INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

Bell & Howell Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600

THE PALO VERDE NEUTRINO OSCILLATION EXPERIMENT

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

By
Lester H. Miller
October 2000

UMI Number: 3000068

UMI Microform 3000068

Bell & Howell Information and Learning Company
300 North Zeeb Road
P.O. Box 1346
Ann Arbor, MI 48106-1346

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Giorgio Gratta
(Principal Adviser)

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Patricia Burchat

I certify that I have read this dissertation and that in my opinion it is fully adequate, in scope and quality, as a dissertation for the degree of Doctor of Philosophy.

Scott Thomas

Approved for the University Committee on Graduate Studies:

Thomas Wasow

Abstract

The $\bar{\nu}_e$ flux and spectrum have been measured at a distance of about 800 m from the reactors of the Palo Verde Nuclear Generating Station using a segmented Gd-loaded liquid scintillator detector. Correlated positron–neutron events from the reaction $\bar{\nu}_e p \rightarrow e^+ n$ were recorded for a period of 353 d including 110 d with one of the three reactors off for refueling. Backgrounds were accounted for by making use of the reactor-on and reactor-off cycles, and also with a technique based on the differences between signal and background under reversal of the e^+-type and n-type portions of the events. Results from the experiment show no evidence for neutrino oscillation. $\bar{\nu}_e \rightarrow \bar{\nu}_x$ oscillation was excluded at 90% CL for $\Delta m^2 > 1.28 \times 10^{-3}$ eV^2 for full mixing, and $\sin^2 2\theta > 0.22$ for large Δm^2. These results support the conclusion that the observed atmospheric neutrino oscillation predominantly does not involve ν_e.

Acknowledgments

The work presented here resulted from a collaborative effort among physicists from four universities, including but not limited to Jerome Busenitz, Joachim Wolf, Andreas Piepke, Kevin Mckinney and Janos Kornis of the University of Alabama; Felix Boehm, Petr Vogel, Herb Henrikson, K.B. Lee, Vladimir Novikov and Brian Cook of the California Institute of Technology; Barry Ritchie and David Lawrence of Arizona State University; and YiFang Wang, Giorgio Gratta, Dillon Tracy and myself of Stanford University. I am especially indebted to the help and guidance given over the years by Giorgio Gratta and YiFang Wang of our group at Stanford. Part of my work on the project was supported by a generous grant from the ARCS Foundation. Finally the support over the years from friends and family has been invaluable, especially during the trials of learning to cope with the vagaries of desert life.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The recent indications of the observation of the tau neutrino announced by the DONUT (E872)[1] experiment marked a milestone for particle physics, verifying directly the existence of the last fermion predicted by the standard model. Since the invocation of neutrinos by Pauli in 1930 to explain the continuous energy spectra observed in nuclear beta decay, experimental efforts to measure neutrino properties have been underway. Positive experimental results until now, however, have been limited basically to ascertaining the existence and interactions of the neutrino. Our knowledge of most properties of the neutrinos remains scarce.

One natural question to address is whether the neutrino has mass. Most current experimental efforts in neutrino physics try to address this question either directly or indirectly. The concept of neutrino mass is complicated by the fact that the neutrino which propagates ("mass eigenstate") is not necessarily the same state as the neutrino which interacts via the weak force ("flavor eigenstate"). The two sets of eigenstates are connected by a mixing matrix

$$|\nu_\ell\rangle = \sum_{i=1}^{N} U_{\ell i}|\nu_i\rangle,$$

where the weak eigenstates $\nu_\ell, \ell = e, \mu, ..N$ are linear combinations of mass eigenstates $\nu_{1,2..N}$ related by a unitary matrix U in direct analogy of the CKM matrix describing mixing in the quark sector. For the case of three family mixing, for example, the

matrix U has four degrees of freedom, three mixing angles and one phase.

Searches for neutrino mass can be performed directly by looking at the kinematic properties of decays involving neutrinos, or indirectly via phenomena requiring mass, such as oscillation due to weak and mass eigenstate mixing.

Direct searches for neutrino mass look for kinematic effects in decays whose final state involves a neutrino. What one really measures as a mass is the effective neutrino mass

$$\langle m_{\nu_\ell} \rangle = \sum m_i U_{\ell i}^2$$

which may in general involve cancellations among the phases of the elements of U. Under the assumption that each weak eigenstate is dominated by a single mass eigenstate; however, one can look for missing energy in the final state. Note that in light of results from oscillation experiments favoring maximal mixing of some neutrino flavors, described below, this may not be a valid assumption. Relaxing this assumption, one can still use the energy spectrum in the final state to look for changes in cross section as the remaining energy available crosses a production threshold for one of the massive neutrino eigenstates, thus changing the possibilities available to the final state.

The most sensitive direct probes of neutrino mass look for modification of the endpoint energy spectrum of tritium beta decays [2, 3] as a probe for electron neutrino mass. Currently the combined upper limit is estimated to be $m_{\nu_e} < 3$ eV at 90% confidence level (CL). However, these experiments tend to see anomalous results giving best fits which prefer negative neutrino mass, and are difficult to interpret. The best limit for the mass of the muon neutrino comes from the decay $\pi \longrightarrow \mu\nu_\mu$ at the PSI experiment [4], $m_{\nu_\mu} < 0.19$ MeV (90% CL). There are several upper limits on the mass of ν_τ, the best being a measurement using the $\tau \longrightarrow 5\pi^\pm\nu_\tau$ decay by the combined LEP experiments [5], $m_{\nu_\tau} < 18$ MeV (95% CL).

Some experiments search for neutrinoless double beta decay, $(A, Z) \longrightarrow (A, Z + 2) + 2e^-$, as another sign of neutrino mass. This transition, with no neutrinos in the final state, is only possible if the electron neutrino is a massive Majorana particle, since by serving as the intermediary particle in the double decay the neutrino quantum

state must have both helicities and be its own antiparticle. Currently the Heidelberg-Moscow germanium experiment [6] gives the best limit on the Majorana neutrino mass as $\langle m_{\nu_e} \rangle < 0.5 - 1.5$ eV, with the uncertainty stemming from the calculation of the nuclear matrix elements involved.

It is worth noting that cosmology also gives bounds on neutrino mass. Models of the big bang predict the existence of a sea of primordial neutrinos, analogous to the cosmic microwave background. These relic neutrinos, if massive, play a role in the formation of structure of the universe and currently are the only known candidate for what is commonly referred to as hot dark matter. While there are large uncertainties in certain fundamental inputs to these models such as the Hubble constant and recent indications of a non-zero cosmological constant, the sum of all species of neutrino mass can still be constrained to be less than approximately 10 eV [7].

A supernova explosion in our galaxy, Supernova 1987A, has also been used to place an upper limit of $M_{\nu_e} < 11$ eV [8]. This limit is derived by assuming that neutrinos of different energies, if massive, will move with different velocities. The 12 events observed over a period of several seconds which are attributed to neutrinos originating from the supernova ranged in energies from 2 to 20 MeV, with no bias for higher energy neutrinos arriving first.

The indirect phenomenon of neutrino oscillation is perhaps the most accessible probe of neutrino mass, and has attracted the most attention experimentally lately. This thesis describes one of many experiments in this area.

1.1 ν oscillation

Neutrino oscillation has a long history beginning with a proposal by Bruno Pontecorvo in 1957 [9] of $\nu \longleftrightarrow \bar{\nu}$ mixing, in analogy with the kaon system oscillation. The idea of neutrino oscillation due to flavor mixing was first posited in 1962 [10] in light of discovery of the muon neutrino, and later expanded to its modern incarnation in analogy with the CKM matrix in 1978[11].

As noted above, one can describe the N ν_ℓ flavor eigenstates as linear superpositions of mass eigenstates ν_i. Since each mass eigenstate will evolve in time according

to the Schroedinger equation, the superposition flavor eigenstate will time-evolve according to

$$|\nu_\ell(t)\rangle = \sum_i e^{-iE_i t} U_{\ell i} |\nu_i\rangle,$$

where E_i is the energy of the respective mass eigenstate components.

The flavor composition of the evolving particle will then oscillate depending on the energy of the neutrino as

$$\langle \nu_{\ell'} | \nu_\ell(t) \rangle = \sum_{i,j} \langle \nu_j | U^\dagger_{\ell' j} e^{-iE_i t} U_{\ell i} | \nu_i \rangle \quad (1.1)$$

$$= \sum_m e^{-iE_i t} U_{\ell i} U^*_{\ell' i}, \quad (1.2)$$

which means the probability of finding a different flavor from the original flavor also follows an oscillatory pattern:

$$P_{\nu_\ell \nu_{\ell'}}(t) = |\langle \nu_{\ell'} | \nu_\ell(t) \rangle|^2 \quad (1.3)$$

$$= \sum_{i,j} U_{\ell i} U^*_{\ell' i} U^*_{\ell j} U_{\ell' j} e^{i(E_i - E_j)t}. \quad (1.4)$$

Under the assumption that the neutrino is ultrarelativistic ($p_\nu \gg m_i$) for all m_i giving a common neutrino momentum p_ν, we can rewrite the energy in the exponent in terms of momentum and distance:

$$E_i = (|\mathbf{p}|^2 + m_i^2)^{1/2} \simeq |\mathbf{p}| + \frac{m_i^2}{2|\mathbf{p}|},$$

where $|\mathbf{p}| \equiv p_\nu$ is the neutrino momentum, and insert this into the oscillation probability (Equation 1.4) to get

$$P_{\nu_\ell \nu_{\ell'}}(x) = \sum_{i,j} U_{\ell i} U^*_{\ell' i} U^*_{\ell j} U_{\ell' j} e^{\frac{m_i^2 - m_j^2}{2E_\nu} L}$$

where t has been replaced by L under the assumption $p_\nu \simeq E_\nu$.

From this result one sees that the probability of oscillation, or the probability of

detecting flavor ℓ' in an initial state of flavor ℓ follows a sum of sinusoidal patterns with characteristic lengths of

$$L_{m,m'} \equiv \frac{4\pi E}{m^2 - m'^2}.$$

There are several subtleties in the above description, namely concerns over the assumptions of the ultrarelativistic limit and whether the state will decohere during propagation (see e.g. [12]), but in all practical applications of neutrino oscillations these issues do not arise.

Neutrino oscillation experiments often analyze their results in a simplified two family mixing scenario. In this case the mixing matrix U has only one free parameter, a mixing angle θ and there is only one mass difference, $\Delta m^2 \equiv m_1^2 - m_2^2$. The unitary mixing matrix U written in terms of θ is

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

and the probability of oscillation is then

$$P_{\nu_\ell \nu_{\ell'}} = \sin^2 2\theta \sin^2 \left(\frac{\Delta m^2 L}{4|\mathbf{p}|} \right) \tag{1.5}$$

$$= \sin^2 2\theta \sin^2 \left(\frac{1.27 \times \Delta m^2\,[eV^2] \times L[m]}{E[MeV]} \right) \tag{1.6}$$

where the latter equation is given in terms of practical units.

Recently a wealth of data on oscillation has been gathered, allowing the analysis of more complex flavor mixing models which will be discussed briefly below after the review of neutrino oscillation experiments.

1.2 Review of ν oscillation experiments

Oscillation experiments can be broadly classified according to the source of neutrinos used. Two sources are extraterrestrial, the sun and cosmic rays which produce solar

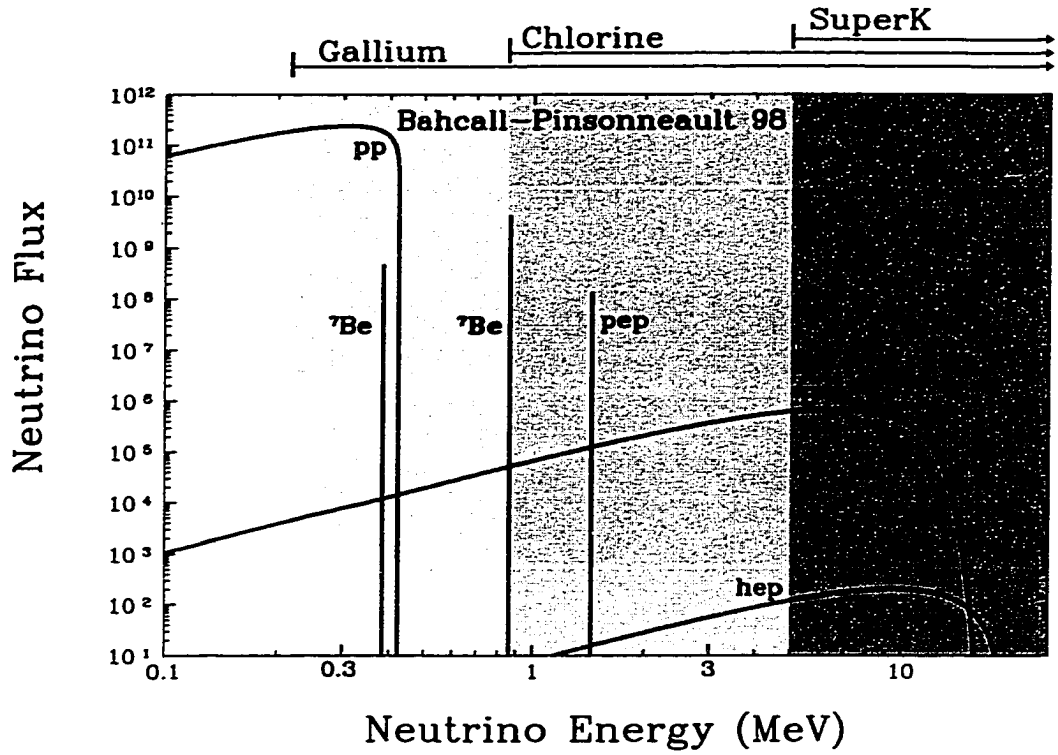

Solar neutrino energy spectrum

Figure 1.1: The different components of the spectrum of neutrinos predicted by the standard solar model (SSM) separated by parent reaction. The thresholds of the three current solar neutrino measurements are shown across the top of the figure. Figure taken from [14].

and "atmospheric" neutrinos respectively. In addition to these sources, neutrinos from two man-made sources are used, those from reactors and those from accelerator beams.

1.2.1 Solar neutrinos

Electron neutrinos originating from the sun, along with helioseismology, are the only known direct probes of the solar core, and offer useful tests of the Standard Solar Model (SSM) [13]. There have been three types of experiment completed thus far which study these neutrinos, all of which have reported fewer events than expected. The different components of the neutrino spectrum predicted by the SSM are shown in Figure 1.1, along with the thresholds of the three experimental methods.

The "solar neutrino problem" was first observed as a deficit in a radiochemical experiment at the Homestake mine in South Dakota [15]. This experiment used the reaction

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

with a threshold of 814 keV. Through periodic extractions of the ^{37}Ar, whose production rate is estimated by counting K-shell electron capture decays, they report observing $(33 \pm 3)\%$ of the expected signal.

This prompted other experiments, SAGE [16] and GALLEX [17], to pursue a similar method using the reaction

$$\nu_e + {}^{71}\mathrm{Ga} \longrightarrow e^- + {}^{71}\mathrm{Ge},$$

which has a lower threshold of 233 keV, allowing observation of "p-p" neutrinos which are products of the primary fusion cycle in the sun. Both experiments reported again a deficit of observed neutrinos. The water Čerenkov detectors Kamiokande [18] and its successor SuperKamiokande [19], with thresholds or around 5 MeV, measured the high energy end of the solar ν_e flux by looking at elastic scattering of ν_e off of electrons, and reported a deficit. These last experiments have the advantage of having directional sensitivity and real time counting, and were able to conclusively prove that the neutrinos being seen do indeed come from the direction of the sun.

Figure 1.2 shows a comparison of the neutrino composition the SSM predicts and the total neutrino flux observed for each of the three experimental thresholds. The discrepancy between SSM prediction and observation cannot be reconciled by tuning the parameters of the solar model such as solar core temperature, since these are well constrained by excellent agreement with helioseismology observations [13]. In addition, solar luminosity measurements place severe constraints on the flux of p-p neutrinos, and the relative abundance of ^{8}B to ^{7}Be neutrinos is difficult to change, the former being a daughter product of the reaction producing the latter. Neutrino oscillations, however, could reconcile the experimental results. In particular, two oscillationn scenarios have been found to be consistent with observation.

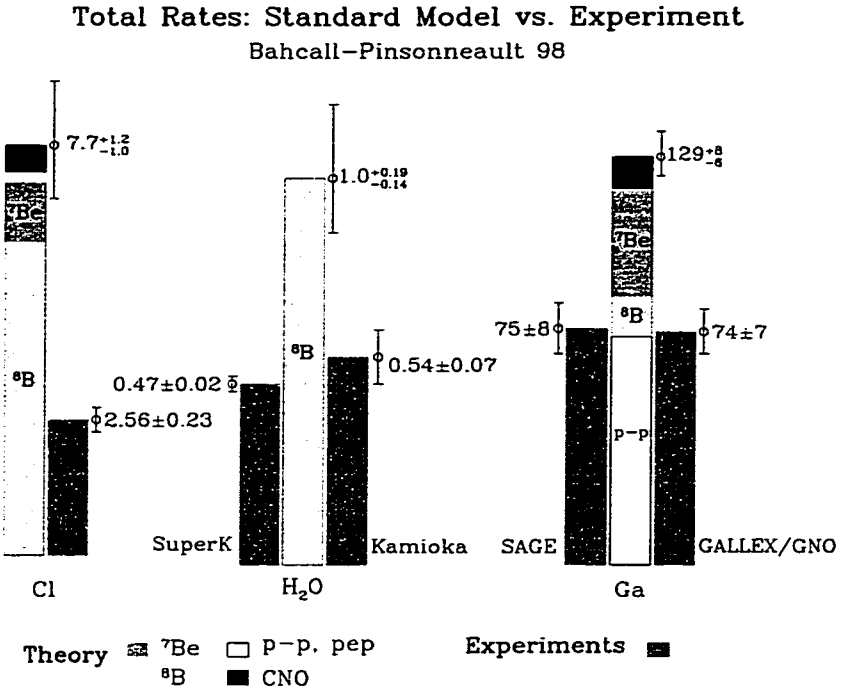

Figure 1.2: A comparison of the possible standard solar model contributions to the fluxes expected for each kind of solar experiment. The solar neutrino unit (SNU) is defined as one interaction per second per 10^{36} atoms, a convenient unit for solar neutrino experiments. Figure taken from [14].

There are two ways to reconcile the experimental results using neutrino oscillations, however. The first, neutrino oscillations in vacuum, also called the "just so" solution, requires almost maximal mixing and fine tuning the mass difference parameters to the orbital radius of the earth. The extremely small mass differences required for this solution are difficult to access in other experimental methods, but one can look for variations in the oscillations as the earth moves from perihelion to aphelion. Superkamiokande has looked for such seasonal variations in the observed flux and seen none [19], although the statistics of the measurement do not entirely rule out the vacuum oscillation solution, whose current allowed solution is shown in the left half of Figure 1.3.

The second oscillation solution to the solar neutrino problem uses a mechanism called the Mikheev-Smirnov-Wolfenstein (MSW) effect [20, 21]. Solar neutrinos, as they propagate from the core of the sun encounter matter which asymmetrically affects the ν_e component relative to the ν_μ or ν_τ component. Electron neutrinos can

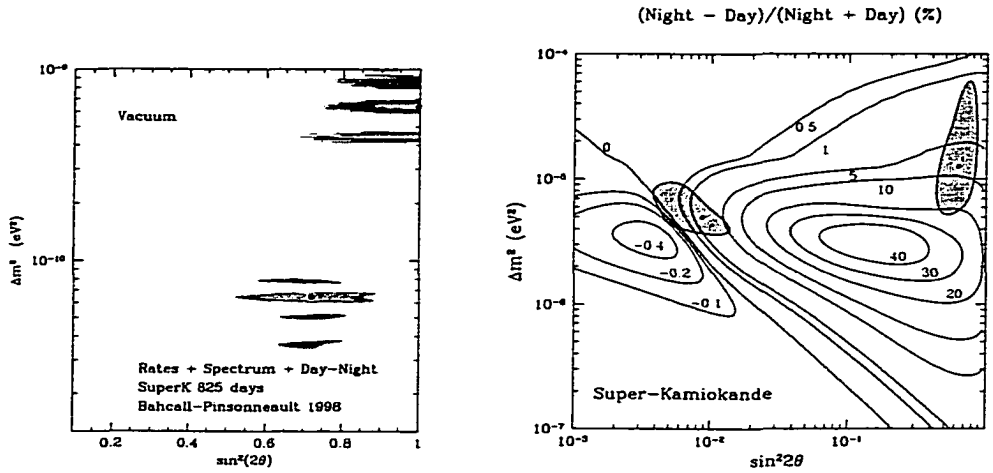

Figure 1.3: Left: The allowed regions for the vacuum oscillation solution for the solar neutrino problem. Less restrictive and more restrictive contours are the 99 and 90% confidence intervals respectively. Right: The allowed MSW solution at 95% CL. The lines overlain are the predicted day/night flux asymmetry due to MSW effects in the earth, a possible method of distinguishing among possible solutions. Figures taken from [14].

scatter elastically off of electrons in the sun via both charged and neutral current interactions, whereas the other neutrino flavors only interact via neutral current weak interactions. This extra interaction produces a different refractive index for the electron neutrino component and can produce a resonant conversion from electron to the other types of neutrinos, in effect amplifying what may be a small mixing in vacuum.

This solution requires mass differences on the order of 10^{-5} eV^2, and has different solutions depending on whether the other components being mixed with are active or sterile flavors of neutrino. The right side Figure 1.3 shows the allowed region for oscillation into an active species of neutrino taking into account current results on solar neutrino flux.

The two distinct allowed solutions can be distinguished (a third solution at $\Delta m^2 = 10^{-7}$, disfavored at 95% level, is not shown in this plot). The small mixing angle solution will cause a noticeable distortion of the ^{8}B spectrum. The large mixing angle

solution, while not having much distortion in the spectrum, would see a variation of total flux depending on whether the solar neutrinos pass through the earth due to MSW effects, the "day-night" effect.

The spectral distortion measurement is hampered by uncertainties in calculating the highest energy "hep" solar neutrinos. The reaction generating these neutrinos,

$$^3\mathrm{He} + p \longrightarrow ^4\mathrm{He} + e^+ + \nu_e,$$

has a very low rate and is not well constrained by theory [22]. Superkamiokande sees a surplus of > 14 MeV events which could correspond either to an oscillation signature or a $\sim$6 times larger than expected "hep" flux, which is still considered to be within the theoretical uncertainty. Currently they see a small "day-night" effect, $\frac{\mathrm{Night-Day}}{\mathrm{Night+Day}}(\%) = 3.5 \pm 2.2(\mathrm{stat.}) \pm 1.3(\mathrm{sys.})$, which is not significant enough to distinguish among solutions.

Several experiments are coming on line or being planned to address the solar neutrino problem and to distinguish among the possible scenarios described above. A 100 ton liquid scintillator experiment at Gran Sasso in Italy, Borexino [23], will be able to have a high statistics measurement of the solar neutrino spectrum with a threshold low enough to see the ^{7}Be neutrino contribution. It can discern vacuum oscillations by looking for seasonal variations in the flux of the ^{7}Be lines, which, being monochromatic, would show very clear oscillation signatures. Borexino will also be able to look for "day-night" variations in ^{7}Be flux.

The SNO [24] experiment in Sudbury, Canada, a Čerenkov detector with a 1000 ton heavy water target, is currently taking data and has the unique ability to measure both charged current ($\nu_e + d \longrightarrow \mathrm{p+p+}e^-$) and neutral current ($\nu_x + d \longrightarrow \mathrm{p+n+}\nu_x$) interactions. It will therefore measure the ν_e and active neutrino components of the solar neutrino flux independently.

Several future experiments to observe solar neutrinos are being considered. Most aim for real-time detection of solar neutrinos at low thresholds and with good energy resolutions, using a variety of technologies including phonon detection in liquid helium, very large helium TPCs, and nuclear transitions from ν interactions. These

experiments should allow a measurement of the true spectrum of solar neutrinos to look for spectral distortions due to oscillations and unambiguously resolve the issue. The KamLAND [25] experiment will have sensitivity to part of the oscillation parameters of the solar neutrino problem using reactor neutrinos at a very long baseline (see Section 1.2.3).

