The SNO collaboration *Sudbury Neutrino Observatory Proposal* SNO reference number SNO-87-12.

[29] The SNO collaboration *Annexes to Proposal* SNO reference number SNO-87-12.

[30] M.E. Moorhead, D.Phil. Thesis

[31] K. Kubodera, S. Nozawa and J.N. Bahcall, Phys. Rev. D **38** 1030 (1988)

[32] R.G.H. Robertson *Neutral-Current Detector Progress Report* SNO-STR-93-046

[33] S.J. Brice *An Overview of the Feedforward Neural Network Technique and its Application to SNO Event Classification* SNO-STR-95-037

[34] M.D. Lay, D.L. Wark and N. West, *It's Alive - the birth of SNOMAN* SNO-STR-91-077

[35] M.D. Lay, D. Phil. Thesis

[36] K. Hirata et al., Phys. Rev. Lett. **58** 1490 (1987). R.M. Bionta et al., Phys. Rev. Lett. **58** 1494 (1987)

[37] A. Burrows, D. Klein and R. Gandhi, Phys. Rev. D **45** 3361 (1992). See also A. Burrows, Ann. Rev. Nucl. Part. Sci. **40** 181 (1990).

[38] Particle Data Book, Phys. Rev. D **50** 1173 (1994)

[39] C.M. Lederer and V.S. Shirley *Table of Isotopes*, 7th edition, Wiley Interscience (1978).

[40] W.Locke, Private communication.

[41] J.F. Duncan and R.G. Richards, New Zealand Journal of Science, 1976, Vol.19, p.185

[42] R.G. Gutman *Membrane Filtration*, Adam Hilger (1987). This is a mainly technical reference text.

[43] E.W. Hooper, B.A. Phillips, S.P. Dagnall, N.P. Monckton UK AEA Harwell, AERE-R 11088 p27, 1984, and references therein.

[44] K. Kawazoe and M. Fujimoto, International Chemical Engineering, 1984, Vol.24, No.3, p.519

[45] C.J. Bland and A. Martin Sanchez, NIM A286, 1990, p.375

[46] R.B. Galloway, Meas. Sci. Technol., 1, 1990, p.725

[47] J.B. Birks *The Thoery and Practice of Scintillation Counting*, Pergamon Press 1964

[48] J.W. McKlveen and W.J. McDowell, NIM A223, 1984, p.372

[49] C. Dodson *Introduction to Alpha/Beta Discrimination*, Beckman Bulletin 7894B.

[50] B.T. Cleveland, NIM **214** 451, 1983

[51] K.H. Howard, *Ion Exchange Properties of HTiO*, Part II Chemistry Thesis, 1994

[52] S. Lilley, M.E. Moorhead and N.W. Tanner, *Seeded tangential filtration: Progress and Prognosis*, SNO-STR-92-089

[53] W. Locke and H. Heron, *Elution of HTiO*, SNO internal report.

[54] M.E. Moorhead, Private communication.

[55] M.E. Moorhead et al., *Design Description for the SUF Plant in SNO.*, SNO-STR-95-003

[56] R.J. Boardman, D. Phil. Thesis

[57] R.S. Lakes and S.K. Poultney, Appl. Opt., **10**, Number 4, 797(1971).

[58] A. Young and R.E. Schild, Appl. Opt., **10**, Number 7, 1668(1971).

[59] A.G. Wright *Sources of noise in photomultipliers.*, Thorn EMI Report.

[60] A. Venema *Thermionic emission* in Handbook of Vacuum Physics, Ed. A.H. Beck, 2, Physical Electronics, Pergammon Press.

[61] P.B. Coates, J. Phys. D, **6**, 1159, 1973