1.2.2 Atmospheric neutrinos

A set of experiments begun in the early 1980's to look for nucleon decay, were the first to measure the atmospheric neutrino flux: the Soudan [26] 1000 ton iron tracking calorimeter, and two Čerenkov detectors, IMB [27] and Kamiokande [28]. They measured the charged current interactions of electron and muon neutrinos by looking for charged tracks from electrons or muons beginning within the detector volume with no apparent entering track. Each noticed an anomaly in the measured ratio of muon to electron type neutrino produced by cosmic rays. This ratio, to first order 2, is a result of the copious π and K production in primary interactions which in turn produce the atmospheric neutrinos. The π^+, for example, decays through

$$\begin{aligned} \pi^- \longrightarrow \mu^- + \bar{\nu_\mu} \\ \searrow e^- + \bar{\nu_e} + \nu_\mu \end{aligned} \tag{1.7}$$

which, like the final state of the π^- and the K, has the simple ν_μ/ν_e ratio mentioned above. While the total ν production rate is poorly known, this ratio can be predicted with reasonable accuracy ($< 5\%$) and is relatively stable up to GeV μ's (higher energy muons tend to reach the ground before decaying more often, increasing the ratio.)

To see if this anomaly is consistent with neutrino oscillation, Kamiokande and later Superkamiokande have looked at the zenith angle dependence of the ν_e and ν_μ flux [29]. They found that the observed rate of muon neutrinos is lower for upward going events, i.e. longer path lengths through the earth. A similar effect was not observed for the electron events. The zenith angle dependence, shown in Figure 1.4 for Superkamiokande, is thus far unable to distinguish what is causing the upward going deficit, but a solution can be given in terms of neutrino oscillation, either to

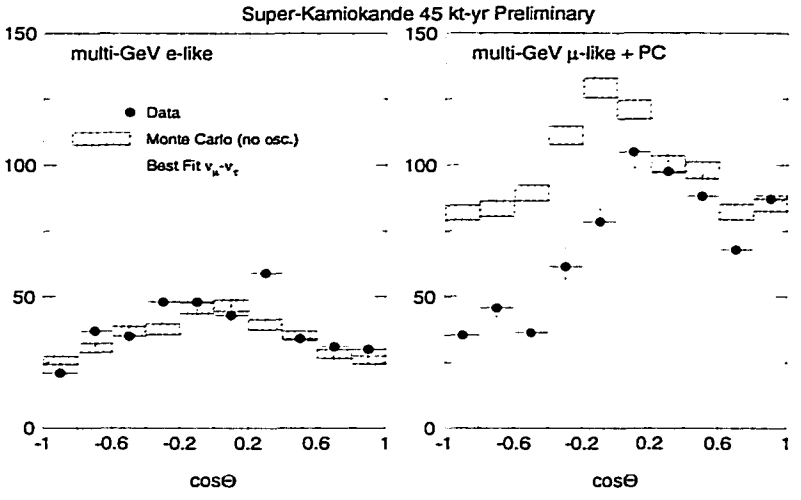

Figure 1.4: The observed zenith angle distribution of ν_e and ν_μ neutrinos. The clear deficit of upward going μ events can be interpreted as an oscillation signature. The data do not seem to have a corresponding distortion of the distribution of electrons, implying the oscillation is between ν_μ and either ν_τ or ν_s, a proposed fourth neutrino which is sterile to the weak interaction.

a τ neutrino or a sterile species. Another possibility not ruled out by the data but theoretically less appealing is ν_μ decay [30].

The atmospheric result implies a neutrino mass difference of $\Delta m^2 \sim \frac{E}{L} \sim \frac{\langle E_\mu \rangle}{R_\oplus} \sim \frac{10GeV}{10^4 km}$, and indeed the best fit for the result is at maximal mixing and $\Delta m^2 \sim 5 \times 10^{-3}$, which is accessible using man-made neutrino sources and has inspired several experiments including the present work at Palo Verde to investigate the phenomenon further. The next generation of experiment should be able to measure ν_μ reappearance as they propagate through more than one oscillation length, which could rule out muon decay and further refine the allowed oscillation parameter space.

1.2.3 Man made neutrino sources

In light of the above anomalies in extraterrestrial neutrino measurements, man made sources of neutrinos, being inherently better understood and easier to control, are an attractive option for investigating oscillation. These sources can either be accelerators or reactors, which offer complementary characteristics. Two figures of merit

are relevant when designing a neutrino experiment. The sensitivity in mixing angle θ is a function of the uncertainty in the final measurement of neutrino interactions seen. The sensitivity to Δm^2, on the other hand, is relatively immune to this uncertainty and is mainly a function of the baseline along which the neutrinos propagate relative to the energy of the neutrino, E/L. Accelerator experiments, with high energy neutrinos, tend to need long baselines but conversely can be energetic enough to create heavier leptons in charged current interactions such as τ's. Appearance experiments can explore very small mixing angles because uncertainties in the interaction rates are limited to the impurities in the original beam and understanding sources of backgrounds. Reactor experiments by comparison use low energy neutrinos but are restricted to looking for a deficit of neutrino interactions relative to a calculation of the expected rate for no oscillation, which limits mixing angle sensitivity.

Accelerator neutrinos

Accelerator experiments generally use high energy ν_μ beams to investigate either appearance of ν_e or ν_τ, or disappearance of ν_μ. Early experiments saw no oscillation signals but more recently evidence for oscillation has been reported by the LSND (LAMPF-1173) [31] experiment looking for $\nu_\mu \longrightarrow \nu_e$ oscillation. LSND uses the decay of stopped π^+ mesons (see Equation 1.7) to achieve a source of neutrinos almost devoid of $\bar{\nu_e}$. They searched for the appearance of $\bar{\nu_e}$ from oscillations of $\bar{\nu_\mu}$ via the inverse beta decay reaction $\bar{\nu}_e + p \longrightarrow e + n$, and seeing a surplus of these events over that expected due to beam $\bar{\nu}_e$ contaminations, deduced an oscillation signal whose allowed region in a two flavor oscillation scenario is shown in Figure 1.5. A separate analysis searching for $\nu_\mu \longrightarrow \nu_e$ using pion decay in flight as a ν_μ source gave a corroborative result with somewhat different systematics.

Much of the allowed region of this result has been excluded by other experiments looking for $\nu_e \longleftrightarrow \nu_\mu$ oscillations, including a very similar experiment at Rutherford Appleton Lab, KARMEN [36, 35]. None of the current generation of experiment is able, however, to fully exclude the result. An experiment at Fermilab, BooNE [37], is being built to confirm or deny this result conclusively and should begin taking data in 2002.

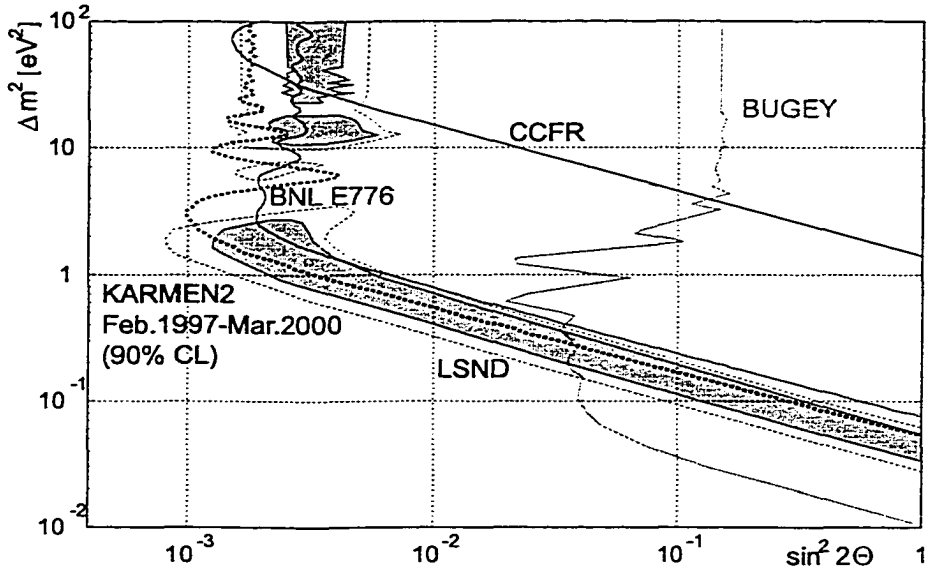

Figure 1.5: The allowed oscillation parameters from the LSND result (shaded region, 90 and 99% CL). Also shown are regions excluded at 90% CL by reactor experiment Bugey [32] and by other accelerator experiments CCFR [33], BNL E776 [34], and KARMEN [35] (dashed line).

Another set of accelerator-based experiments are aimed at looking for $\nu_\mu \longrightarrow \nu_\tau$ oscillation. NOMAD and CHORUS [38] were two early experiments at CERN which used iron tracking calorimeter and plastic emulsions respectively as targets for the τ appearances. Neither saw any evidence for oscillation.

The first true long baseline experiments are currently starting to take data or are being built. The K2K [39] experiment, a beam from the KEK lab in Japan to the Superkamiokande detector (L=235 km, $\langle E \rangle$ = 1.4 GeV), has begun taking data and is looking for signs of ν_μ disappearance. It should be sensitive to parts of the oscillation solution of the atmospheric neutrino anomaly and may have early indications of a deficit of events. The NuMi beam from Fermilab to the MINOS [36] iron tracking detector at the Soudan mine (L=734 km, $\langle E \rangle$ = 17.6 GeV) will have a better Δm^2 sensitivity. By counting both neutral current and charged current interactions, they should distinguish between ν_μ mixing with sterile versus τ neutrinos at the atmospheric neutrino solution parameters. Finally, the NGS beam from CERN to the ICANOE and OPERA [40, 41] experiments at Gran Sasso (L=732 km, $\langle E \rangle$ =

26.7 GeV) should offer complementary results to those of MINOS by looking for ν_τ appearance. The latter two detectors use liquid argon TPC/iron tracking and emulsion technologies respectively.

Reactor neutrinos

In a series of experiments beginning in 1953, Reines and Cowan first measured inverse beta decay induced by anti-neutrinos from a reactor [42]. Reactors have been used as intense sources of $\bar{\nu}_e$'s since then, being useful for having the best Δm^2 sensitivity among experiments with man-made sources as explained above. None of these experiments has seen any sign of oscillations. Using previous short ($\lesssim$10 m) baseline experiments such as Gösgen [43] and Bugey [32] as a high statistics source measurement, the longer baseline experiments Chooz [44] and the Palo Verde [45] experiment described below have measured the flux at $\sim$1 kilometer baselines. At $\langle E_\nu \rangle \sim 4$ MeV for reactor neutrinos, these experiments are sensitive to $\Delta m^2 \sim 10^{-3}$, and were designed to confirm the oscillation scenario seen in the atmospheric neutrino anomaly.

The next generation of reactor neutrino experiment, KamLAND [25], has baselines $L \sim$100 km and will begin taking data in 2001. It will be the first experiment using man made sources to explore the solar neutrino allowed oscillation parameter space, being sensitive to the large mixing MSW solution described above.

1.3 Summary

The allowed and excluded regions of most of the experiments described above are shown in Figure 1.6. There are three oscillation signals, two from the extraterrestrial sources, and the LSND accelerator result. None of the signals have an overlap.

Three neutrino oscillation signals with separated and independent allowed Δm^2 regions cannot be accounted for in a three family mixing model, as only two independent Δm^2 mass differences are free in that model. For this reason there have been proposals of four family models [7] involving a fourth, sterile, neutrino. The fourth neutrino species is usually assumed to be sterile to weak interactions in light of the LEP result limiting the number of interacting neutrino families with good accuracy

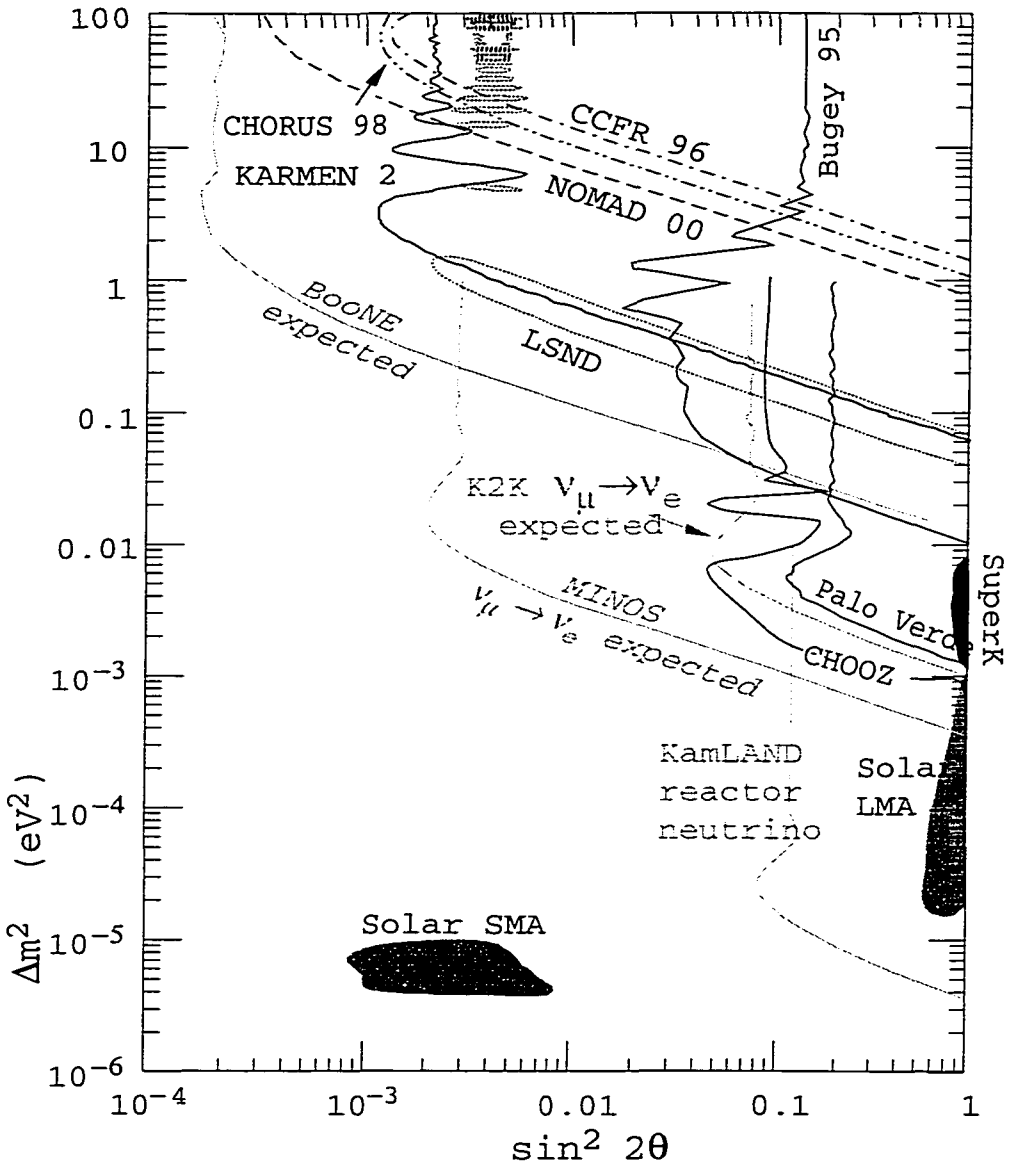

Figure 1.6: The allowed regions for the LSND, solar and atmospheric neutrino results (although the latter is not ν_e, ν_μ mixing). Also shown are the exclusion regions of other experiments described above looking for either ν_e or ν_μ disappearance or $\nu_e \longleftrightarrow \nu_\mu$ oscillation.

to three [46] based on the contributions of neutrinos to the observed Z boson total width. These models find a mixing matrix and set of mass eigenstates which satisfy the combined cosmological constraints and experimental results described above.

These scenarios tend to predict two pairs of almost degenerate, maximally mixed neutrino mass eigenstates. The heavy pair, ν_3 and ν_2, would correspond to the μ and τ mixing observed in the atmospheric neutrino anomaly, while the light pair ν_1 and ν_0 would correspond to the solar neutrino oscillation solution. The heavy and light pairs are mixed minimally at the large mass difference ($\Delta m^2 \sim 2$ eV) part of the allowed LSND region.

While this model does satisfy all the observed constraints, it seems not to respect the normal hierarchical nature of the three families found in other parts of the standard model, which would lead one to expect $m_{\nu_1} \ll m_{\nu_2} \ll m_{\nu_3}$ or $\Delta m^2_{1,2} \ll \Delta m^2_{2,3}$. If this were the case, one could describe the solar neutrino problem as ν_1, ν_2 mixing, $\Delta m^2_{1,2} \lesssim 10^{-4}$, and address the atmospheric neutrino oscillation solution with a simplified three parameter model to which the Palo Verde and Chooz results can add useful input. This model is explored further in the conclusion below.

Chapter 2

The Experiment

The Palo Verde experiment, a reactor neutrino experiment located at a nuclear power plant in Tonopah, Arizona, was designed to be sensitive to the neutrino oscillation solution found by Kamiokande based on observed zenith angle distribution of the atmospheric neutrino flux [28].

2.1 The detector

The Palo Verde Nuclear Generating Station, the largest nuclear power plant in the USA, consists of three identical pressurized water reactors with a total thermal power of 11.63 GW. The detector was located at a distance of 890 m from two of the reactors and 750 m from the third at a shallow underground site. Figure 2.1 shows an aerial view of the power plant site marked with the location of the laboratory. The 32 meter-water-equivalent overburden entirely eliminated any hadronic component of cosmic radiation while reducing the cosmic muon flux to 22 $m^{-2}s^{-1}$. In order to reduce the ambient γ-ray flux in the laboratory all materials in and surrounding the detector were selected for low activity. The laboratory walls were built with an aggregate of crushed dolomite, selected for its low content of natural radioisotopes. Concentrations of 170, 750, and 560 ppb for ^{40}K, ^{232}Th, and ^{238}U were measured in the concrete resulting in a tenfold reduction of γ-ray flux when compared with locally available aggregate (see Table 2.1). A low ^{222}Rn concentration of about 20 Bq/m^3 in the lab

Figure 2.1: The Palo Verde Nuclear Generating station seen from the air. The white "X" marks the location of the neutrino detector lab. The reactors are housed in the dome shaped buildings on the left.

air was maintained by forced ventilation. Temperature and humidity were controlled to ensure stable detector operation.

The segmented detector, shown in Figure 2.2, consisted of a 6×11 array of acrylic cells dimensioned at 900 cm×12.7 cm×25.4 cm and filled with a total of 11.34 tons of liquid scintillator. A 0.8 m long oil buffer at the ends of each cell shielded the central detector from radioactivity originating in the photomultiplier tubes (PMTs) and laboratory walls. The cells were made by cutting and bonding large 0.62 cm thick acrylic sheets. The total acrylic mass in the detector was 3.48 tons. Each cell was individually wrapped in 0.13 mm thick Cu foil to ensure light-tightness and was viewed by two 5-inch low activity PMTs [47], one at each end, housed in mu-metal boxes. The PMTs, developed by Electron Tubes Inc. in collaboration with us for the experiment, are detailed in Appendix B. The target cells were suspended on rollers held in place by thin sheet metal hangers. All structural materials were dimensioned as lightly as possible to minimize dead material between cells and radioactive contamination. Each cell could be individually removed from the mechanical structure for maintenance.

Material	Size	Isotope	Purity Level		Est. Background Rate [Hz]	
			ppb	Bq/kg	Singles	Triple coinc.
Laboratory Walls	392 tons	^{238}U	560	6.9	27	4.9
		^{232}Th	750	3.0	39	5.3
		^{40}K	170	53	14	1.9
Water in Buffer	100 tons	^{238}U	0.001	$1 \cdot 10^{-5}$	0.01	$2 \cdot 10^{-3}$
		^{232}Th	0.001	$4 \cdot 10^{-6}$	0.006	$1 \cdot 10^{-3}$
		^{40}K	$3 \cdot 10^{-5}$	$1 \cdot 10^{-5}$	0.002	$2 \cdot 10^{-4}$
Steel Buffer Tanks	1968 kg	^{238}U	0.1	0.0012	0.03	0.003
		^{232}Th	0.1	0.0004	0.19	0.03
		^{60}Co		0.001	1.04	0.17
Steel support columns	240 kg	^{238}U	0.6	0.0074	0.03	0.003
		^{232}Th	4	0.016	0.93	0.13
		^{60}Co		4.3	0.39	0.06
I-beams	257 kg	^{238}U	0.5	0.0062	0.02	0.003
		^{232}Th	2	0.0081	0.5	0.07
		^{60}Co		22	2.03	0.34
Steel Support Strips	475 kg	^{238}U	0.1	0.0012	0.023	0.006
		^{232}Th	0.1	0.0004	0.12	0.03
		^{60}Co		0.001	0.47	0.15
Stainless steel axles	12.4 kg	^{238}U	7	0.086	0.04	0.01
		^{232}Th	20	0.081	0.63	0.16
		^{60}Co		73	0.97	0.31
Photo–tubes	264 tubes	^{238}U	1.2 Bq/PMT		0.72	0.10
		^{232}Th	0.8 Bq/PMT		0.78	0.12
		^{40}K	5 Bq/PMT		0.37	0.06
Radon	6165 l	^{222}Rn	20 mBq/l		22	5.8
Acrylic Tanks	3535 kg	^{238}U	0.017	$2 \cdot 10^{-4}$	0.19	0.055
		^{232}Th	0.021	$9 \cdot 10^{-5}$	0.087	0.029
		^{40}K	$3 \cdot 10^{-4}$	$1 \cdot 10^{-4}$	0.014	0.003
TOTAL Background Rate:					112	19.7

Table 2.1: The measured activities of the detector and lab components. The natural radioactivity of all materials used in the detector and lab construction were measured to ensure a low random background rate for the experiment. Also listed are the estimated contributions to the trigger (described below) in the form of single hits and triply coincident hits in any 3×3 subset of the segmented target array. (The measured triple rate was 50 Hz, due to our using a looser definition of triple coincidence in the trigger). The two largest contributions of γ backgrounds to the detector are the lab walls and Radon.

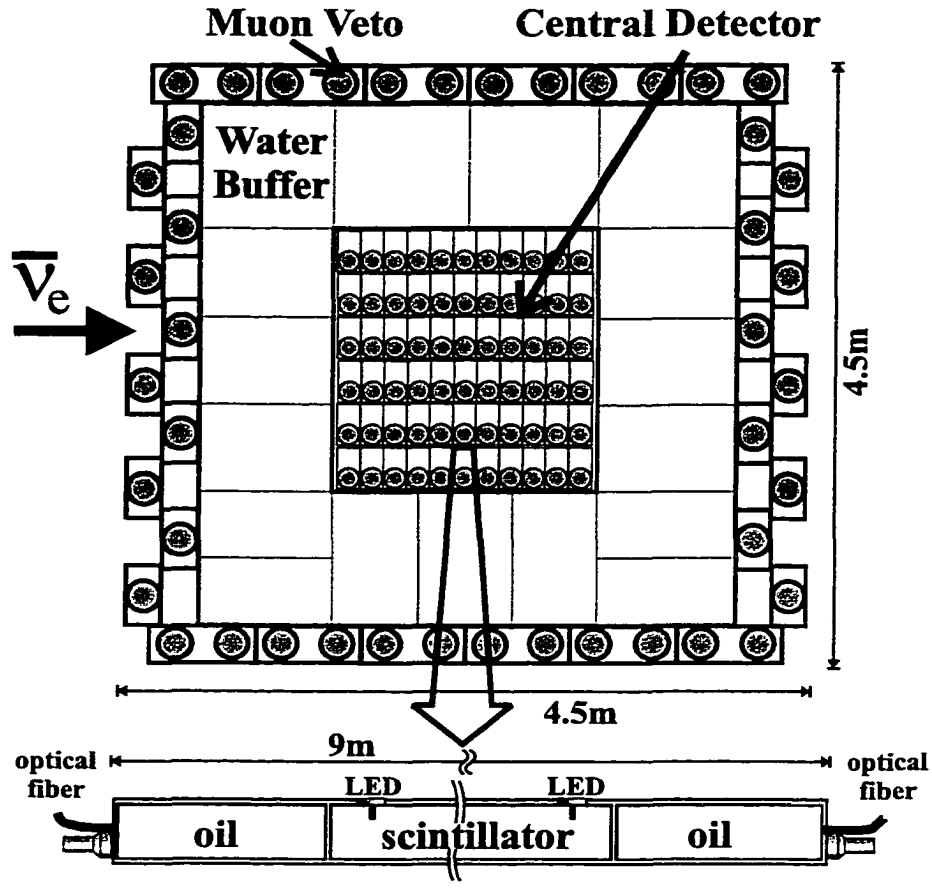

Figure 2.2: The Palo Verde detector. One of the 66 target cells with PMTs, calibration LEDs, and optical fiber flashers is shown lengthwise at the bottom.

The detector was oriented such that the $\bar{\nu}_e$ flux was approximately perpendicular to the long axis of the cells.

The target liquid scintillator, developed by Bicron Radiation Measurement Products and our collaboration [48], was composed of 36% pseudocumene, 60% mineral oil, and 4% alcohol, and was loaded with 0.1% Gd by weight. This formulation was chosen to yield long light transmission length for scintillation light, good stability, high light output, and long term compatibility with acrylic. The optimal efficiency of the PMT photocathodes is at 440 nm, and a small amount of the wavelength shifter was added to the mixture to shift the scintillation light to this wavelength. The loading of Gd into organic scintillator required a large development effort. While early

batches were unstable and showed rapid degradation of light attenuation length, the scintillator filled into the central detector cells performed remarkably well having an average attenuation length upon filling of 11.5 ± 0.1 m at 440 nm, as shown in the bottom of Figure 2.3, and losing only ~10% of this transmission length over the two years of detector operation (see Section 3.2 below). The attenuation length of the final formulation as a function of wavelength for several early batches of the scintillator is also shown in the top of Figure 2.3.

The central volume of the detector was surrounded along the sides by a 1 m thick buffer of high purity deionized water (about 105 tons) contained in steel tanks which, together with the oil buffers at the ends of the cells, served to attenuate gamma radiation from the laboratory walls and outer detector components as well as neutrons produced by cosmic muons passing outside of the detector. The low Z of water minimized the neutron production by nuclear capture of stopped muons inside the detector and had a high efficiency for neutron thermalization.

The outermost layer of the detector was an active muon veto counter, providing 4π coverage. It consisted of 32 twelve meter-long PVC tanks (from the MACRO experiment[49]) surrounding the detector longitudinally, and two endcaps. The endcaps were mounted on a rail system to allow access to the central detector. The horizontal tanks were read out by two 5-inch PMTs at each end; the vertical tanks were equipped with one 8-inch PMT at each end while the endcaps used 3-inch PMTs. The liquid scintillator used in the veto was a mixture of 2% pseudocumene and 98% mineral oil, with a light attenuation length at 440 nm in excess of 12 m. The detector is pictured in Figure 2.4 with the veto endcap both closed and rolled away to reveal the target cells and water buffer tanks.

A schematic of one channel of the central detector's front-end electronics is shown in Figure 2.5. Each channel could be digitized by either of two identical banks of electronics. The $\bar{\nu}_e$ signal (described below) consists of a pair of time correlated sub-events. The dual bank system allowed both sequential parts of the event to be recorded with no deadtime by switching between banks. Due to the large dynamic range of energy in the data of interest (30 keV to 10 MeV, or 1 to 400 photoelectrons typically), each PMT had both a dynode and anode output connected to ADCs, as

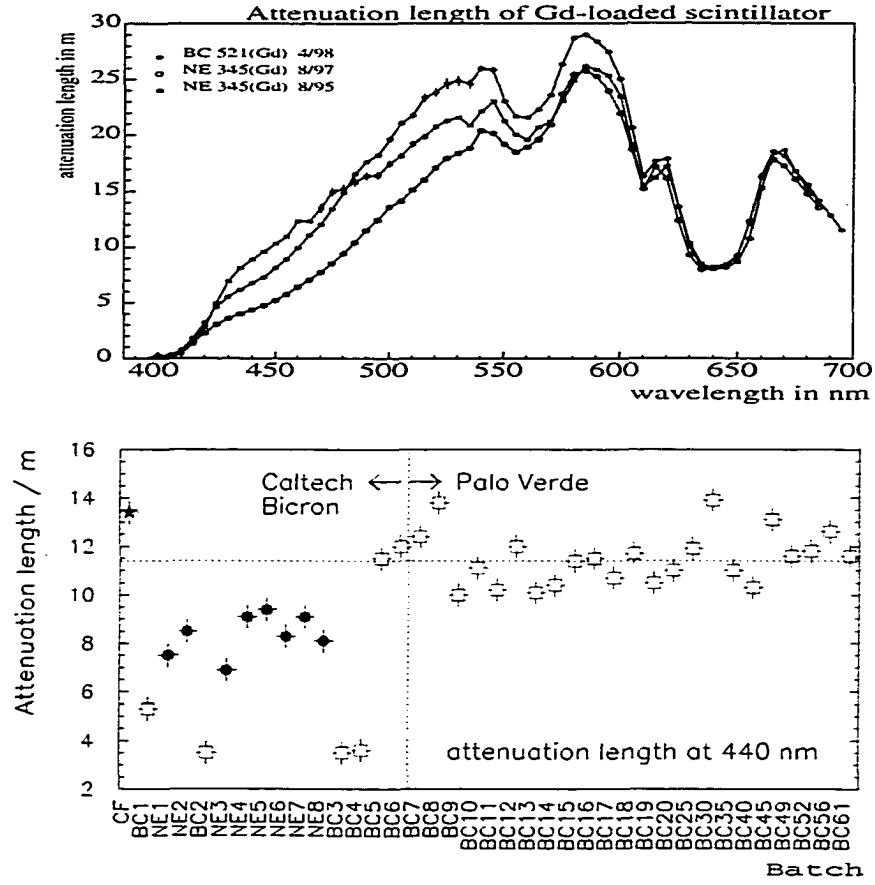

Figure 2.3: Top: The liquid scintillator attenuation length for light as a function of wavelength for three developmental batches. The scintillation light of the scintillator was 440 nm to match the optimal PMT photocathode efficiency. Bottom: The liquid scintillator attenuation length for all the batches mixed. The batches on the left were mixed at CalTech and Bicron during development, while the batches on the right were mixed at Palo Verde. The column farthest left "CF" is a mixture of mineral oil and pseudocumene used as a reference.

well as a $\times 2$ amplifier feeding three discriminator thresholds for the trigger and TDCs. The higher TDC threshold served to avoid crosstalk from large signals in adjacent channels while the lower threshold allowed timing information to still be available at the single photoelectron level. The relative time of arrival from each end of a cell was used to reconstruct longitudinal position. The measured PMT pulse charge at each end, corrected for light attenuation based on the distance traveled in the cell, allowed

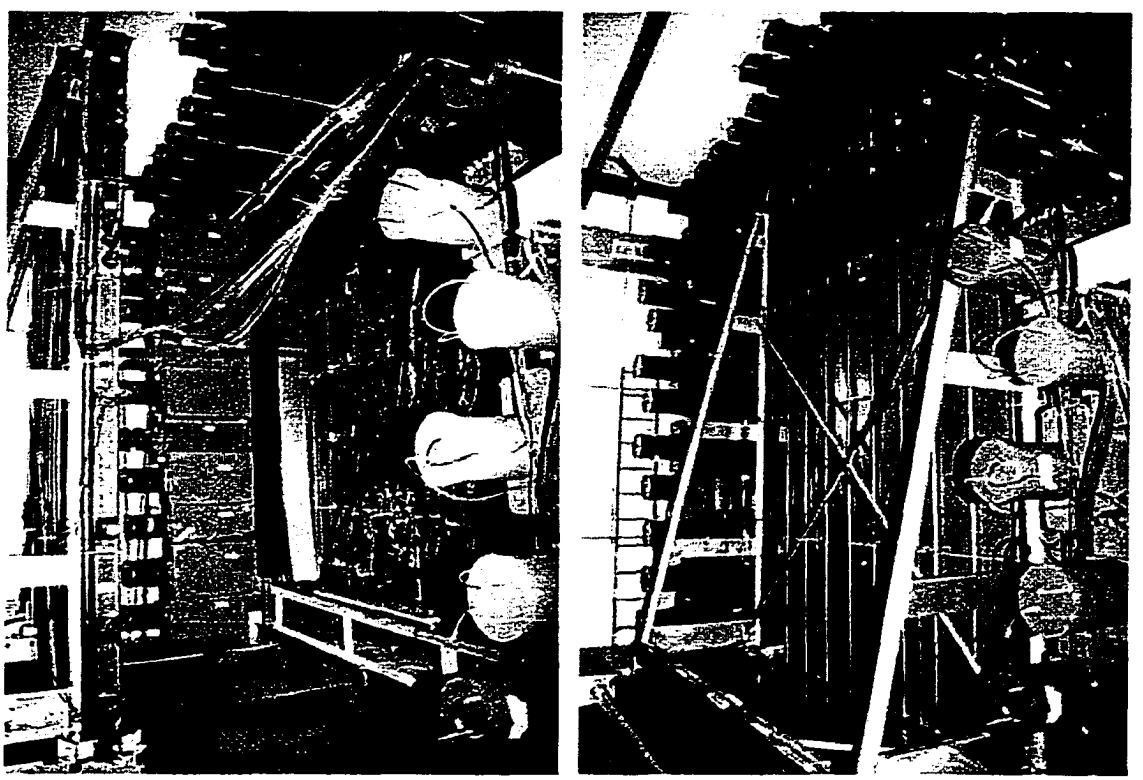

Figure 2.4: Looking at the detector towards the near end. Left: The detector with the veto endcap open. The cabling was instrumenting the ends of the target cell array, which was surrounded by the water buffer tanks. Right: The detector with the veto endcap closed. The 4π veto was composed of a barrel of MACRO tanks and the roll-away endcap assemblies.

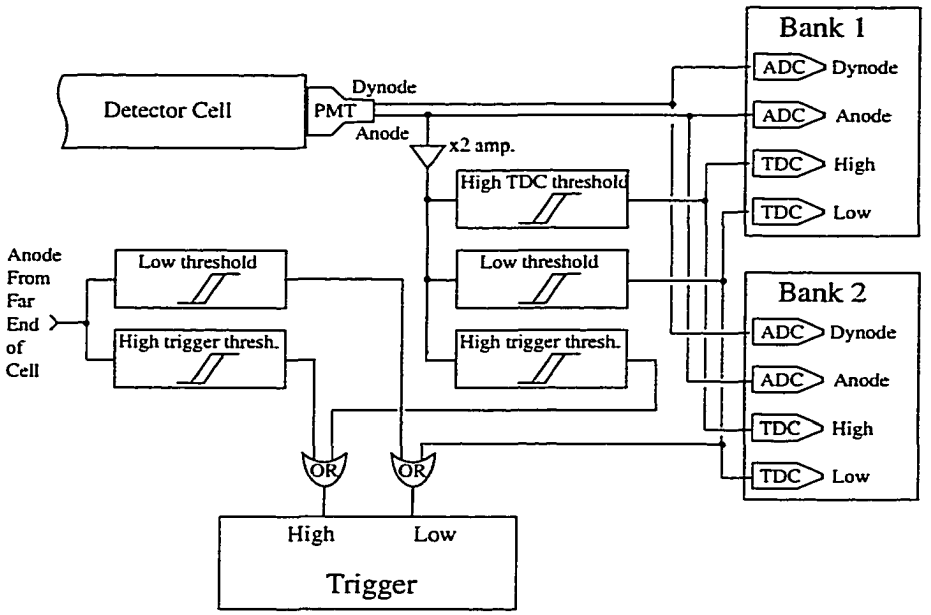

Figure 2.5: Schematic of the front-end electronics. The charge and timing of PMT pulses were read out by two banks of ADCs and TDCs. The timing information was discriminated with two thresholds to avoid crosstalk and retain dynamic range. Trigger discriminator signals from the two PMTs in each cell were OR'ed and input to the trigger circuit.

energy reconstruction. In order to simplify the design while maintaining good signal linearity, no active components were used in the signal path to the ADCs.

The data acquisition system (DAQ) consisted of a FASTBUS system with an event builder designed for the ALEPH experiment [50], and online data recording was handled by a MicroVAX using in-house software. A separate PC controlled the calibration and high voltage systems.

Each cell was connected to the trigger via the OR of the discriminated signals from the two PMTs. Signals were tagged according to two thresholds: a *high* threshold corresponding to ~600 keV for energy deposits in the middle of the cell and a *low* threshold corresponding to ~40 keV, or one photoelectron at the PMT. The *low* trigger threshold also served as the lower TDC threshold.

The trigger, which had a decision time of around 40 ns for each of the two sequential sub-events, used a Field Programmable Gate Array (FPGA) to search for

patterns of energy deposits in the central detector, and could be reprogrammed easily to change trigger conditions as needed for calibrations[51]. The calibration trigger conditions normally used included an override which could produce a trigger on demand (for LED, fiber flasher and pedestal events), a trigger on any single *high* hit, and a trigger on any single *low* hit. The pattern searched for to trigger on $\bar{\nu}_e$ events is discussed below. The trigger was tested for efficiency using the LED flashers connected to the target cells and found to have $< 0.1\%$ deadtime.

Each veto tank-end was connected to a discriminator whose threshold was approximately at the single photoelectron level. A veto signal, defined as a coincidence of hits at both ends of a veto tank, disabled the central detector trigger for 10 μs in order to avoid most related activity related to the passing muon. The threshold of the veto tanks was estimated to be around 3 MeV near the ends of the side veto tanks, the location with the worst efficiency for registering a muon, as estimated by looking at the energy spectra of the veto tanks. Typical veto rates were $\sim$2 kHz, of which $\sim$ 1700 are muon passages and the remainder mainly γ's from the lab walls. With each event, the time and hit pattern of the previous muon registered in the veto counter was recorded along with information as to whether or not the muon passed through the target cells. The operations of looking for veto hit coincidences, recording of hit patterns, and timing were performed by an FPGA in a reprogrammed spare trigger module.

Knowledge of the veto inefficiency for both stopping muons (*one hit missed*) and through-going muons (*two hits missed*) is useful for understanding and tracking some sources of backgrounds in the anti-neutrino data set. The stopping muon inefficiency, $4 \pm 1\%$, was found by measuring the percentage of hits found in the bottom set of veto counters with no corresponding entering hit. This measurement was cross-checked by using measured veto tank energy thresholds for seeing a hit in combination with a Monte Carlo model of expected cosmic muon angular distributions. Most of this inefficiency was due to the geometry of the side tanks of the veto where there was little overlap among adjacent tanks.

The relative *one hit missed* inefficiency was tracked in the data set by looking for un-vetoed Michel electron events from stopped muons decaying in the central

detector, enabling a measurement of relative veto performance during data taking.

The through-going muon inefficiency was measured using the rate of muon tracks found crossing the detector target without a veto hit, 0.45 Hz, which after correcting for detector livetime gives a *two hits missed* inefficiency of (0.07±0.02)%. Both Michel electron and muon track events were easily tagged in the data set by their large energy deposits relative to other kinds of events. The smallness of the *two hits missed* relative to the *one hit missed* inefficiency is due to correlations between incoming and outgoing muons, i.e. the high probability of muons going through the sides to also pass through either of the highly efficient top or bottom veto planes.

2.2 The $\bar{\nu}_e$ signal

The $\bar{\nu}_e$ signal was detected via the reaction $\bar{\nu}_e + p \rightarrow n + e^+$ as illustrated in Figure 2.6. Signal events consist of a pair of time-correlated sub-events: (1) the positron kinetic energy ionization and two annihilation γ's forming the prompt part and (2) the subsequent capture of the thermalized neutron on Gd forming the delayed part. By loading the scintillator with 0.1% Gd, which has a high thermal neutron capture cross section, the neutron capture time was reduced to ~27 μs from ~170 μs for unloaded scintillator. Furthermore, Gd de-excites by releasing an 8 MeV γ cascade, whose summed energy gives a robust event tag well above natural radioactivity. In contrast, neutron capture on protons releases only a single 2.2 MeV γ.

Background was rejected at trigger level using the detector segmentation by looking for coincidences of energy deposits matching the pattern of inverse beta decay. Each of the sub-events of a $\bar{\nu}_e$ signal was triggered by scanning the detector for a pattern of three simultaneous hits in any 3×5 subset of the cell array. This threefold coincidence, called a *triple*, must consist of at least one *high* trigger hit, due to either the positron ionization or neutron capture cascade core, and at least two additional *low* trigger hits, resulting from either positron annihilation γ's or neutron capture shower tails. The use of identical trigger requirements for the two *triples* was found to give rise to close to an optimal signal to noise ratio. Five μs after finding an initial *triple*, the trigger began searching for a delayed *triple*. The blank time is of the same

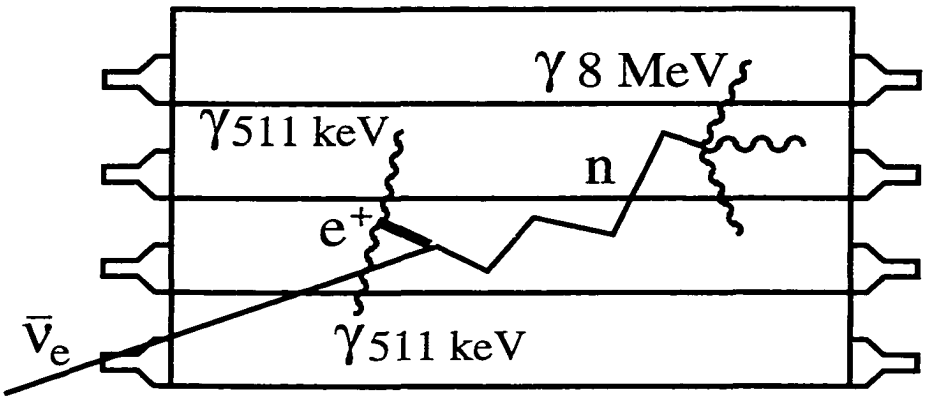

Figure 2.6: A schematic of the inverse beta decay signal. The ionization and annihilation of the positron forms a prompt triple coincidence of energy deposits, followed by the thermalization and capture of the neutron forming a delayed triple coincidence.

order of magnitude as the neutron thermalization time and suppressed possible false signals from PMT afterpulsing. If two *triples* are found within 450 μs of each other, the candidate $\bar{\nu}_e$ event was digitized for offline analysis. The time span between the prompt and delayed portions of an event is referred to as the *inter-event* time.

Along with the timing and charge of each PMT hit, the scaled rate of *low*, *high* and *triple* trigger hits, veto information, *inter-event* time, and the status of the LED and fiber calibration systems were recorded.

2.3 Expected $\bar{\nu}_e$ interaction rate

In order to calculate the expected $\bar{\nu}_e$ interaction rate in the detector, the status of the three reactors was tracked daily, and the fission rates in the cores were calculated based on a simulation code provided by the manufacturer of the reactors. This code used as input the power level of the reactors, various parameters measured in the primary cooling loop, and the original composition of the core fuel elements. Of the four isotopes — ^{239}Pu, ^{241}Pu, ^{235}U, and ^{238}U — whose fissions produce virtually all of the thermal power and as well as neutrinos, measurements of the neutrino yield per fission and energy spectra exist for the first three[52, 53]. The ^{238}U yield, which contributes 11% to the final $\bar{\nu}_e$ rate, was calculated from theory[54]. Further details

of the anti-neutrino flux calculation and the associated uncertainties are given in Appendix A.

This calculated $\bar{\nu}_e$ flux is then used to compute the expected rate of $\bar{\nu}_e$ candidates $N_{\bar{\nu}_e}$ at the detector as a function of the oscillation parameters Δm^2 and $\sin^2 2\theta$:

$$N_{\bar{\nu}_e} = n_p \int dE_{\bar{\nu}_e} \sigma(E_{\bar{\nu}_e}) \eta(E_{\bar{\nu}_e}) \times \sum_{i=1}^{3} \frac{\mathcal{I}_{\bar{\nu}_e,i}(E_{\bar{\nu}_e}) \left(1 - P_{osc,i}(\Delta m^2, \sin^2 2\theta, L_i, E_{\bar{\nu}_e})\right)}{4\pi L_i^2} \quad (2.1)$$

where $\sigma(E_{\bar{\nu}_e})$ is the inverse beta decay cross section[55], $\eta(E_{\bar{\nu}_e})$ is the (energy dependent) detector efficiency, n_p is the number of free protons in the target, and $\mathcal{I}_{\bar{\nu}_e,i}$ is the source strength of reactor i at distance L_i with oscillation probability $P_{osc,i}$. The actual interaction rate in the detector target before detector efficiency, which will be referred to here as $R_{\bar{\nu}_e}$, is obtained by setting $\eta(E_{\bar{\nu}_e}) = 1$ (see Figure 2.7). The energy spectrum actually measured in the detector is the energy of the positron created by the inverse beta decay. This spectrum is approximately $E_{\bar{\nu}_e} - 1.8$ MeV, slightly modified by the kinetic energy carried away by the neutron (~50 keV).

Previous short baseline experiments which measured the rate of $\bar{\nu}_e$ emission by reactors have found good agreement between calculated and observed neutrino flux by using largely the same method of calculation. A high statistics measurement at Bugey[32], in particular, found excellent agreement both in spectral shape (χ^2/n.d.f. = 9.23/11) and in absolute neutrino yield (agreement better than 2%, dominated by systematic errors). These previous generation experiments prove that the reactor antineutrino spectrum, i.e. the $\bar{\nu}_e$ flux at the distance $L = 0$, is well understood.

The expected $\bar{\nu}_e$ interaction rate in the whole target, both scintillator and the acrylic cells, is plotted in Figure 2.8 for the case of no oscillation from July 1998 to July 2000. Around 220 interactions per day were expected with all three units at full power. The periods of sharply reduced rate occurred when one of the three reactors was off for refueling, the more distant reactors contributing each approximately 30% of the rate and the closer reactor the remaining 40%. The short spikes of decreased rate are due to short reactor outages, usually less than a day. The gradual decline in

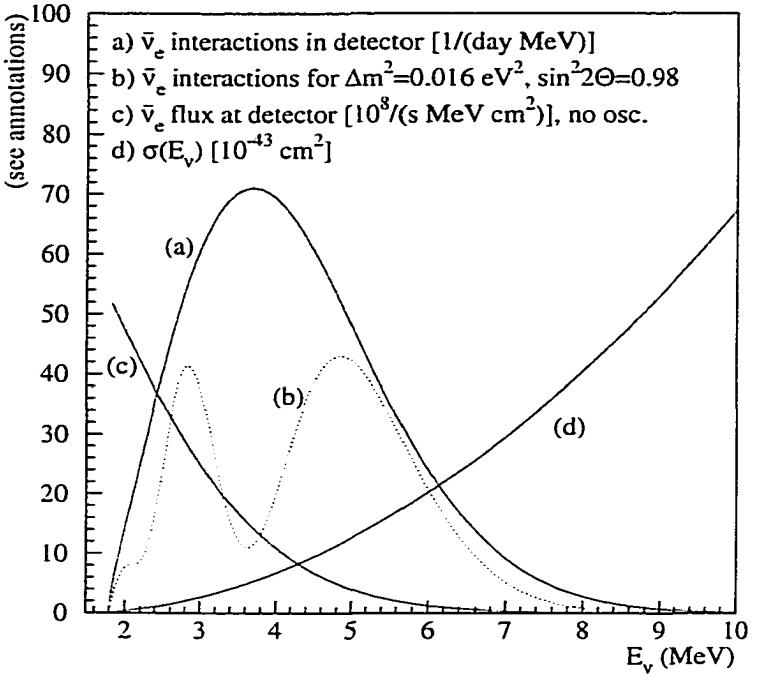

Figure 2.7: $\bar{\nu}_e$ flux, inverse beta decay cross section, and $\bar{\nu}_e$ interaction spectrum at the detector for the case of no oscillations. Also included is the expected interaction rate for the best fit of the original Kamiokande atmospheric result.

rate between refuelings is caused by fuel burnup, which changes the fuel composition in the core and the relative fission rates of the isotopes, thereby affecting slightly the spectral shape of the emitted $\bar{\nu}_e$ flux.

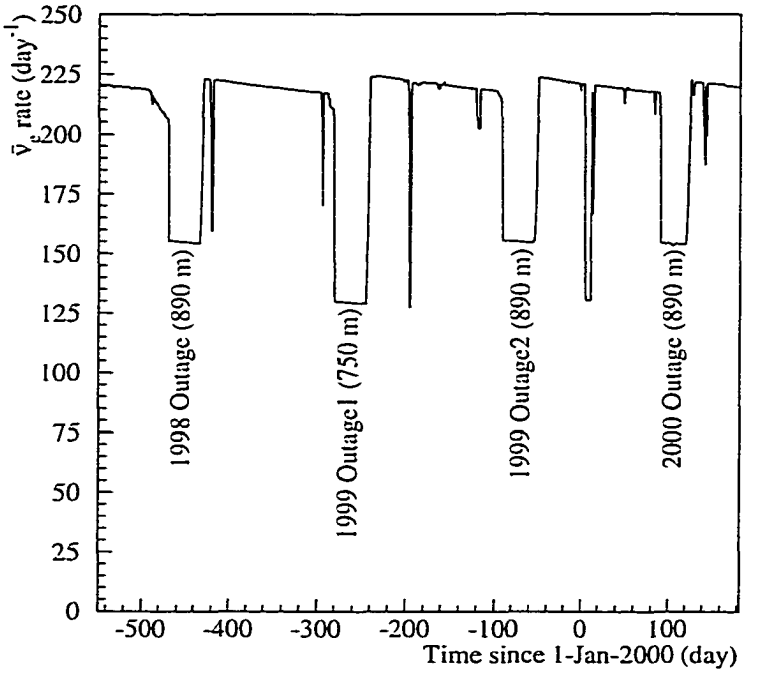

Figure 2.8: The calculated $\bar{\nu}_e$ interaction rate in the detector target. The long periods of reduced flux from reactor refuelings are used for background subtraction in one analysis of the data. The decreasing rate during the full power operation is a result of the changing core composition as the reactor fuel is burnt.

Chapter 3

Calibration

In order to maintain constant data quality during running, a program of continuous calibration and monitoring of all central detector cells was followed. Blue LEDs installed inside each cell were used for relative timing and position calibration. Optical fibers at the end of each cell, also illuminated by blue LEDs, provided information about PMT linearity and short term gain changes. LED and fiberoptic scans were performed once a week. Radioactive sources were used to map the light attenuation in each cell, for absolute energy calibration, and to determine detection efficiencies for positrons and neutrons. A complete source scan was undertaken every 2–3 months.

3.1 LED and optical fiber calibrations

As seen in Figure 2.2, every cell of the central detector had two LEDs, one at each end at a distance of 90 cm from the PMTs. These blue LEDs, which provided fast light pulses with a rise time comparable to scintillation light, were used for timing calibrations needed for position reconstruction along the cell's axis.

The difference in pulse arrival time between the two PMTs of a cell Δt was described as a function of the position z with an effective speed of light c_{eff}, an offset z_0 and a small nonlinear correction $f(1/Q)$:

$$\Delta t = (z - z_0)/c_{\text{eff}} + f(1/Q_n, 1/Q_f). \tag{3.1}$$

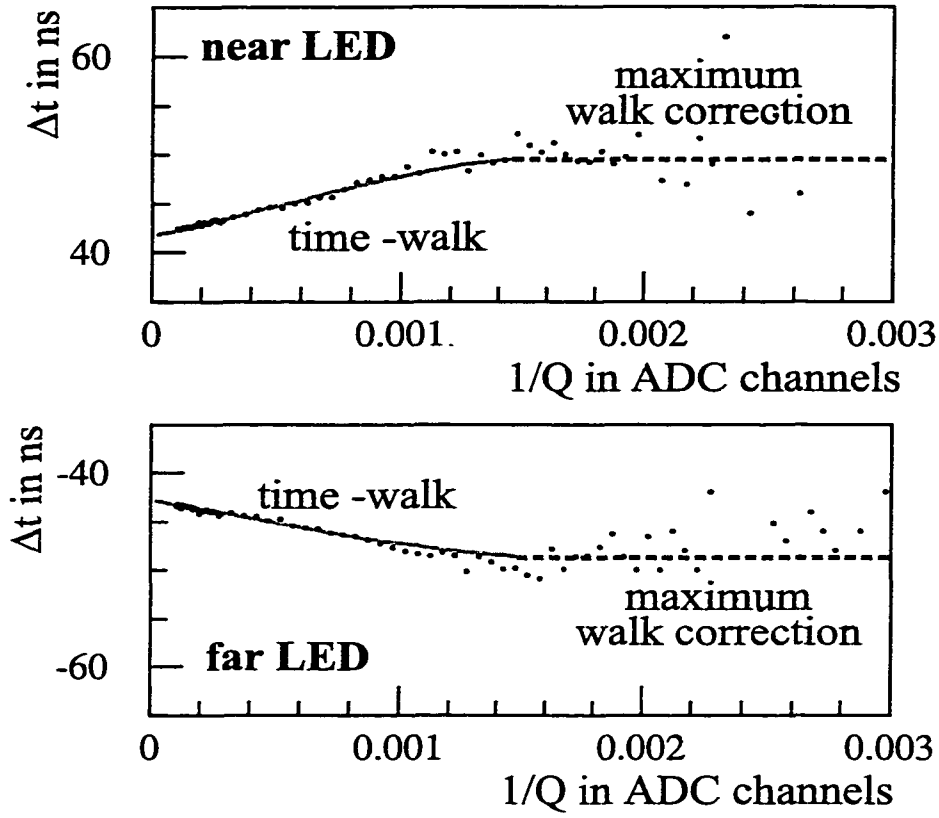

Figure 3.1: Time difference Δt as function of $1/Q$ for near- and far-end LEDs of a cell, illustrating the *time-walk* effect. For values of Q close to the discriminator threshold the *time-walk* correction was kept at a constant (maximum) level.

The correction $f(1/Q_n, 1/Q_f)$, a function of both near and far PMT pulse charge Q, describes the dependence of the pulse height due to *time-walk* in the leading edge discriminators used in the front-end electronics. To extract these calibration parameters and compensate for the *time-walk* effect, a third order polynomial was fit to Δt versus $1/Q$ (see Figure 3.1). The intercepts at $1/Q = 0$ for the two LED positions provide c_{eff} and z_0, while the slopes were used to parameterize the *time-walk* correction.

In order to check the suitability of longer wavelength 470 nm LED light to measure timing properties of ~425 nm scintillation light, data taken with a ^{228}Th source at several longitudinal positions were reconstructed with the LED timing calibration

parameters. Comparing the reconstructed positions with the actual source positions, the effective speed of light measured with the LED system was found to be on average 3.6% lower than that with the sources. A simulation of the light transport in a cell with various indices of refraction and attenuation lengths of the scintillator suggested that the small discrepancy in Δt between LED and scintillation light was due to the difference in attenuation length. The correction factor was found to be constant over several months. Weekly LED scans were therefore used to correct for short term variations in Δt and a constant correction factor was applied to the effective speed of light.

The fiberoptic system included 15 blue LEDs, each illuminating a bundle of 12 fibers. The light output of each LED was measured in two independent reference cells with PMTs checked to be linear over the whole dynamic range of the LEDs. By taking a run which scanned through all light intensities and mapping each PMTs response relative to the reference cells, the nonlinear energy response of the PMTs was calibrated. Low intensities were used to determine the single photoelectron gain of each PMT, which was used to correct for changes from the nominal gain setting of 4×10^7.

3.2 Scintillator transparency and energy scale calibration

In addition to weekly LED and fiberoptic calibrations, the energy response of the scintillator was measured every three months using a set of sealed radioactive sources. Eighteen 2.4 mm diameter teflon tubes ran along the length of the detector allowing insertion of the sources adjacent to any cell at any longitudinal position. the response of each PMT as a function of position was measured by recording the Compton spectrum from the 2.614 MeV γ of a ^{228}Th source at seven different locations along each cell.

Monte Carlo simulation found that the upper half maximum of a Gaussian function fitted to the Compton spectrum is relatively independent of resolution; this point

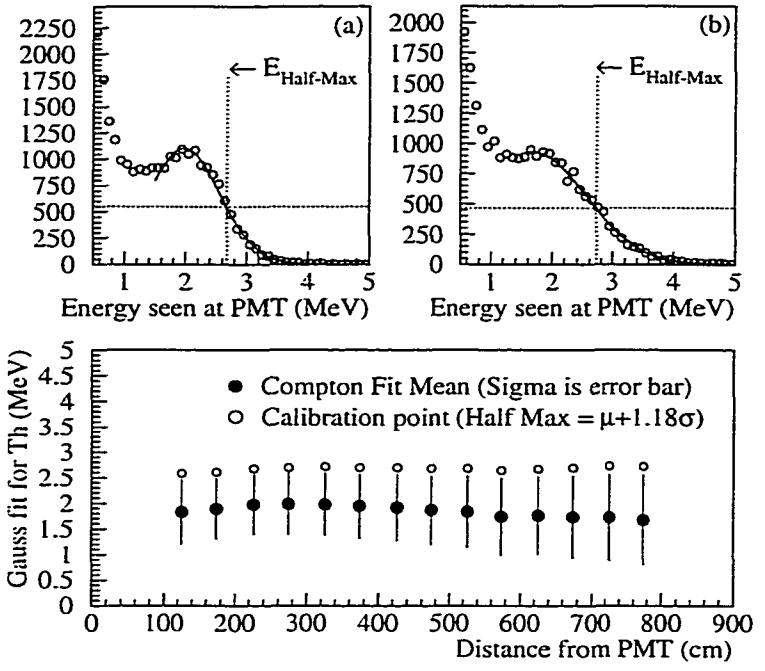

Figure 3.2: To find a feature in the Compton spectra of the ^{228}Th source scans, we verified, using Monte Carlo simulation, that the upper half-max of a gaussian fit to the spectra were independent of resolution. Top left and right: simulated spectra and fitted half-maxima for a location near and far from the PMT ((a) best and (b) worst resolution). Bottom: upper half-max of fitted spectra as a function of distance from the PMT, showing the relative independence of the benchmark as a function of resolution. The response of 12 PMTs at 14 distances were simulated. Note that the overall energy scale was not set using the response to the Th source scan, but was tuned to reconstruction of a 1.275 MeV γ from ^{22}Na source scan (see text).

was therefore used as the benchmark of the cell response (see Figure 3.2). The response versus distance from the PMT, shown in Figure 3.3 for one cell, was then fit to the phenomenological function $\exp(p_0 + p_1 z) + \exp(p_2 + p_3 z)/z$, where z is source longitudinal distance from the PMT. The effective attenuation length of the scintillator (including multiple total reflection on the acrylic walls) was generally between 3–4 m and over a year was found to change on average $\lesssim$1 mm/day, demonstrating that the Gd scintillator was remarkably stable.

The overall energy scale was determined from the position of the 1.275 MeV peak

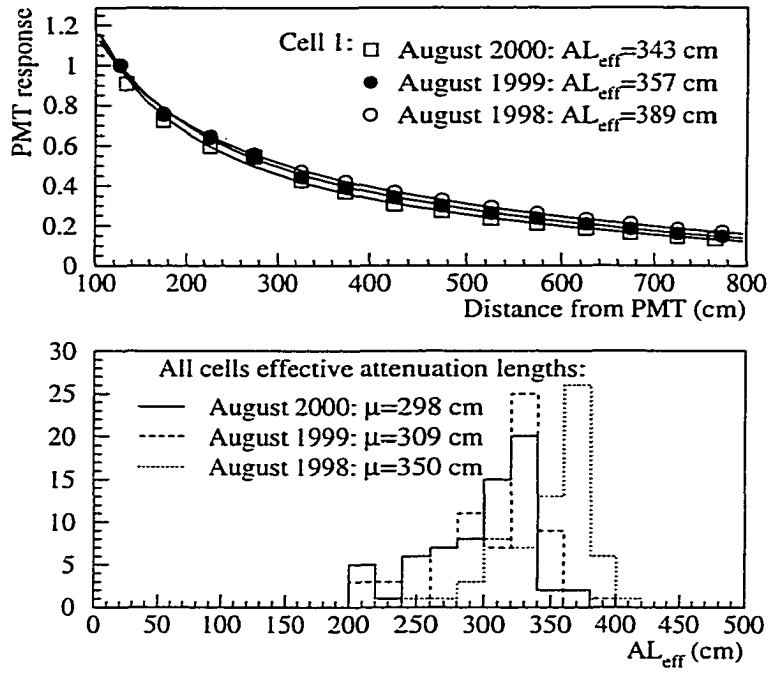

Figure 3.3: Effect of aging on Gd loaded scintillator. Top: ^{228}Th Compton edge position at seven different longitudinal locations along a typical cell. The three curves are from calibrations taken at the beginning, middle and end of the data set. The curves are normalized at the location nearest to the PMT. Bottom: Effective attenuation lengths for all 66 cells from the three calibrations.

of a ^{22}Na source, and then verified by taking data with several γ sources in different energy ranges: ^{137}Cs (0.662 MeV), ^{65}Zn (1.351 MeV), ^{228}Th (2.614 MeV), and the capture of neutrons (8 MeV) from an Am-Be source. The gamma cascade from neutron capture was modeled according to measurements of the emitted spectrum[56]. In contrast to homogeneous detectors which measure total absorption energy peaks, 25% of the detector target mass consisted of the inert acrylic of the cell walls, which absorbs some energy. The Monte Carlo simulation was therefore used to find the correct final distributions of energy detected from single and multiple scattering of the γ's. The total energy reconstructed for data and Monte Carlo for each source is plotted in Figure 3.5. The data were matched with Monte Carlo simulation for the ^{22}Na spectrum in Figure 3.15 to find the overall energy scale and to the spectra

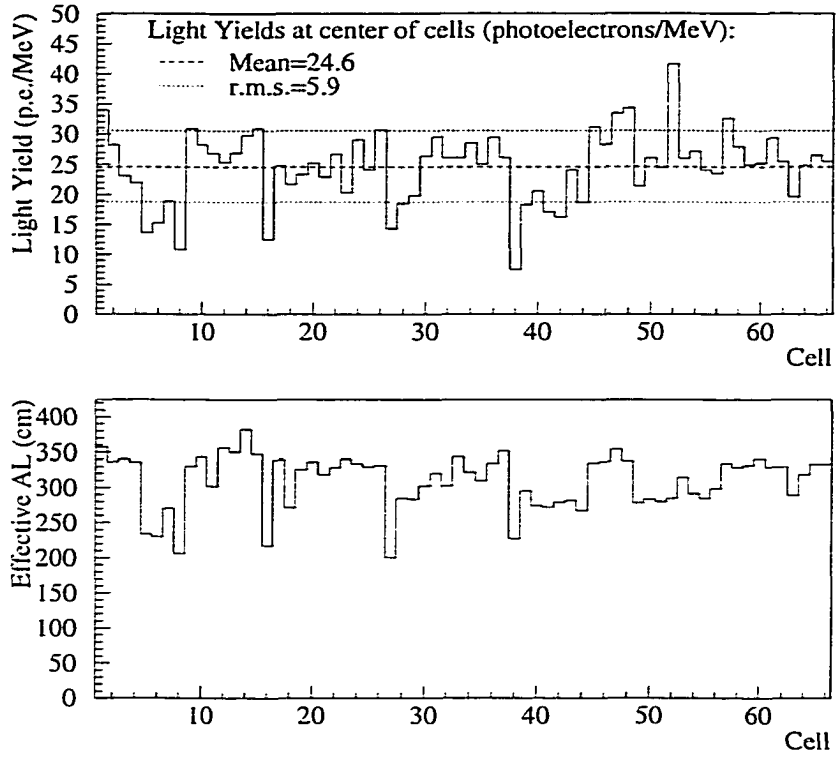

Figure 3.4: Top: The light yield at the center of the 66 target cells, after PMT photocathode quantum efficiency. Bottom: Effective attenuation lengths of the cells at the time of the light yield measurement. The variation of light yield from cell to cell is mainly due to variations of liquid transparency from cell to cell.

in Figure 3.5 to assure that the scintillator response was linear over the energies of interest. The light yield after PMT quantum efficiency, plotted for all the target cells in Figure 3.4, was found to be ~25 photoelectrons per MeV in the center of the cells, but varied substantially from cell-to-cell because of different liquid transparencies. The agreement for three of the four sources in Figure 3.5 is good, the exception being ^{228}Th, in which the data have a consistently higher Compton scattering peak than Monte Carlo predicts. This discrepancy, studied extensively but never understood, was consistent across data taken at all locations in the detector target and therefore did not affect the scintillator transparency calibrations.

To check the quality of the energy calibration, we reconstruct the same ^{228}Th

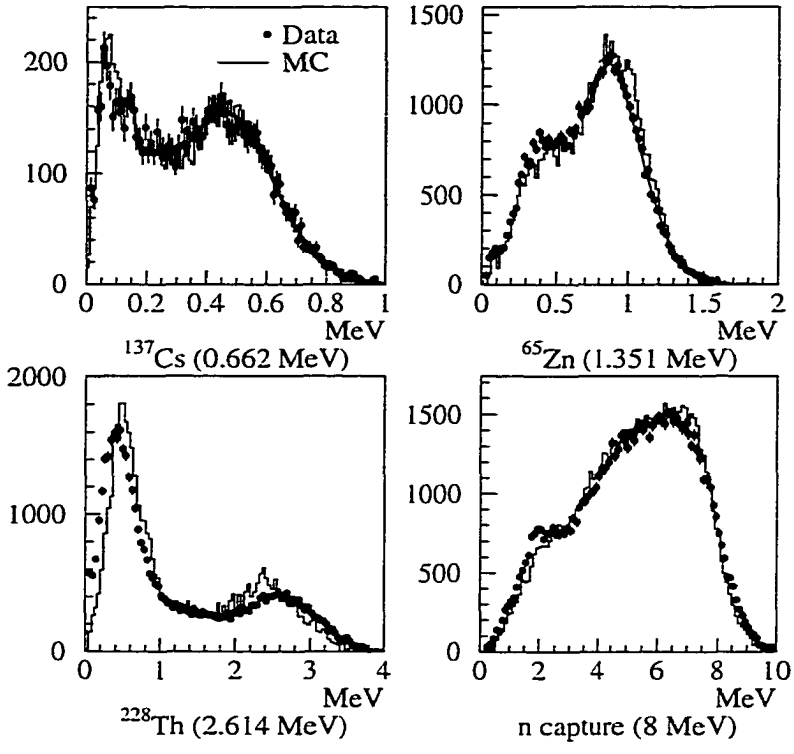

Figure 3.5: The total energy reconstructed for various sources compared for data and Monte Carlo.

scan used for the calibration to see the reconstructed Compton scattering energies. This effectively closes the loop of the energy scale calibration, in that we should get back exactly the energies we defined from the Compton scattering fits as our energy reconstruction parameters. As shown in Figure 3.6, the reconstruction shows little bias with respect to longitudinal position, z, and a spread of almost 10% from location to location. The lack of a bias shows that the chosen attenuation function is reasonable. One can see that the accuracy of the reconstruction gets worse towards the edges of the cell nearest the PMTs, a natural consequence of these locations being at the steepest part of the attenuation correction function. As mentioned above the reconstructed ^{228}Th scans were also used to correct the position reconstruction parameters for the speed of scintillation light in the cells.

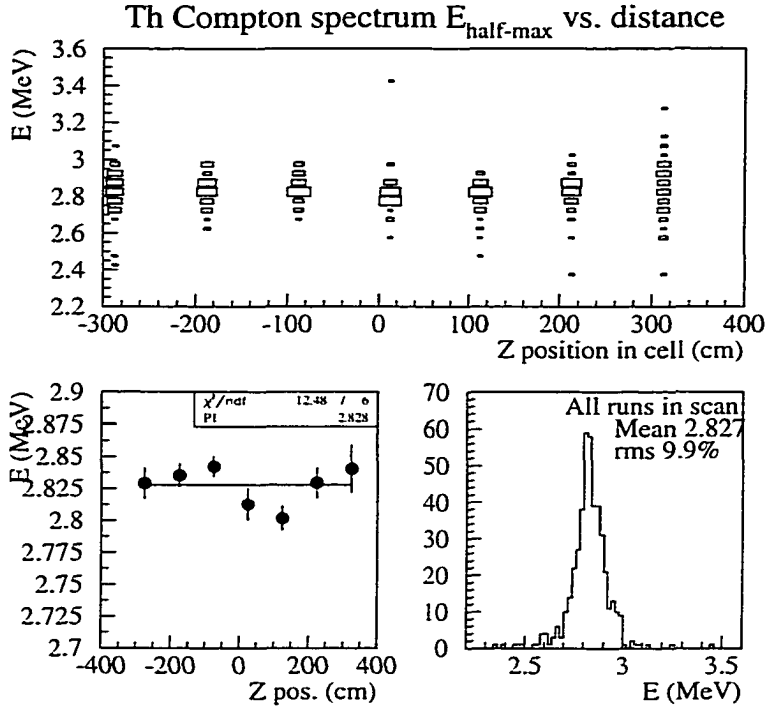

Figure 3.6: The reconstructed Compton scattering energies ($E_{\mathrm{half-max}}$) from the May 1999 Thorium source scan (7 longitudinal positions in all 18 source insertion tubes). Since the attenuation length and individual target cell energy responses were measured using Thorium scans, this serves as useful feedback on the internal consistency of the energy calibration.

3.3 Monte Carlo simulation

The $\bar{\nu}_e$ efficiency of the detector is a relatively strong function of event location and, to a lesser extent, of time due to scintillator aging. A further complication comes from the trigger efficiency being a function of threshold (voltage) while only energy (charge) was recorded. For this reason a Monte Carlo model which includes a detailed simulation of the detector response, including the PMT pulse shape, is used in the estimate of the overall efficiency for $\bar{\nu}_e$ detection. A variety of measurements were performed to cross-check that the Monte Carlo accurately models the detector response.

The physics simulation program is based on GEANT 3.21[57]. This code contains

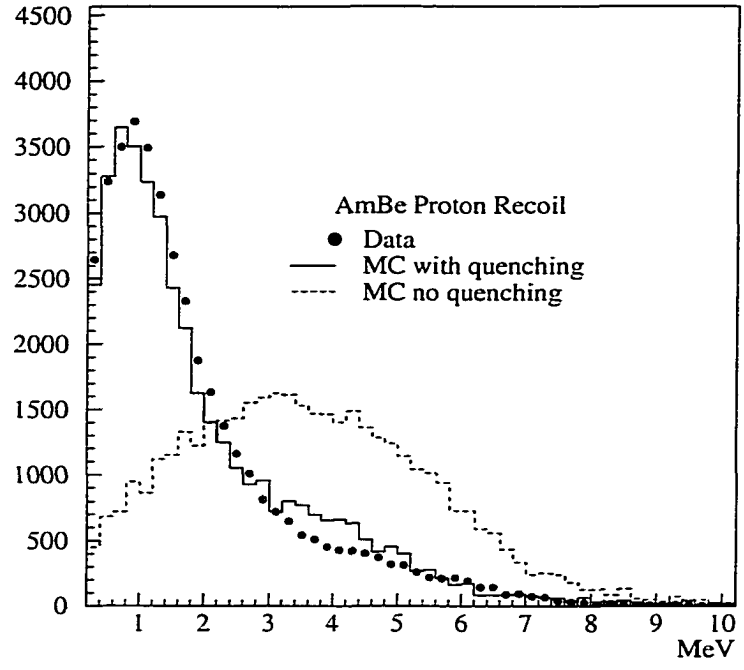

Figure 3.7: The effect of scintillator light quenching in proton recoil events from the Am-Be source. Accurate simulation of proton recoils is useful for understanding $\bar{\nu}_e$ backgrounds resulting from fast neutrons.

the whole detector geometry and simulates the energy, time, and position of energy deposits in the detector. Hadronic interactions are simulated by GFLUKA[58] and the low energy neutron transport by GCALOR[59].

Scintillator light quenching, parameterized as a function of ionization density, is included in the simulation[60]. Since quenching is important for simulation of proton recoils from fast neutrons, a possible source of backgrounds, we checked the simulation parameters using data taken with an Am-Be neutron source (described below in Section 3.4.2). This source emits neutrons with kinetic energies up to 10 MeV, creating proton recoils in the detector scintillator. By digitizing any energy deposits seen during the neutron release, the high ionization density of these recoiling protons was used in cross-checking the parameters which control scintillator light quenching in the simulation, as shown in Figure 3.7

The event reconstruction program reads the output of this physics simulation

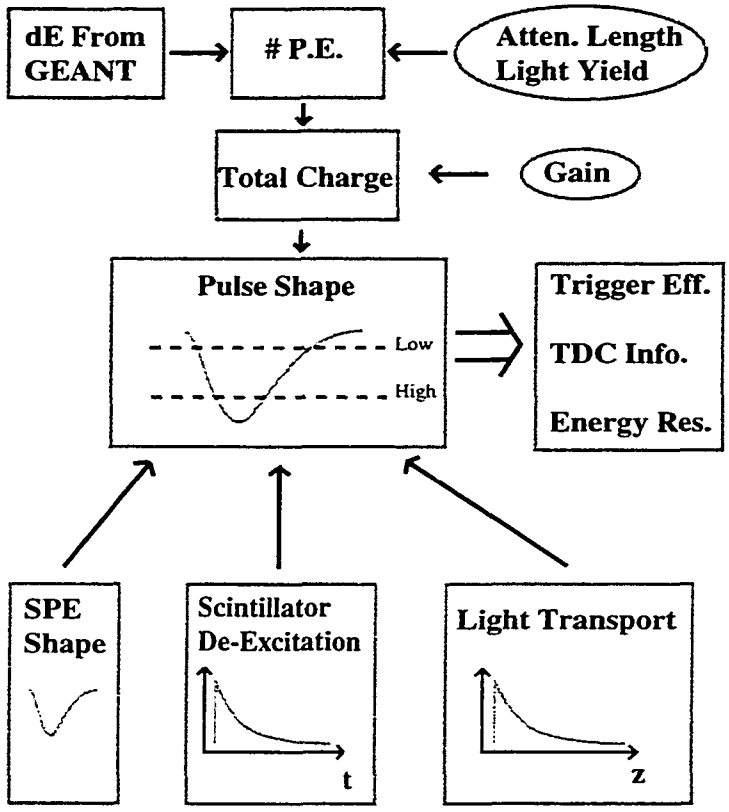

Figure 3.8: Schematic of the Monte Carlo detector simulation program. In order to convert energy deposited into accurate threshold simulation, calibration information from the detector was used to reconstruct PMT pulses.

and then applied the second step of the Monte Carlo, the simulation of the detector response as PMT pulses which were then converted into time and amplitude digitizations and trigger hits. A logical scheme of this detailed detector simulation is shown in Figure 3.8. The simulation of PMT pulses uses the calibrations of light yield, attenuation length, and PMT gain in combination with parameterizations of light arrival times and PMT single photoelectron characteristics to construct the pulses. This procedure is described in detail in this section.

The calibrations discussed above empirically provide the scintillator light yield (photoelectron/MeV) and attenuation function for each cell, which in turn provide the number of photoelectrons $\overline{N}_{p.e.}$ expected for a simulated energy deposit. A 10%

uncertainty is then introduced to the measured energy to account for the energy scale uncertainty found with the Thorium scan cross-check described above in Figure 3.6. The total charge of the pulse is then modeled by sampling a Poisson distribution with mean $\overline{N}_{p.e.}$. This number of simulated photoelectrons is then folded with the PMTs measured gain from the fiberscan data, modified by a PMT first stage gain variance ($10N_{p.e.}$). This latter quantity takes into account the fact that a given PMT pulse charge is a statistical function of the number of electrons produced in the first amplification stage of the PMT, typically around 10.

To simulate the pulse shape, an arrival time is assigned to each photoelectron, and individual photoelectron pulses (whose shape is derived from real data) are summed into a final pulse. The calculated arrival time of each photon is a combination of two processes, scintillator de-excitation and propagation along the cell.

The latter distribution is parameterized by the distance traveled to the PMT, larger distances giving larger variances, using a light transport simulation of 2×10^7 photons. The results of this ray tracing simulation are shown in Figure 3.9. The arrival time of photons becomes clearly broader as the distance traveled along the cell increases. The distributions of arrival times follow an exponential fall since the probability of photon extinction due to scintillator attenuation is an exponential function of the path length in the scintillator. In the case shown the attenuation length of the liquid is 11.0 m. The simulation also found that the distribution of geometric path lengths were monotonically decreasing without a discernible rise during the early arrivals: direct light is the most efficiently collected. It is interesting to note that the simulation confirmed the observed double exponential form of the effective light attenuation of the acrylic cells.

The final parameter of the pulse simulation is the time structure of the scintillator light emission or de-excitation. This is also reasonably modeled as an exponential decay, since scintillation begins in organic liquid scintillators within a few hundred picoseconds of the excitation. This parameter can be tuned by using large pulses, which remove the effects of photoelectron statistics, leaving the photon arrival time function as the main generator of the pulse features. We compared simulations of a large (14 MeV) pulse to oscilloscope traces of a PMT pulse from a muon traversing

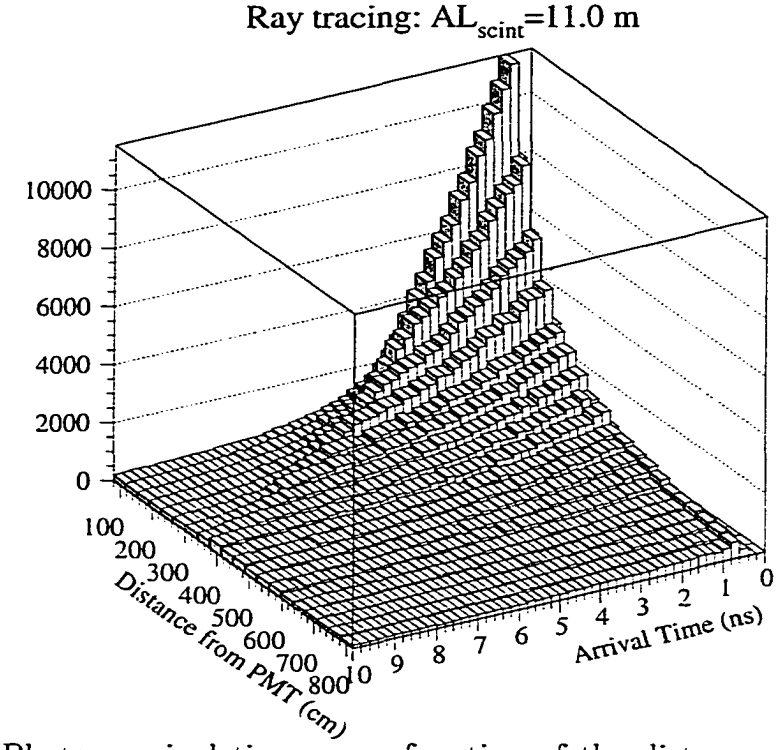

Figure 3.9: Photon arrival times as a function of the distance of an energy deposition from the PMT from a ray tracing simulation. This simulation was used to parameterize the effects of cell geometry on pulse shapes.

a target cell. The rise time and peak location of the simulation matches those of the observed pulse reasonably well. The observed pulse has a slightly longer tail distribution than the simulation, which is probably due to afterpulsing in the PMT (see Appendix B); whereas the slight feature in the oscilloscope trace tail is a signal reflection, an artifact of a cabling impedance mismatch.

After summing the single photoelectrons into a PMT anode signal, the resulting pulse is reshaped and attenuated. The signals of the detector travel through delay cables for storage during the trigger decision time, which tend to cause this reshaping. The pulse is then be stepped through to look for discriminator hits. While the nominal thresholds of the TDC and trigger discriminators are known, variations from cell to cell and over time need to be accounted for. The data set itself was used to track and account for these variations by looking at TDC hits in each channel. By plotting the charge collected in each event versus the efficiency for producing a TDC hit,

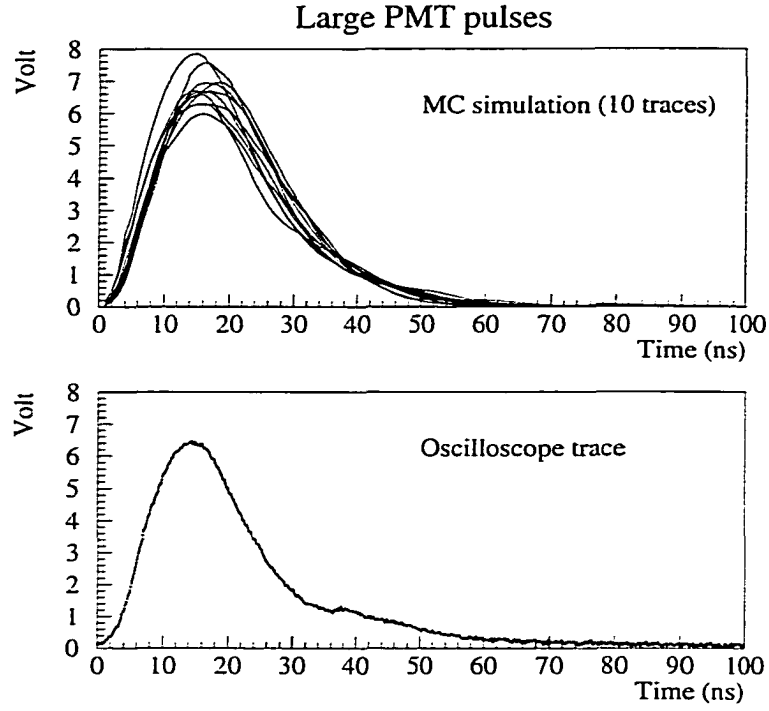

Figure 3.10: The simulation of large pulses isolates the scintillator de-excitation and light transport parameters. At bottom is an oscilloscope trace of a large PMT pulse from a muon traversing a target cell. At top are the Monte Carlo simulations of several 14 MeV pulses.

an effective TDC threshold was measured for that channel as shown in Figure 3.11. Variations in this efficiency are mainly due to changes in the front end power supply voltages. Since PMT pulse heights are linear with respect to charge in this range, the relative variations were used to correct the nominal threshold values of the simulation. High TDC threshold corrections were also applied to the *high* trigger thresholds.

The trigger threshold simulation for each cell was compared to data taken with a ^{22}Na β^+ source near the center of each cell. This was done for all cells on two occasions, in February 1999 and January 2000. The trigger conditions were loosened for these data, a single *low* hit producing a trigger and the event tagged if a *high* threshold was crossed. By plotting the reconstructed energy for each event versus the efficiency for a *high* trigger tag as done for the TDC threshold tracking, an effective *high* trigger threshold in MeV for that location in the cell was determined. The

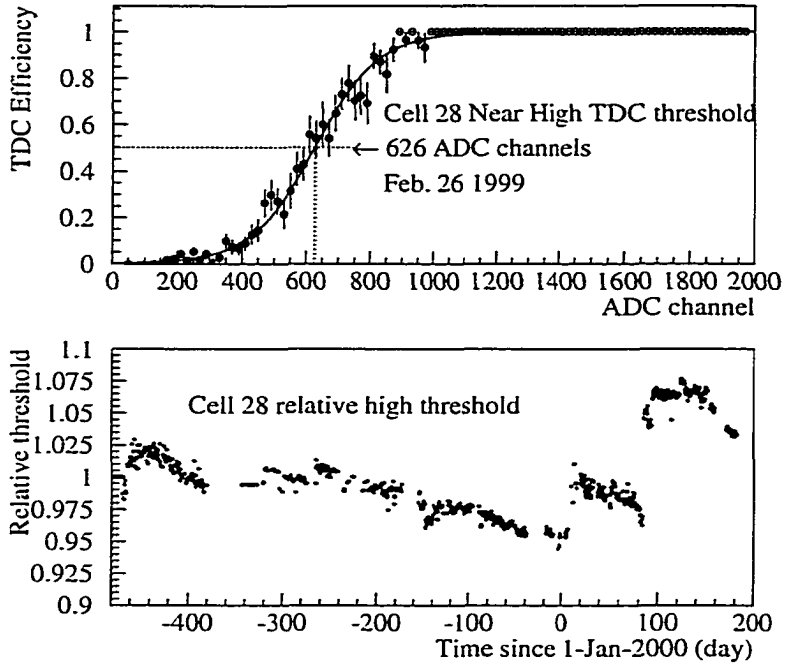

Figure 3.11: TDC threshold for a cell as monitored in the data over the two years of data taking by the method described in the text. These relative corrections were used by the detector response simulation to adjust TDC thresholds.

low threshold was measured similarly. The trigger threshold, defined as the energy at 50% efficiency, is plotted for all 66 cells in Figure 3.12 for the 1999 calibration runs. The value of the nominal *high* trigger threshold used in the simulation was tuned to get the average measurement here to agree, but the cell-by-cell variation was not used for tuning. These variations are due to differences in the scintillator light yield and attenuation lengths across cells, and is reasonably well reproduced by the simulation. Cell 38, for example, had the scintillator with the worst performance among target cells and therefore has the highest effective trigger threshold. The cell-by-cell agreement between simulation and measurement, ~6%, is expected from the known inaccuracy of each cell's energy scale. The *low* thresholds have much worse agreement from cell to cell. The *low* threshold, nominally 30 mV, is at the single photoelectron level; the difficulty in accurately reproducing *low* thresholds PMT by PMT is an inherent limitation of the single photoelectron simulation which must

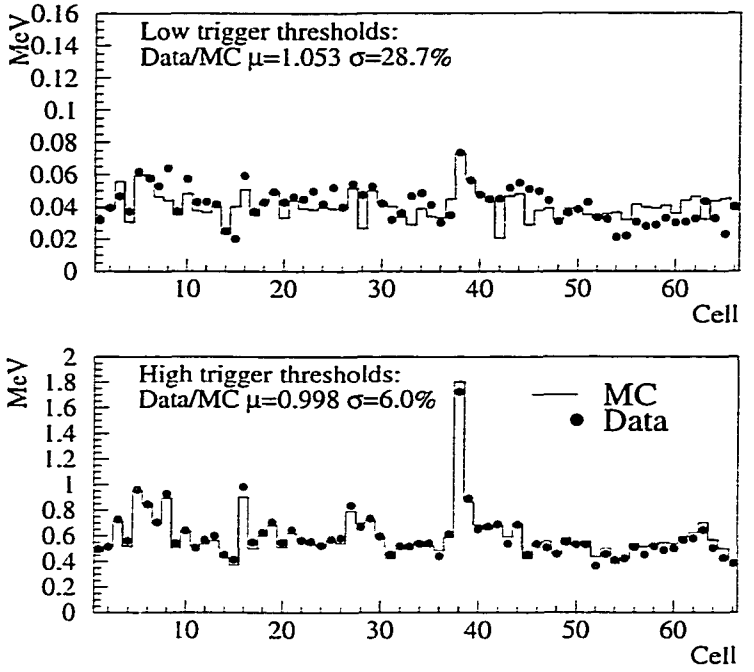

Figure 3.12: A comparison of the *high* and *low* trigger thresholds from data and Monte Carlo. The data were taken with a ^{22}Na source at the center of each cell.

assume that all such pulses are shaped similarly. The 2000 threshold calibration runs give a similar result.

TDC thresholds were checked by the same algorithm, plotting the threshold hit efficiency versus reconstructed energy. A more direct check of the TDC simulation, however, compares the position reconstruction for data and Monte Carlo simulation. Figure 3.13 shows the longitudinal position of the third largest energy deposit in each event for a ^{22}Na calibration run, representing the position reconstruction of the energy deposited by one of the two positron annihilation γ's. Since these energy deposits tend to be small ($\sim$100 keV), some fraction have one or both PMTs responses below the TDC *low* threshold. These events constitute the tails of the distribution in Figure 3.13 since only the relative signal amplitude was used for position reconstruction. The narrower central peak is populated by events with TDC information available. One can see a few percent overpopulation of the central peak for the Monte Carlo simulation, indicating perhaps a tendency towards overestimating the *low* TDC hit

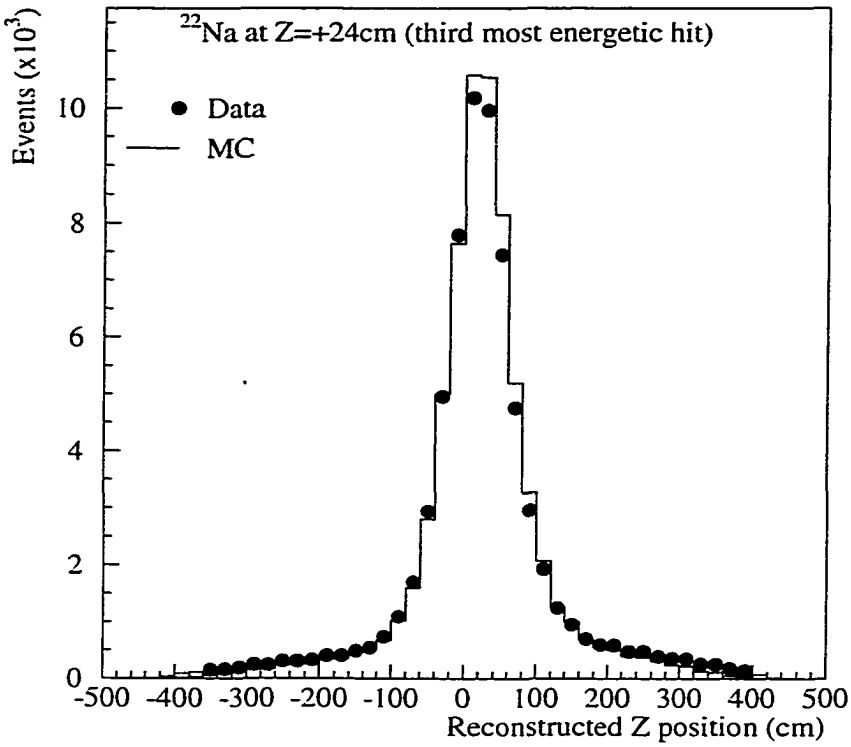

Figure 3.13: A comparison of the reconstructed z position of an annihilation gamma energy deposit from a ^{22}Na run, between data and Monte Carlo. The long tails are due to lower energy deposits which do not have TDC for position reconstruction, in which case relative signal amplitudes are used, resulting in worse resolution. Statistical errors in the plot are negligible.

frequency. There also is indication of a ~1 cm bias towards the near end. The resolution of both tail and central peaks are in good agreement, however. The discrepancies are deemed small enough to not seriously affect the quality of the Monte Carlo efficiency prediction, and may be the result of a slight tilt of the detector. This tilt, the far end being approximately 1 cm lower than the near end, could cause these biases by shifting some of the target liquid towards the negative Z end of the detector. Such a shifting of mass would have only second order effects when averaged over the entire target mass.

3.4 $\bar{\nu}_e$ detection efficiency

While direct calibration of the anti-neutrino efficiency of the detector is impossible, close analogues of the prompt and delayed portions of the anti-neutrino signal are available in the form of ^{22}Na positron and Am-Be neutron sources. These sources differ in important aspects to that of the true $\bar{\nu}_e$ signal events in that the neutrons and positrons from the calibration sources originate externally of the target cells and have different kinematics than those of inverse beta decays; we therefore rely on the Monte Carlo simulation to bridge this gap and provide the estimate of the true $\bar{\nu}_e$ efficiency. The absolute efficiency of the detector for positron annihilations and neutron captures was measured using these sources, the results of which were then tested against the Monte Carlo simulation. In this way we verify the efficacy of the Monte Carlo.

3.4.1 Positron efficiency

The ^{22}Na source emits a 1.275 MeV primary γ which is accompanied 90% of the time by a low energy positron which annihilates in the source capsule. The primary γ can mimic the positron ionization of a low energy $\bar{\nu}_e$ event. This deposit, in conjunction with the positron's annihilation γ's, closely approximates the positron portion of a $\bar{\nu}_e$ event near the trigger threshold.

In four rounds of data taking the ^{22}Na source was inserted into the central detector at various locations chosen to provide a sampling of distances from the PMTs and edges of the fiducial volume. The source activity was known to 1.5%, allowing determination of an absolute efficiency. After applying the offline selections used for $\bar{\nu}_e$ prompt *triples* (see Section 4.1) and correcting for detector DAQ deadtime, the measured absolute efficiency was compared with the Monte Carlo prediction; the results are summarized in the top portion of Figure 3.14. Good agreement is seen in the average efficiency over all runs, and run-by-run agreement was 10%.

The energy spectra predicted by the simulation and measured in the data for the ^{22}Na runs were compared. The total energy seen in all cells and the energy detected in the three most energetic hits are plotted in Figure 3.15. The trigger thresholds

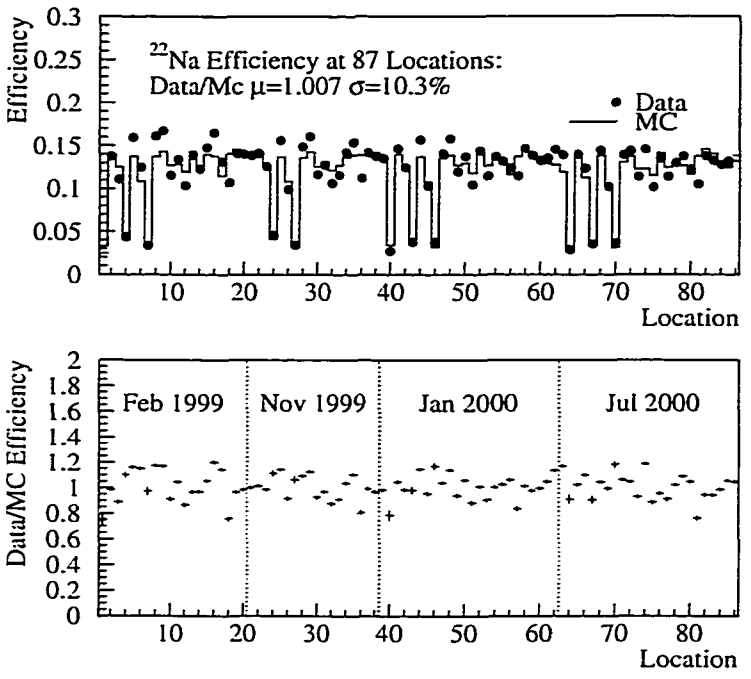

Figure 3.14: Comparison of data and Monte Carlo ^{22}Na detection efficiency for source runs taken at various locations in the fiducial target region. The runs were taken in four calibration sets spanning the two years of $\bar{\nu}_e$ data taking. Top: The absolute detection efficiencies from the measurements and the simulation. Bottom: The ratio of measured to calculated efficiencies for the 87 runs. Locations at the edges of the detector tend to have lower efficiencies.

can be seen in the spectra: the *high* trigger threshold is the rising edge at around 0.5 MeV in the spectrum of the most energetic hit (E_1), and the *low* trigger threshold is the rising edge at around 50 keV of the third most energetic hit (E_3).

We investigated further the large run-by-run spread in data versus Monte Carlo simulation positron efficiency estimates. Two possible causes of this variation are the simulation using inacurrate trigger thresholds or the uncertainty in the cell-by-cell energy scale. For two of the rounds of ^{22}Na efficiency runs, we also have measurements of the *high* trigger thresholds. Either of the causes listed above for the efficiency inaccuracy could cause the differences in the trigger thresholds of data and Monte Carlo simulation. The threshold discrepancies were therefore used to tune either the

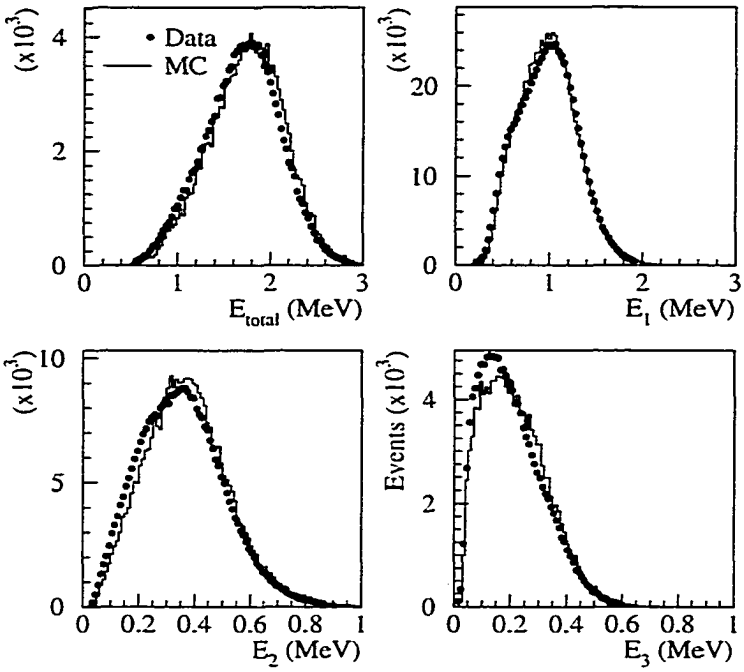

Figure 3.15: The Monte Carlo simulation and the data compared for the spectra of total energy and first, second, and third most energetic hit (E_{total}, E_1, E_2, and E_3) digitized in the ^{22}Na calibration runs.

Monte Carlo trigger thresholds cell-by-cell or to artificially introduce the same energy scale inaccuracies to the Monte Carlo, and the simulation ^{22}Na efficiencies were then rechecked against data to look for improvement. While both tuned versions of the simulation returned *high* trigger thresholds which agreed exactly with data, neither showed improved agreement with data for the efficiency estimates. The case of energy scale tuning is shown in Figure 3.16.

3.4.2 Neutron efficiency

A similar procedure was used to check the neutron capture detection efficiency. The Am-Be neutron source is attached to one end of a thin (7.5mm) NaI(Tl)-detector, which tagged the 4.4 MeV γ emitted in coincidence with a neutron. The device, pictured in Figure 3.17, can be inserted between target cells up to five meters into

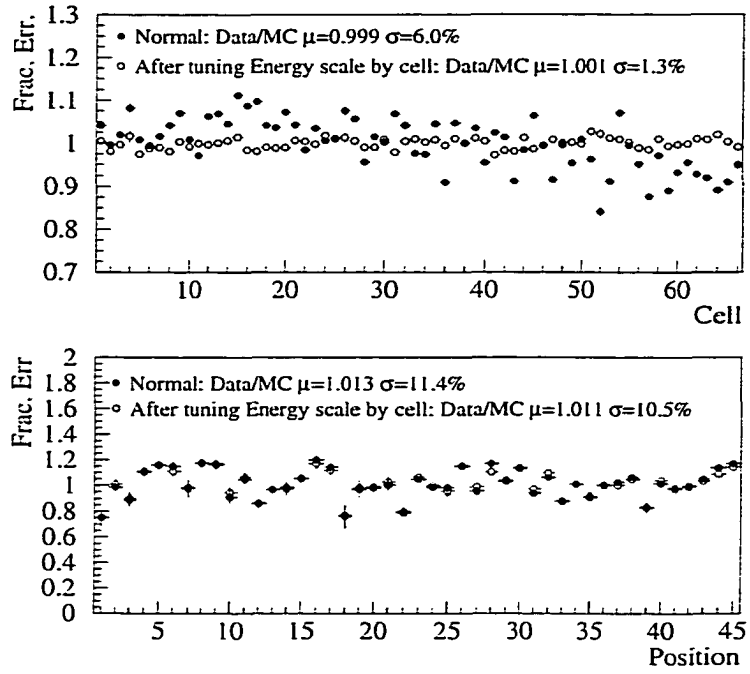

Figure 3.16: ^{22}Na calibration runs before and after using the trigger threshold measurements to tune the cell-by-cell energy scales. Top: The high trigger thresholds measured versus Monte Carlo prediction. Bottom: The ^{22}Na detection efficiency measured versus Monte Carlo simulation. No improvement is seen in the Monte Carlo efficiency estimates.

the detector. To take calibration data we modify the trigger to accept a NaI(Tl) tag as a forced prompt trigger. The energy seen by the NaI detector is digitized in a spare ADC channel and calibrated using various source Compton scattering spectra as is done for the energy calibrations of the detector target. The NaI(Tl) tag forces the digitization of the 4.4 MeV γ as the prompt part of an event and opens a 450 μs window for neutron capture; this is the same coincidence window used in the $\bar{\nu}_e$ runs. By counting the percentage of time correlated triggers follow the NaI event, one is able to unfold an Am-Be neutron detection efficiency. We took Am-Be calibration data in 58 locations throughout the detector in four sets of runs.

There are a few corrections which need to be accounted for in this calibration. The most important is the effect of the neutrons emitted by the Am-Be source itself.

Figure 3.17: A NaI detector with an Am-Be source attached (top right) at the end can be introduced up to 5m deep into the central detector. The left picture shows the unfolded detector and the bottom right picture shows the folded device.

With an activity of the source corresponding to ~500 Hz, there is a sizeable chance for a previously emitted neutron to be present and interfere with the detection of the capture of the neutron truly associated with the NaI trigger. This effect is large, and can be seen in the plot of the time between the NaI trigger and the neutron captures (*inter-event* time) for the cumulative Am-Be runs plotted in the top of Figure 3.18. There is a surplus of events at long *inter-event* times relative to what one would see for a simple single neutron capture (dashed line in the plot). This surplus comes from neutrons or γ's from the source being detected which are not correlated to the NaI trigger. To account for this and get the correct *inter-event* time structure, a code was written to simulate the Am-Be decays and the associated activity in the detector. It accounted for the Monte Carlo prediction of neutron capture efficiency and the characteristic time structure of the capture, including a predicted ~5 μs thermalization time. The code also accounted for the activity of the source itself by tracking how many neutrons were likely to be present during the process, and how likely these were to be captured, as well as how likely the γ and proton recoils accompanying a neutron emission were to create a *triple*. Finally it accounted for the 2% chance of veto activity suppressing the trigger. The simulation *inter-event* time

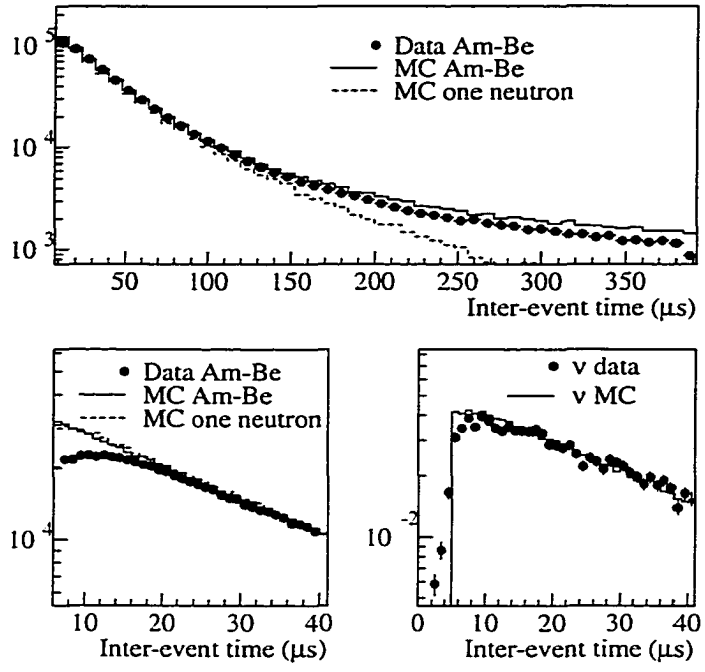

Figure 3.18: Top: The Am-Be data and simulation *inter-event* times. Included as well is the single neutron capture time prediction of the Monte Carlo. There is a discrepancy at both long and short *inter-event* time scales. Bottom left: The short *inter-event* times for data, single neutron simulation, and Am-Be simulation, showing an apparent shortage of capture events in the data at *inter-event* times less than ~ 20 μs. Bottom right: The neutron capture times of $\bar{\nu}_e$ candidates in the data and the Monte Carlo simulation for the same time scale shows no similar discrepancy. Note that the data *inter-event* variable has an intrinsic ~ 1 μs jitter in it due to the $\bar{\nu}_e$ trigger conditions.

spectrum for the runs is shown in Figure 3.18 as a solid line. Other backgrounds are not important, since they represent a very small contribution to the overall Am-Be induced *triple* rate after the neutron selection cuts (see Section 4.1).

The simulation gave back a time structure which agrees only marginally with the data. At longer *inter-event* times the simulation predicts more events than are seen. The Monte Carlo also overestimates the number of neutron captures at *inter-event* times less than ~ 20 μs.

The discrepancy at large *inter-event* times could be created by a gross misunderstanding of the Am-Be source strength. While nominally 500 Bq, if its true activity were a few kBq, this extra activity would tend to produce more triggers earlier and hence reduce the number of long *inter-event* times. We cross-checked the source strength by measuring the *triple* trigger rate seen in the detector, 625 Hz, when the source is inserted near the middle of the detector. The average Am-Be decay at that location of insertion has a 55% *triple* efficiency upon emission of a neutron and 67% efficiency for neutron capture at this particular location, leading to an expectation of 1.2 *triple*'s per Am-Be decay. After subtracting the native 50 Hz *triple* trigger rate, due to other activities present during normal data taking, this gives a 470 Hz source activity. Conversely a similar misunderstanding of the trigger activity induced by the Am-Be source, e.g. a true efficiency of 100% for the γ and proton recoils accompanying neutron emissions, would not produce this result. A neutron capture trigger efficiency which was incorrectly estimated would show up as a disagreement in the number of capture events following a NaI trigger. The possibility of false triggers in the NaI was checked and found to be negligible.

At shorter *inter-event* times, the long rise of the capture time seen in the data could either be a flaw in the simulation of Am-Be neutron thermalization time or it could be some detector effect seen only in the special trigger conditions, the latter being more likely. The former possibility is not very worrisome, since no evidence of such a large thermalization time is seen in the $\bar{\nu}_e$ data *inter-event* time structures (see the bottom right of Figure 3.18). The data and Monte Carlo simulation agree much better for $\bar{\nu}_e$ events in spite of the fact that the data includes a substantial background, most of which comes from fast neutrons. The data therefore span both extremes of Am-Be neutron energies, from 50 keV neutrons originating in inverse beta decay events to upwards of 100 MeV for cosmically induced neutrons.

In light of these discrepancies, to get an estimate of the agreement between data and Monte Carlo for neutron efficiency we tried looking at the area around 20-200 μs. This still covers a large percentage of the neutrons and avoids the discrepant parts. The resulting run-by-run efficiency estimates are plotted in Figure 3.19 for all Am-Be data taken. The Monte Carlo and data efficiencies agree over the runs remarkably

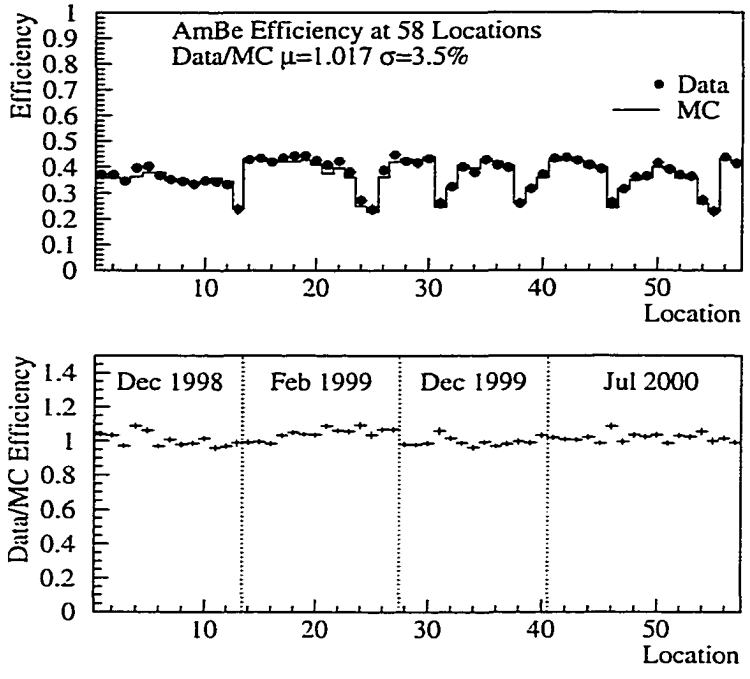

Figure 3.19: Top: Comparison of data and Monte Carlo detection efficiency for Am-Be neutron captures at various locations in the detector target. The runs were taken in four calibration sets spanning the two years of $\bar{\nu}_e$ data taking. A cut was applied to select *inter-event* times between 20 and 200 μs. Locations at the edges of the detector tend to have lower efficiencies. Bottom: The ratio of Monte Carlo and data detection efficiencies over the runs. Agreement overall is within 2%, with a small run-by-run spread of around 3.5%.

well, including at locations near the edge of the detector where the efficiency is reduced substantially due to neutrons escaping the fiducial volume. To see how much uncertainty there is regarding where exactly to cut on the *inter-event* time, a variety of different regions were tried, including 10-100 μs, 10-200 μs, 15-150 μs, 20-200 μs, which gave back average data/Monte Carlo ratios of 1.012, 1.003, 1.019, and 1.017 all with very similar variances of data/Monte Carlo efficiencies of around 3.5%.

As with the ^{22}Na runs, the energy spectra predicted by the simulation and measured in the data were compared. The total energy seen in all cells and the energy detected in the three most energetic hits is plotted in Figure 3.20. Note the small peak in E_{total} at ~2 MeV arising from neutrons being captured on hydrogen. The

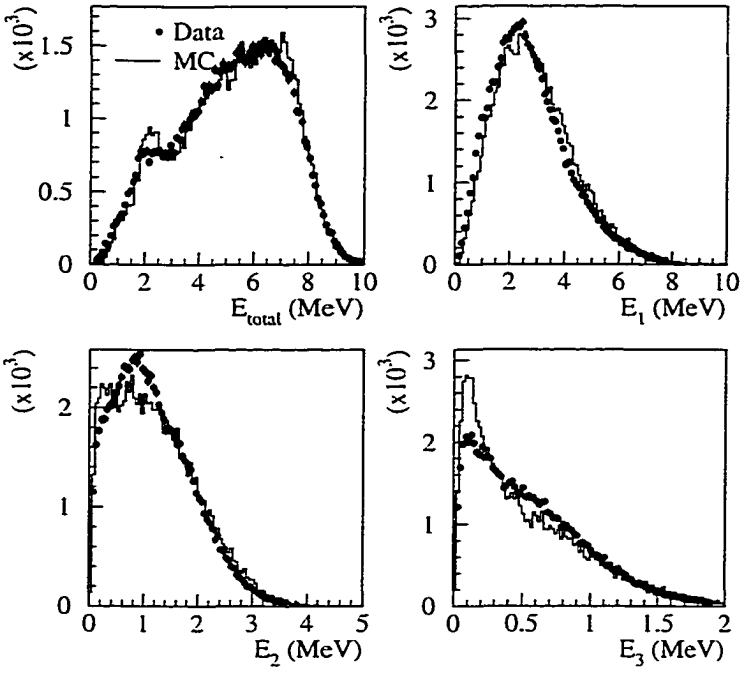

Figure 3.20: The Monte Carlo simulation and the data compared for the spectra of total energy and first, second, and third most energetic hit (E_{total}, E_1, E_2, and E_3) digitized in the Am-Be calibration runs.

differences in data versus Monte Carlo spectra for ^{22}Na and Am-Be are taken into account in estimating systematic errors by varying the $\bar{\nu}_e$ energy selection cuts.

The above cross-checks verify our ability to accurately generate the events, model the detector response, reconstruct the events, and correctly calculate the livetime of the DAQ system. Taken together these procedures complete the task of estimating our $\bar{\nu}_e$ efficiency.

The Monte Carlo simulation for $\bar{\nu}_e$ events models the expected interactions throughout the entire target, including the acrylic walls of the cells, since there is significant efficiency for inverse beta decay originating in the acrylic. The Monte Carlo simulation yields an average efficiency over all the detector as a function of $\bar{\nu}_e$ energy. The efficiency from the simulation is folded with the incident $\bar{\nu}_e$ spectrum (which may be distorted by oscillation depending on the hypothesis tested), to get the effective efficiency.

Chapter 4

$\bar{\nu}_e$ Selections and Backgrounds

The Palo Verde experiment took data from July 1998 through July 2000 in 759 runs, each on average about 12 hours long, for a total of 353 days of data. For the purposes of analysis, the data set was split into 13 periods, each having approximately a month of data. This grouping allows for reasonable statistics while giving some ability to track the data over time. Four of the periods represent data taken with one of the reactors off for refueling, three with one of the 890 meter baseline reactors off and one with the closer 750 meter distant reactor off. Thus, the entire data set has three different baseline combinations and neutrino fluxes. Before presenting the analysis and methods of extracting the $\bar{\nu}_e$ signal from the data, the signal selections and sources of backgrounds are discussed here.

4.1 $\bar{\nu}_e$ selection

The trigger rate for time-correlated events (two *triples* occurring within 450 μs) is $\sim$1 Hz. Most of those events are random coincidences of two uncorrelated *triple* hits, which occur individually at a rate of $\sim$50 Hz, mainly from natural radioactivity. While the two triple coincidences required by the trigger represent a rather strict online selection, the following additional offline cuts are applied in order to select neutrino events:

- The energy reconstructed in both prompt and delayed *triples* has at least one hit

with E >1 MeV and at least two additional hits with E >30 keV. No single hit is allowed to be greater than 8 MeV. These selections confirm the *triple* trigger condition, and are generally above the trigger thresholds to remove uncertainties stemming from threshold energy scales.

- The prompt *triple* resembles a positron, i.e. annihilation γ's each less than 600 keV, and together less than 1.2 MeV. This cut is the only one which treats the two *triples* asymmetrically, a feature which is used in the second analysis described in the next chapter. As seen in Figure 4.1, these selections are highly efficient yet reject backgrounds from fast neutrons (backgrounds discussed in the following sections).

- At least one of the two *triples* in the event has more than 3.5 MeV of reconstructed energy for rejection of γ backgrounds, most of which are from natural activities around the detector below 3 MeV. These low energy backgrounds can be clearly seen in Figure 4.2 as the low energy edge of the total energy spectrum of the $\bar{\nu}_e$ data. These γ's form most of the trigger rate through random coincidences.

- The prompt and delayed portions of the event are correlated in space and time (within 3 columns, 2 rows, one meter longitudinally, and 200 μs). Most random backgrounds will not have a correlation in space or time, whereas virtually all neutrons from $\bar{\nu}_e$ events tend to capture within a meter of the interaction vertex, with characteristic times of ~30 μs.

- The event started at least 150 μs (~5 neutron capture times) after the previous veto-tagged muon activity.

The $\bar{\nu}_e$ efficiency is ~33% at the trigger level and ~17% after the above selection cuts.

In addition to corrections for selection cut efficiency and trigger efficiency, detector livetime is a substantial correction to the number of neutrinos seen and deserves some comment. Deadtime comes from two sources, the DAQ and the muon veto. DAQ livetime is the ratio of the number of *triples* the DAQ was available for digitizing

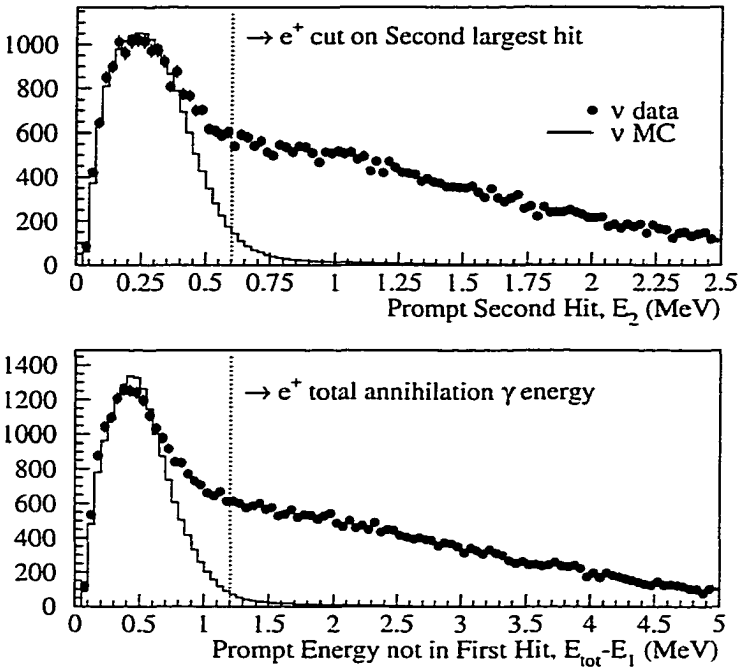

Figure 4.1: The selections for positrons for both Monte Carlo of pure $\bar{\nu}_e$ events and data candidates. These selections take advantage of the low energy annihilation γ's relative to higher energies for backgrounds in the second most energetic prompt hit. The data have a large contributions in the prompt portions from spurious neutron captures and proton recoils from fast neutrons (see Section 4.2) which are strongly rejected by these cuts.

to the total number of *triples* the trigger saw. These numbers are available from trigger scalers. The trigger livetime was measured to be >99.9%. The DAQ livetime varies with the *triple* rate, and was determined to be ~74% and ~92% for 1998 and 1999-2000 data sets respectively. The higher livetime in 1999 is due to improvements made in the trigger conditions. In 1998 all *triple* triggers were digitized to aid in diagnostics. For 1999-2000 the trigger was modified to accept only time correlated events in order to reduce the substantial DAQ deadtimes of 1998. The *high* trigger thresholds were reduced by 30% to increase the trigger efficiency. A misrouted trigger cable in the 1998 data set also reduced the trigger efficiency and was responsible for almost all of the change in Monte Carlo efficiency between the two years.

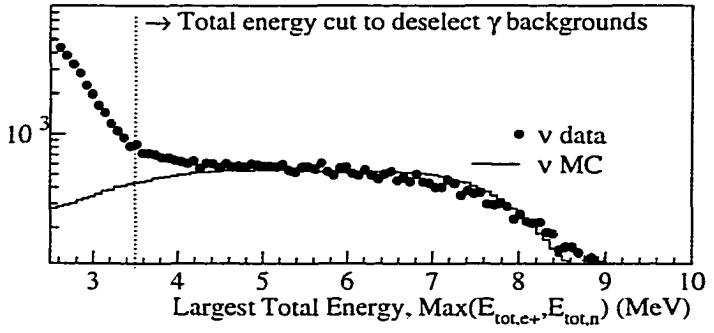

Figure 4.2: The triggers resulting from random coincidences of natural γ activity, which have generally lower energy, are strongly rejected by requiring that the total energy of either the prompt or delayed portion of a $\bar{\nu}_e$ candidate event have more than 3.5 MeV. Here the selection is displayed relative to Monte Carlo simulation of pure $\bar{\nu}_e$ signal and the data sample.

The muon deadtime can be further divided into two contributions: 150 μs of deadtime caused by each muon, which at an average rate of $R_\mu = 1990$ Hz left the detector live 74.2% of the time but varied with each data period's average veto rate; and muons which interrupted a neutrino event between the positron and the neutron capture, which estimated from the fit parameters of the Monte Carlo capture time left 92.5% of events uninterrupted. The total uncertainty in the calculation of detector deadtime is less than 1%. A precise tabulation of the data set including livetimes and efficiencies is shown in Table 4.1.

4.2 Backgrounds

Backgrounds can be separated into two types: correlated and uncorrelated. Uncorrelated background events are due to unrelated *triple* hits which randomly coincided in the time window allowed. Although most of the events collected were random coincidences, almost all of this type of background is removed by requiring at least one sub-event to have more than 3.5 MeV of reconstructed energy. These events do not have a time correlation between prompt and delayed sub-events (*inter-event* time)

Period	Reactor Status	Days of Data	Livetime	Eff_{MC}	$\text{Eff}_{\text{total}}$
1998 1	890 m Off	30.11	0.5201	0.1494	0.0777
1998 2	Full Power	30.37	0.5179	0.1500	0.0777
1999 1	Full Power	22.61	0.6528	0.1812	0.1183
1999 2	750 m Off	22.25	0.6512	0.1805	0.1175
1999 3	Full Power	30.43	0.6500	0.1802	0.1171
1999 4	Full Power	27.21	0.6513	0.1808	0.1178
1999 5	Full Power	27.73	0.6558	0.1806	0.1184
1999 6	890 m Off	29.63	0.6544	0.1808	0.1183
1999 7	Full Power	21.04	0.6422	0.1804	0.1159
2000 1	Full Power	29.74	0.6283	0.1779	0.1118
2000 2	Full Power	26.64	0.6297	0.1776	0.1118
2000 3	890 m Off	27.45	0.6284	0.1723	0.1083
2000 4	Full Power	27.88	0.6252	0.1721	0.1076
All		353.10	0.6213	0.1736	0.1084

Table 4.1: The data set as divided into approximately month long periods.

characteristic of neutron capture. They have instead a longer time correlation determined by the probability that the veto detected no muon between the prompt and delayed random *triples*. At a 2 kHz muon rate, this background is seen as a 500 μs tail under the normal neutron capture distribution. By looking at the *inter-event* times of the candidate $\bar{\nu}_e$ events at longer time scales, this background can be measured.

The *inter-event* time distribution after all neutrino selections (except the time correlation cut) is shown in Figure 4.3. The Monte Carlo for a pure neutron capture sample is empirically fitted to the sum of two exponentials. There are two time constants due to the inhomogeneity of the target: neutrons which remain in the scintillator have a 27 μs capture time, whereas those which enter the acrylic have a longer capture time due to the absence of Gd. The data *inter-event* time distribution is fitted to a function of three exponentials with fixed time constants consisting of the Monte Carlo fit τ's multiplied by a third time constant of 500 μs. Integrating the resulting 500 μs exponential of the uncorrelated background in the signal region gives an estimate of 4.3$\pm$0.2 events per day, or 10% of the $\bar{\nu}_e$ candidates being uncorrelated background events.

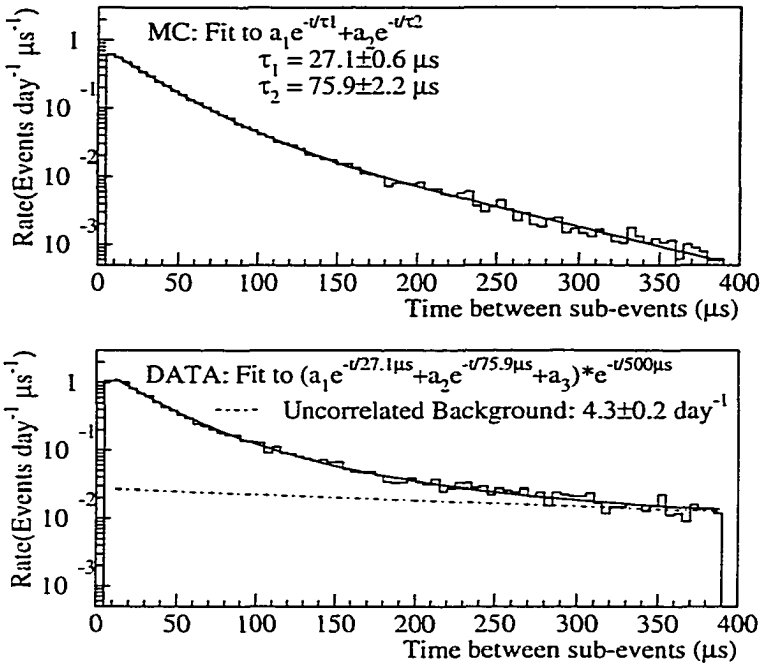

Figure 4.3: The time elapsed between the prompt and delayed portion of neutrino candidate events for Monte Carlo simulation and data. The Monte Carlo is fit to two exponentials. The data are fit to three exponentials, constrained to have the time structure of the simulation and an additional contribution due to uncorrelated background (dashed line).

To measure the uncorrelated background in smaller parts of the data set, the statistical accuracy of the three exponential fit method becomes unacceptably poor. A simpler method is therefore used in conjunction with the above fit. For *inter-event* times longer than 200 μs, the $\bar{\nu}_e$ candidates are dominated by uncorrelated background. The integrated number of candidates from 200–400 μs is scaled to estimate the number underneath the signal region (<200 μs). Using the scaling from the fit of the entire data set shown in Figure 4.3, the uncorrelated background was measured for each of the analysis periods as shown in Figure 4.4. The low energy of the uncorrelated background events makes the detector acceptance to these events more prone to instability, as evidenced by the non-statistical fluctuations seen. The variations observed in this background still represent less than 2% of the total $\bar{\nu}_e$ candidate rate,

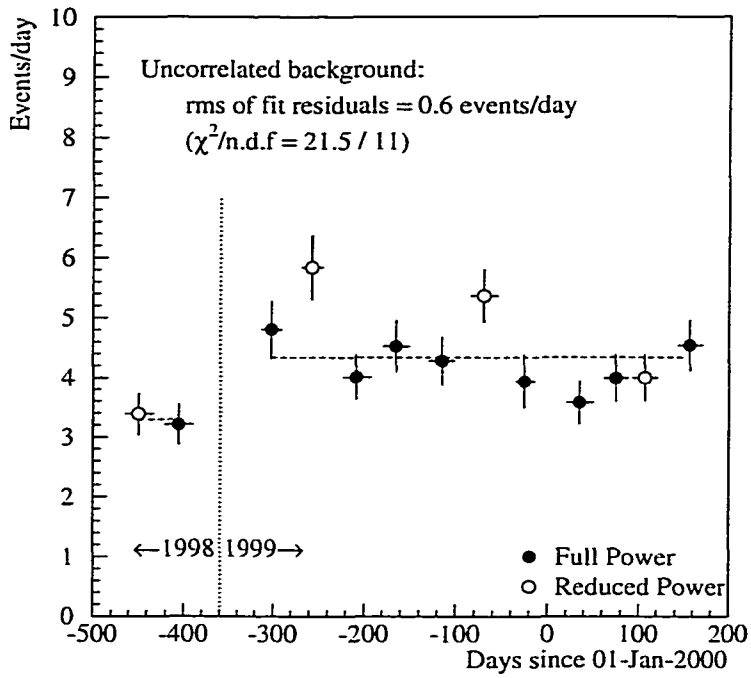

Figure 4.4: The uncorrelated background rate as measured in approximately month-long segments of the data. Overlaid is a fit to a constant within each year's data. While the spread in rates is non-statistical, the *rms* of the residuals to the fit is less than 2% of the total $\bar{\nu}_e$ candidate rate. (Horizontal error bars in this and following periodic rate plots represent the number of real days of data in the corresponding period.)

however, and can be accounted for in data analysis by statistical subtraction.

Correlated backgrounds have the neutron capture *inter-event* time structure of the $\bar{\nu}_e$ candidates. These events come mainly from cosmic muon induced fast neutrons from spallation or muon capture, as shown schematically in Figure 4.5. These fast neutrons can either (1) induce more neutrons via spallation, two of which can be captured in the detector with one capture mimicking a positron signature; or (2) they can cause proton recoil patterns in the central detector which appear as a positron signature and then get captured. Spallation neutrons originate from muons passing through the walls of the lab without being registered in the veto detector or from muons passing through the detector shielding undetected by the veto. Muon capture

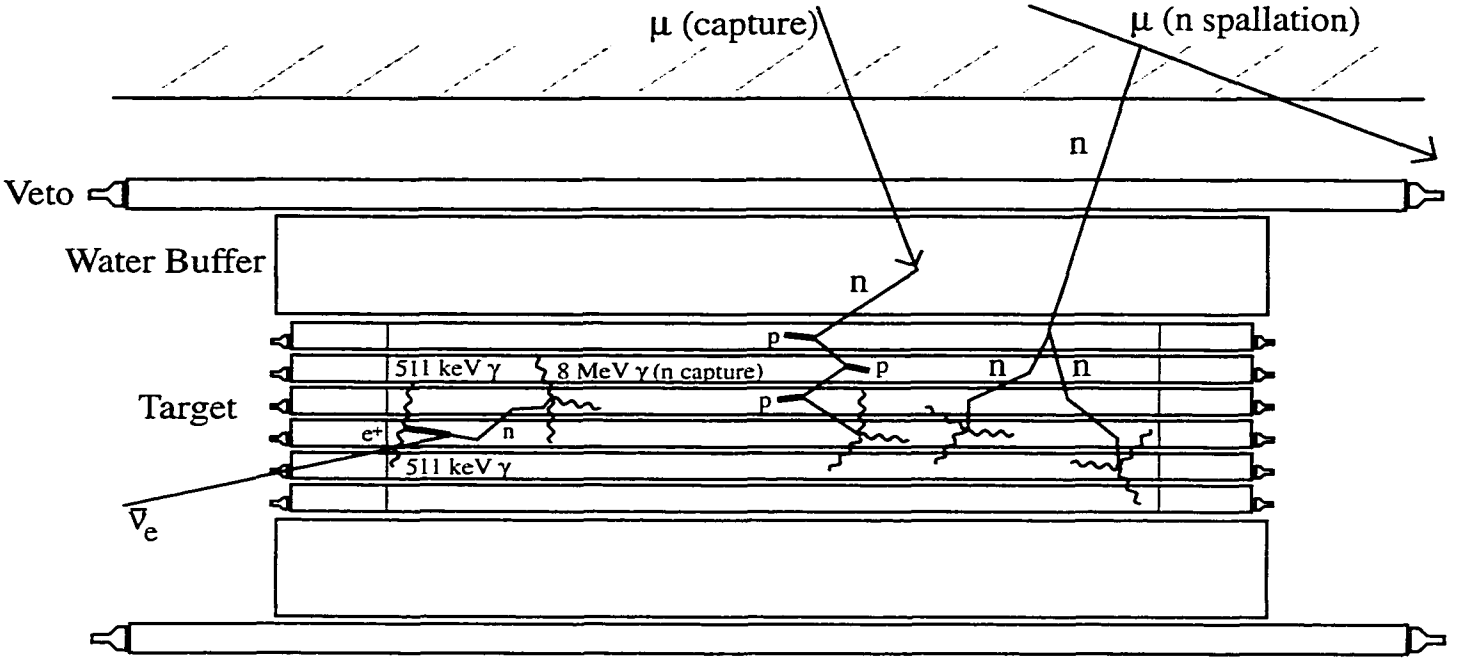

Figure 4.5: A schematic drawing of the detector showing a signal event and two examples of interactions which contribute to the correlated background. Neutrons which cause correlated backgrounds come mainly from muons undetected by the veto and captured in the water buffer (left), and spallation from muons in the lab walls (right). These neutrons in turn can either induce more neutrons via spallation (as shown at right) or cause three coincident proton recoils as they thermalize (as shown at left).

neutrons mainly originate from muons stopping in the water buffer without registering in the veto.

To illustrate some properties of correlated background, Figure 4.6 shows the time elapsed since the previous veto hit for $\bar{\nu}_e$ candidates, with all selection cuts applied except that of previous muon timing. This distribution is fit to a three exponential function analogous to that used for the *inter-event* time fits. The two time constants for neutron capture are not identical to those for $\bar{\nu}_e$ events, but tend to be smaller since after passage of a muon there are often more than one neutron in the detector to be captured. The third exponential time constant is again constrained to 500 μs as expected in a random sampling of events unrelated to the previous muon. Since

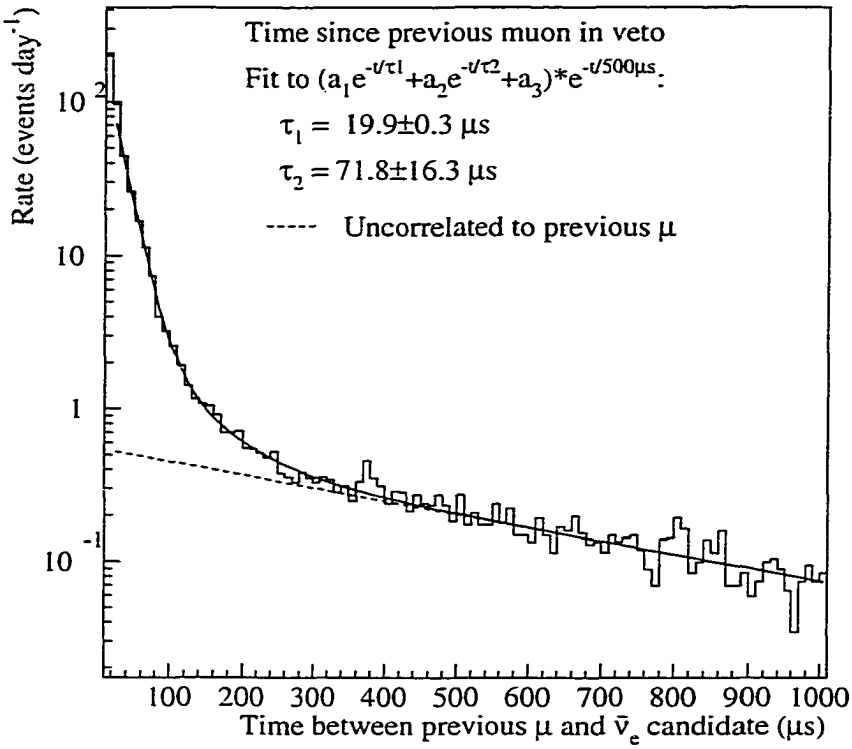

Figure 4.6: The time between the previous muon traversing the veto counter and a neutrino candidate event trigger. The fit is to a sum of three exponentials. The dashed portion denotes events uncorrelated to the previous muon, the rest are correlated background events induced by cosmic muons. The neutron capture time constants, shorter than the simulation prediction for $\bar{\nu}_e$ events, imply multiple neutron production.

at very short times there are other contributions such as muon decay, times less than 15 μs are excluded from the fit. Muon-induced-neutron backgrounds dominate the candidates in the first 150 μs after the previous tagged muons, motivating the selection cut on μ timing.

In order to show that the correlated background is constant in time, the previous muon time cut is disabled and a plot is made of the $\bar{\nu}_e$ candidate rate versus time, as shown in the top of Figure 4.7. The rate shown, corrected for detector deadtime and for the relative efficiency for each period, is reasonably stable, with around 1% variation over time. This shows that the detector efficiency for background remained stable over the data taking period.

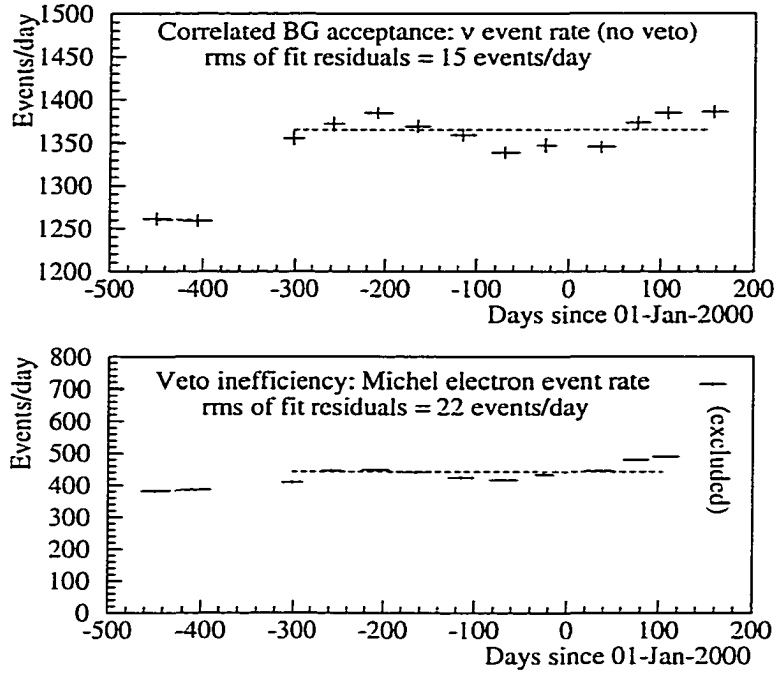

Figure 4.7: Top: The monthly rates of correlated background events (all $\bar{\nu}_e$ selections except the previous muon timing). Bottom: The rate of Michel electrons found in the detector from decaying stopped muons, a good indicator of relative muon veto efficiency. The veto apparently lost efficiency in the last data taking period. Both are fit to flat lines, separately for the 1998 data set due to trigger changes.

Aside from the detector efficiency for background, however, a loss of veto efficiency could also cause a fluctuation in background. (The rates fit in Figure 4.7 are with the muon timing selection disabled, and hence do not vary with veto inefficiency.) To track veto efficiency, we select events containing Michel electrons from stopped muons decaying in the detector target, which, due to the 10 μs trigger suppression following veto activity, occur only when the muon enters the target undetected in the veto. This measure of the veto *one hit missed* inefficiency was stable within 5% during the data taking with the exception of the last monthly period, when some loss of veto efficiency apparently occurred. The veto *two hits missed* inefficiency should have the the same stability as the two forms of inefficiency are correlated as mentioned earlier. Note that both the Michel rate and muon-induced uncorrelated background rates may

have some seasonal variation in them.

4.3 Neutron–$\bar{\nu}_e$ direction correlation

The neutrons produced in the inverse beta decays will have momenta slightly biased away from the source, whereas no correlation is expected for background. This effect is the consequence of momentum conservation which requires that the neutron be always emitted in the forward hemisphere with respect to the incoming $\bar{\nu}_e$. Such correlation has been observed already in the Gösgen experiment[43] and again at Chooz[61]. The theoretical treatment of the effect can be found in [55].

The signal to background ratio can be measured using this effect. The $\bar{\nu}_e$ source is to the left of the detector in Figure 2.2. The relative horizontal location (relative column in the target cell array) of neutron capture cascade cores versus positron ionizations for data and the simulation of the $\bar{\nu}_e$ signal are plotted in Figure 4.8. Defining the asymmetry $A_{\rm data} = \frac{R-L}{R+L}$ in terms of the number of neutrons captured one column away from the source R and one column toward the source L, a slight asymmetry 0.051 ± 0.015 is found in the data, at 3.4 σ significance. Using the Monte Carlo simulation which gives $A_{\rm MC} = 0.134$ to estimate the portion of the data consisting of $\bar{\nu}_e$ signal and assuming the background to be symmetric in this variable, an effective signal to noise ratio

$$\frac{S}{N} = \frac{A_{\rm data}}{A_{\rm MC} - A_{\rm data}} = 0.6^{+0.4}_{-0.3} \tag{4.1}$$

is found. This value agrees well with the ratio of 0.83 ± 0.02 found with the *swap* analysis method described below.

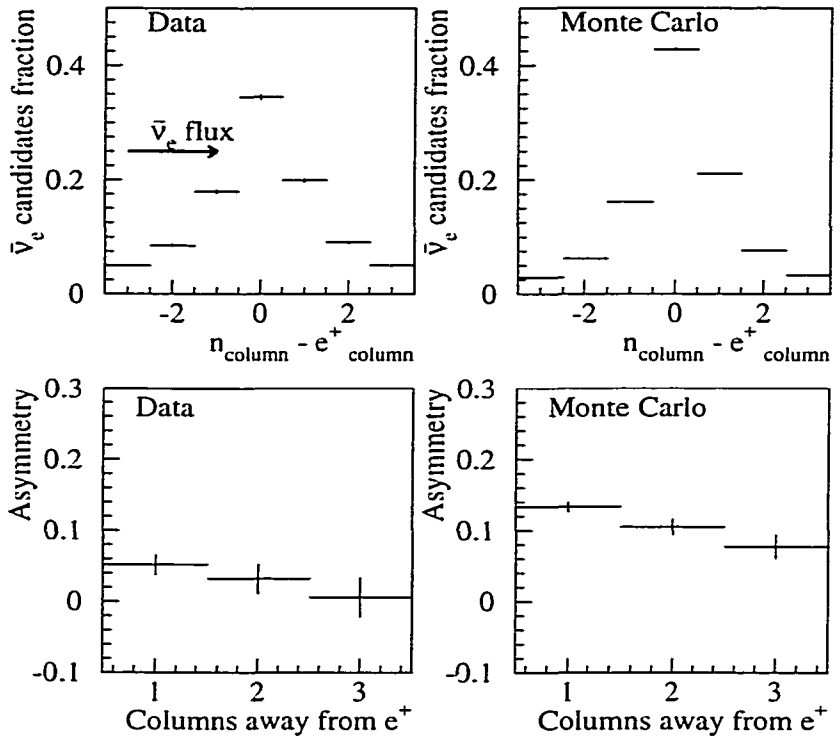

Figure 4.8: Top: The relative horizontal location of the of the neutron capture and positron ionization in the detector target. The kinematics of the inverse beta decay cause a slight bias in the momentum of the neutron from inverse beta decay, seen here as an asymmetry in the populations of neutrons captured one column away or towards the $\bar{\nu}_e$ source. Backgrounds should not exhibit such asymmetry. The $\bar{\nu}_e$ data set is on the left and the Monte Carlo simulation of pure $\bar{\nu}_e$ signal is on the right. Bottom: The asymmetry as a function of distance away from the e^+ column, for data and Monte Carlo. The largest value, that found for one column away or towards the reactors, is the asymmetry quoted in the text.

Chapter 5

Analysis

As discussed in the previous chapter, after all selection cuts there is still substantial background in the remaining data set. The correlated background, coming mainly from muon induced neutrons, is difficult to predict and subtract. The yield and spectrum of neutron spallation is a function of muon flux and energy, which in turn is a function of depth. While some measurements of fast neutron spectra and fluxes have been done in the past [62, 63], there is no model which can consistently predict the fast neutron production [64, 65]. In the following two methods used to extract the $\bar{\nu}_e$ signal from data are presented.

5.1 Analysis with the *on-off* method

The conceptually simpler method of subtracting background is to take advantage of periods of reduced power levels of the reactor source. Ideally all three reactors would be down at once allowing for a direct measurement of the background. However, in practice only one of the three Palo Verde reactors was refueled at any given time. These reduced power periods occurred twice annually for about a month. The data set for 1998 is treated independently, subtracting 1998 *off* from 1998 *on* and the three remaining reduced power data sets from the 1999-2000 full power data, since the efficiency of the detector changed between the two years. Combining the data across the change would require accurate knowledge of the relative background efficiency in

1998 and 1999-2000. By subtracting these data taken at reduced flux from the full flux data, a pure neutrino sample is retrieved albeit containing the statistical power of only a small portion of the potential data set: the subtraction is limited by the lower statistics of the refueling times and treats the $\bar{\nu}_e$ flux from the two reactors still at full power as background.

The primary concern arising from use of this method, aside from the loss of statistics, is guaranteeing that the background rates during the *on* and *off* periods were stable. Both correlated and uncorrelated backgrounds were carefully tracked to ensure stability as discussed above. The last period of full power data was excluded from this analysis in light of the lower veto efficiency during that period. The contributions of the remaining fluctuations in veto efficiency to this analysis' uncertainty is discussed, along with the other systematic errors, further below.

The uncorrelated component of the background can be subtracted from the data directly since it is measurable with reasonable accuracy. This subtraction will increase the statistical uncertainty but reduces considerably the potential biases the fluctuations in uncorrelated background rates can introduce.

The total efficiency corrected $\bar{\nu}_e$ rates for the various full power and reduced power data sets are shown in Figure 5.1. The seven full power periods of 1999-2000 are combined into one average rate. One can see a clear reduction in the rates during the refueling periods.

After efficiency correction the four reduced flux rates are then subtracted from their respective full power rates to find observed neutrino interaction rates in the detector. The results are tabulated in Table 5.1 for both the case where the raw $\bar{\nu}_e$ candidate rates are used and where measured uncorrelated background is subtracted from the raw rates first. The agreement between the calculated and observed $\bar{\nu}_e$ interaction rates improves substantially when the uncorrelated background is subtracted, due to the upward fluctuation in this background during the second and third reduced flux periods. The ratios of observed to calculated interaction for both sets of results are plotted in Figure 5.2. The average ratio of observed to calculated rates reported here takes into account the correlation of the 1999-2000 full power measurement common to the latter three reduced flux subtractions. No significant deviation from the

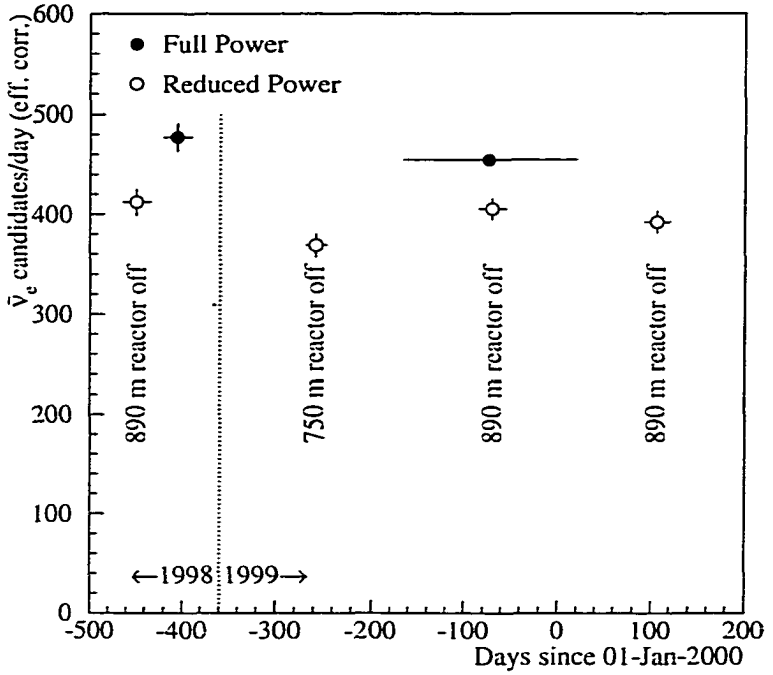

Figure 5.1: The neutrino rates after efficiency correction for the four refueling periods and the 1998 and 1999-2000 full power periods. Uncorrelated background has not been subtracted.

expected neutrino interaction rates is seen in either case.

In order to test the results for oscillation hypotheses in the two flavor $\Delta m^2 - \sin^2 2\theta$ plane, a χ^2 analysis is performed comparing the calculated $R_{\mathrm{calc},i}$ and observed $R_{\mathrm{obs},i}$ interaction rates for each refueling period i. The $\bar{\nu}_e$ interaction rates from the analysis accounting for uncorrelated background were used. At each point in the oscillation parameter plane, taking into account the changes in detector efficiency due to distortions of the neutrino spectrum, the quantity

$$\chi^2 = \sum_{i=1}^{2} \frac{(\alpha R_{\mathrm{calc},i} - R_{\mathrm{obs},i})^2}{\sigma_i^2} + \frac{(\alpha - 1)^2}{\sigma_{\mathrm{syst}}^2} \tag{5.1}$$

is computed, where α accounts for possible global normalization effects due to systematic uncertainties (discussed below) across both periods and σ_i is the statistical uncertainty in each period. The function is minimized with respect to α. The point in

Period	1998 1	1999 1	1999 2	2000 1
L (m)	890	750	890	890
ON N_{cand} (day^{-1})	37.08 ± 1.10	52.64 ± 0.53		
OFF N_{cand} (day^{-1})	32.04 ± 1.03	44.08 ± 1.41	47.96 ± 1.27	42.48 ± 1.24
ON-OFF N_{cand} (day^{-1})	5.03 ± 1.51	8.56 ± 1.51	4.68 ± 1.38	10.16 ± 1.35
R_{obs} (day^{-1})	64.8 ± 19.5	79.6 ± 12.8	49.3 ± 11.7	62.3 ± 12.4
R_{calc} (day^{-1})	63.4	89.0	63.4	64.2
Combined Obs/Calc	0.89 ± 0.09			
Uncorrelated Background Subtraction Analysis				
Unc. BG ON	3.21 ± 0.34	4.14 ± 0.16		
Unc. BG OFF	3.38 ± 0.35	5.84 ± 0.54	5.37 ± 0.45	4.00 ± 0.40
R_{obs} (day^{-1})	67.1 ± 20.5	93.5 ± 13.7	58.9 ± 12.4	63.5 ± 13.0
R_{calc} (day^{-1})	63.4	89.0	63.4	64.2
Combined Obs/Calc	1.00 ± 0.10			

Table 5.1: The results of the *on-off* analysis, showing candidate rates N and efficiency corrected $\bar{\nu}_e$ interaction rates R observed and expected. The results are also shown for the case where the uncorrelated backgrounds are accounted for in the candidate rates before the *on-off* subtraction. Errors are statistical only. Systematic errors are estimated to be 6.9%.

the physically allowed parameter space with the smallest chi-square χ^2_{best} was found, which represents the oscillation scenario best fit by the data. For the *on-off* analysis with uncorrelated background subtracted, the non-oscillation hypothesis has a χ^2/n.d.f. $= 0.27/3$. The analysis has a χ^2_{best}/n.d.f. of $0.27/1$ and α_{best} of 1.003, at the oscillation parameters $\sin^2 2\theta = 0.69$ and $\Delta m^2 = 3.0 \times 10^{-4}$. Clearly a non-oscillation hypothesis is acceptable. We then find a 90% confidence level (CL) acceptance region in the plane of two flavor mixing parameters.

The 90% confidence level acceptance region is defined according to the procedure suggested by Feldman and Cousins[66] by:

$$\Delta\chi^2 = \chi^2(\Delta m^2, \sin^2 2\theta) - \chi^2_{\text{best}} > \Delta\chi^2_{\text{crit}}(\Delta m^2, \sin^2 2\theta) \tag{5.2}$$

where $\chi^2(\Delta m^2, \sin^2 2\theta)$ is the fit quality minimized with respect to α at the current point in $\Delta m^2 - \sin^2 2\theta$ space and $\Delta\chi^2_{\text{crit}}$ is the CL χ^2 cutoff. Due to the sinusoidal

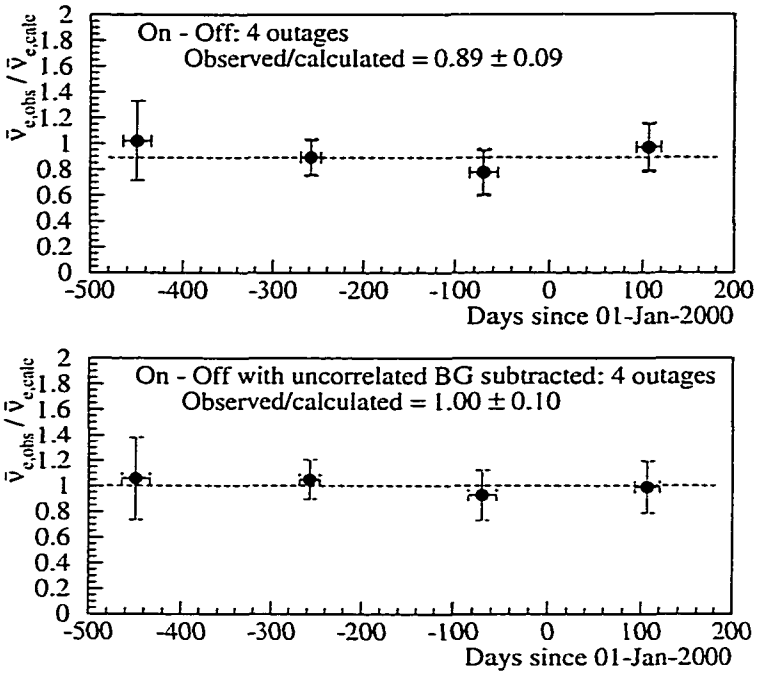

Figure 5.2: The ratio observed/calculated neutrino interaction rates for the four reduced flux periods of the data set for the raw data (top) and the data with uncorrelated background accounted for (bottom). Errors are statistical only. Systematic errors are estimated to be 6.9%.

dependence of the expected rates on the oscillation parameters and the presence of physically allowed boundaries to those parameters, the cutoff is not simply the $\Delta\chi^2$ one would analytically find for a three parameter minimization but has to be calculated for each point in the plane. To find the $\Delta\chi^2_{crit}$ for a point, the experiment is simulated 10^4 times under the assumption that the oscillation hypothesis represented by that point is true. For each simulated data set, a χ^2_{best} is extracted and a $\Delta\chi^2$ found for the point. These 10^4 $\Delta\chi^2$, the simulations' fit qualities to the hypothesis, are then ordered. The $\Delta\chi^2$ of which 90% of the simulations are a better fit is a 90% CL and therefore that oscillation hypothesis' $\Delta\chi^2_{crit}$. The region excluded by the analysis is shown in curve (a) of Figure 5.6.

The direct subtraction of the reduced power from the full power data set yields a spectrum of the $\bar{\nu}_e$ interactions observed in the detector from the reactor which is being

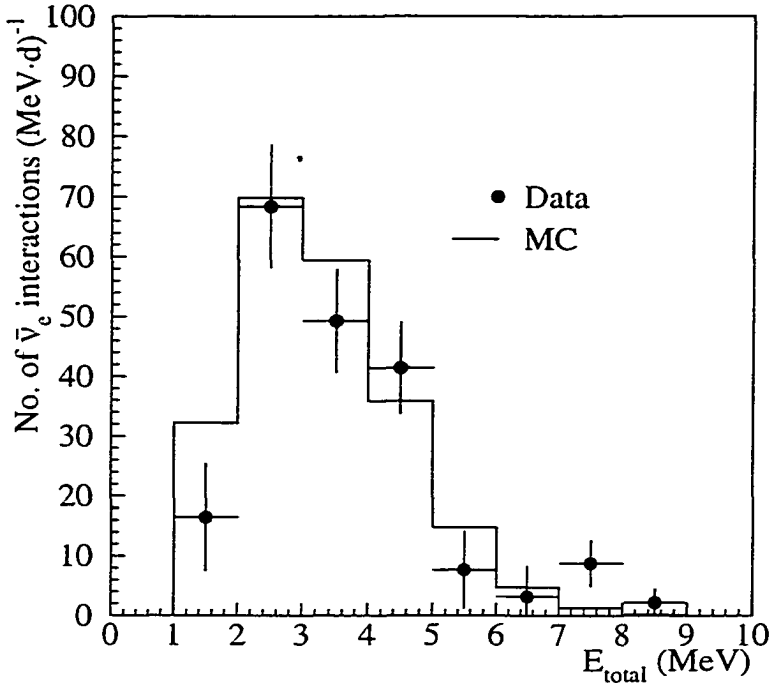

Figure 5.3: Total energy of the prompt sub-event ($\sim E_{\bar{\nu}_e} - 1.8$ MeV) in 1 MeV wide energy bins (observed and expected for no oscillation) from the *on-off* analysis, summed over the four refueling periods. Uncertainties are statistical only. Agreement is acceptable, with χ^2/n.d.f. = 10.6/8. Uncorrelated background is not subtracted.

refueled. Since the *on-off* analysis supports a no-oscillation hypothesis, the spectra of the four refueling periods can be added and the summed spectrum plotted, as seen in Figure 5.3 along with the Monte Carlo expectation. The observed neutrino spectra in principle provide additional information for testing oscillation hypotheses by showing potential spectral distortions. In practice, however, the systematic uncertainties as a function of neutrino energies are very difficult to estimate. For example, while the contribution of the uncorrelated background to the total observed $\bar{\nu}_e$ rate can be measured and accounted for as done in the analysis above, the unknown energy spectrum of this background precludes doing the same for a spectral analysis of the data.

5.2 Analysis with the *swap* method

A second analysis was used which has the potential of using the full statistical power of the neutrino data set by subtracting background directly. The method, discussed in detail in[67], takes advantage of the asymmetry of the prompt (positron) and delayed (neutron capture) sub-events of the neutrino signal. The data selection and trigger treat the two portions of the event identically with the exception of two cuts designed to isolate events with annihilation-like γ's in the prompt *triple*.

The candidates remaining after the selection cuts can be written as:

$$N = B_{\text{unc}} + B_{\text{nn}} + B_{\text{pn}} + S_{\nu} \tag{5.3}$$

where B_{unc}, B_{nn}, and B_{pn} are uncorrelated, two-neutron, and proton-recoil–neutron-capture backgrounds respectively; and S_{ν} is the neutrino signal. Applying the same neutrino cuts with the positron cuts reversed, or *swapped*, (such that the positron cuts are now applied to the delayed *triple*) gives:

$$N' = B_{\text{unc}} + B_{\text{nn}} + \epsilon_1 B_{\text{pn}} + \epsilon_2 S_{\nu} \tag{5.4}$$

Since the uncorrelated background and two neutron capture backgrounds are symmetric under exchange of the prompt and delayed *triples*, their efficiencies with the reversed cuts remain the same. The parameters ϵ_1 and ϵ_2 denote the relative efficiency change for proton recoils and neutrino signal under the swap, respectively.

The positron cuts are highly efficient for positron annihilation events but have poor efficiency for neutron captures. The Monte Carlo simulation is used to estimate $\epsilon_2 = 0.159$. Subtracting (5.4) from (5.3) leaves the majority of the neutrino candidates and only proton recoil background:

$$N - N' = (1 - \epsilon_1) B_{\text{pn}} + (1 - \epsilon_2) S_{\nu}. \tag{5.5}$$

The neutrons which cause the proton recoil background are created either by muon capture or spallation in the laboratory walls, or by muons entering the veto counter

undetected. Each of these cases was investigated for possible contributions to $B_{\rm pn}$. It turns out that the dominant contribution at Palo Verde is muon spallation in the lab walls.

To estimate $(1-\epsilon_{\rm l})B_{\rm pn}$ due to muon spallation, it is noted that the proton recoil spectrum extends beyond 10 MeV, well above the positron energies of the neutrino signal and other sources of background. These measured high energy events can be used to normalize the $B_{\rm pn}$ background in the signal using the Monte Carlo ratio:

$$r \equiv \frac{B_{\rm pn}^{\rm MC}({\rm E_{l,e^+}} < 8{\rm MeV})}{B_{\rm pn}^{\rm MC}({\rm E_{l,e^+}} > 10{\rm MeV})}, \tag{5.6}$$

where $B_{\rm pn}^{\rm MC}({\rm E_{l,e^+}} < 8{\rm MeV})$ is the fraction of simulated $B_{\rm pn}$ events passing the normal $\bar{\nu}_{\rm e}$ selections, and $B_{\rm pn}^{\rm MC}({\rm E_{l,e^+}} > 10{\rm MeV})$ is the fraction of simulated events in the high energy background region, i.e. where the most energetic prompt hit $E_{\rm l,e^+}$ is greater than 10 MeV. Multiplying the ratio r by the measured high energy proton recoil rate gives the $B_{\rm pn}$ background contribution:

$$B_{\rm pn} = r B_{\rm pn}^{\rm data}({\rm E_{l,e^+}} > 10{\rm MeV}). \tag{5.7}$$

To illustrate, Figure 5.4 shows the prompt most energetic cell spectra for the simulated $\bar{\nu}_{\rm e}$ signal and for the $\bar{\nu}_{\rm e}$ candidates in the data after all neutrino selections except the $E_{\rm l,e^+} < 8$ MeV cut. The data sample has a significant high energy tail which is the proton recoil background. This figure also has a simulation of proton recoils from fast neutrons which is used in estimating r as described below. Fast neutrons from muon capture can also contribute to r, and a correction is applied to account for this contribution, as discussed below.

The spectrum of the fast neutrons from muon spallation is not well understood and a wide range of possibilities are found in the literature. However, such spectrum can be decoupled somewhat from the resulting proton recoil spectrum. Among a range of spectra we conservatively chose four to sample from: the two extreme cases $E^{-0.5}$ and $E^{-2.0}$, a spectrum reported by the KARMEN experiment, $e^{-E/39}$ [68], and one measured at accelerators for photo-nuclear interactions $E^{-1.86}$ [69]. In the first

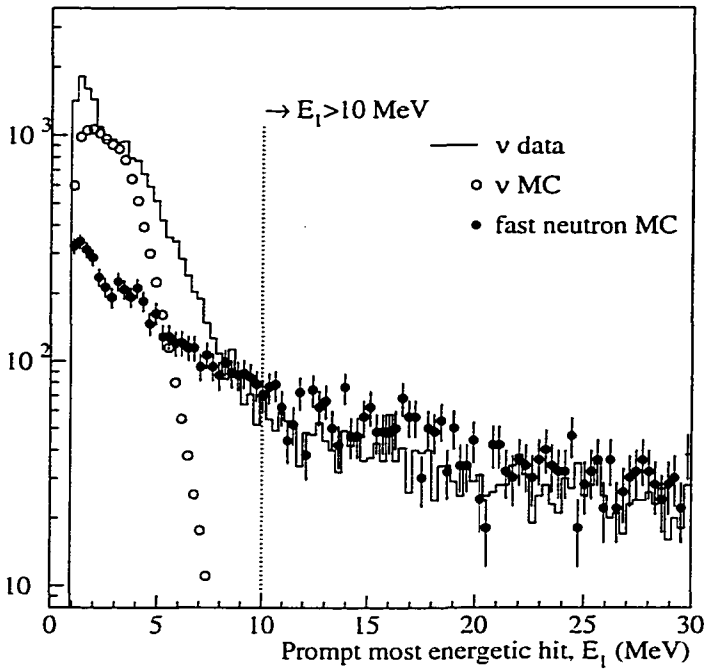

Figure 5.4: The most energetic prompt hit spectrum for the $\bar{\nu}_e$ Monte Carlo, the data, and a simulation of fast neutron backgrounds. We use the high energy range of this spectrum to estimate the proton recoil component of the background due to fast neutrons from muon spallation.

three cases we assume the neutrons are emitted isotropically.

In the latter spectrum, $E^{-1.86}$, we computed the neutron spectrum by producing cosmic muons in the energy range 0.01 MeV$< E_\mu <$500 GeV according to the proper energy and angular distributions [70]. We then generated real bremsstrahlung γ's according to the distribution $1/E_\gamma$ in the range 10 MeV$< E_\gamma < E_\mu$. Neutrons are then produced from photo-nuclear processes with a spectral shape $E^{-1.86}$ and an angular distribution from [71] in the energy range 10 MeV$< E_n < E_\gamma$.

To investigate the contribution to B_{pn} from muon spallation in the lab walls, we generated neutrons in a one meter shell of the lab walls with a lower energy cutoff of 10 MeV. Neutrons produced deeper in the lab walls or with less energy are completely absorbed before reaching the detector target. We first chose a tracking limit in the water buffer of 10 keV to save computation time. Neutrons in the target were tracked

until completely thermalized. The resulting ϵ_1 and r for the four simulated spectra are shown in the top half of Table 5.2. There is little dependence of either ϵ_1 or r on the fast neutron spectrum. Each simulation found that the $B_{\rm pn}$ events from muon spallation are almost symmetric under the exchange of prompt and delayed portions ($\epsilon_1 \sim 1$).

Also tabulated are the predicted number of high energy events $B_{\rm pn}^{\rm MC}(E > 10$ MeV) for each model. These were calculated by using the simulation efficiency for the neutrons to produce a high energy event and a neutron yield of $6 \times 10^5 \mu^{-1}{\rm g}^{-1}{\rm cm}^{-2}$, obtained by rescaling the measurements of neutron yield found in [62] to our depth of 32 m.w.e. These rates can be compared to the rate found in the data, $B_{\rm pn}^{\rm data}(E >$ 10 MeV)=11.8/day. Note that the simulated rates vary widely, an order of magnitude above and below that found in the data.

We performed the simulations again for three of the four models with neutron tracking in the water in addition to all other media having a thermal energy cutoff, shown in the lower part of Table 5.2. There is some effect on the average value for ϵ_1 but again the results are decoupled from the spectra. The change is apparently due to almost thermalized neutrons exiting and then re-entering the target. The softest spectrum, E^{-2} was not resimulated due to computation constraints. We take as the value to be used in the analysis the average of the three results, $(1-\epsilon_1)r = -0.02 \pm 0.05$, implying that this source of background is almost completely eliminated in the swap analysis subtraction.

The same calculation was performed for neutrons created in the passive detector shielding by untagged muons. In spite of a substantially higher efficiency for these neutrons to create backgrounds, we only expect 1600 muon spallation neutrons per day to be created by through-going untagged muons. The expected $B_{\rm pn}$ from these neutrons is therefore much smaller, being only a couple of percent of that from the walls.

The yield and spectrum of neutrons from muon capture are reasonably well understood. Since these neutrons tend to be lower in energy $E \lesssim M_\mu \sim 100$ MeV, only those created in the vicinity of the detector have any efficiency for creating background. Knowing the veto inefficiency to miss stopping muons ($4 \pm 1\%$), the capture

$E_{\text{n,spall.}}$ Model	ϵ_1	r	$(1-\epsilon_1)r$	$B_{\text{pn}}^{\text{MC}}(\text{d}^{-1})$ $E > 10$ MeV
Tracking cutoff in water 10 keV:				
$E^{-0.5}$	1.16 ± 0.07	0.69 ± 0.04	-0.11 ± 0.05	155
$E^{-2.0}$	1.20 ± 0.11	0.67 ± 0.07	-0.13 ± 0.07	1.7
$e^{-E/39}$	1.06 ± 0.07	0.77 ± 0.05	-0.05 ± 0.06	17
$E^{-1.86}$	1.15 ± 0.06	0.76 ± 0.04	-0.11 ± 0.04	32
Tracking cutoff in water 10^{-5} eV:				
$E^{-0.5}$	1.08 ± 0.06	0.65 ± 0.06	-0.05 ± 0.04	267
$E^{-2.0}$	–	–	–	–
$e^{-E/39}$	0.94 ± 0.04	0.86 ± 0.04	0.05 ± 0.04	24
$E^{-1.86}$	1.07 ± 0.04	0.71 ± 0.06	-0.05 ± 0.04	26
Average			-0.02 ± 0.05	

Table 5.2: The *swap* analysis correction due to backgrounds from fast neutrons produced by spallation in the laboratory walls for various possible neutron energy spectra. The top results are for simulations with a tracking cutoff in the detector water buffer of 10 keV. The bottom set of results are from simulations with neutron tracking to thermal energies in all media. $E^{-2.0}$ was not simulated with this case due to the large computation time required. The rightmost column shows the expected $B_{\text{pn}}^{\text{MC}}(\text{E}_{\text{l,e+}} > 10\text{MeV})$ high energy proton recoils per day predicted by the simulation, to compare with 11.8 events per day in the data.

rate in water surrounding the detector and its contribution to the background can be estimated using Monte Carlo simulation, as detailed in Table 5.3. These neutrons also create a small contribution to $B_{\text{pn}}^{\text{MC}}(\text{E}_{\text{l,e+}} > 10\text{MeV})$ which modifies the r found above by $\sim$10%. Overall this proton recoil background appears to be symmetric as well, $\epsilon_1 = 0.74 \pm 0.04$, meaning that the subtraction also strongly rejects this background. We scale the muon capture background correction $(1-\epsilon)B_{pn,cap.}$ by the observed relative Michel rate to account for varying veto efficiency for each data period. The uncertainty of the residual background $(1-\epsilon_1)B_{\text{pn}}$ correction from spallation and capture contribute systematic errors of 3.2 and 1.3% respectively to the *swap* analysis observed neutrino rates.

The results of this analysis are summarized in Table 5.4. The total $\bar{\nu}_{\text{e}}$ rates observed and the ratio of observed to expected $\bar{\nu}_{\text{e}}$ interactions returned by the analysis

Step	Quantity	Uncertainty (%)
Stopped μ Rate	150 ± 50 Hz	30%
Veto Inefficiency	4 ± 1 %	25%
μ^-/μ_{total}	0.43	~few%
Capture on ^{16}O	18%	~few%
n per capture	0.98	~few%
Total capture n/day	$39k \pm 16k$ n/day	40%
ϵ	0.74 ± 0.04	6%
$B_{pn,cap.}$	2.8 ± 0.1/(39k n)	4%
$(1-\epsilon)B_{pn,cap.}$	0.73 ± 0.07/(39k n)	17%
Total correction: (livetime corrected)	0.7 ± 0.35	50%

Table 5.3: The details of calculation of the muon capture background correction and its associated uncertainty.

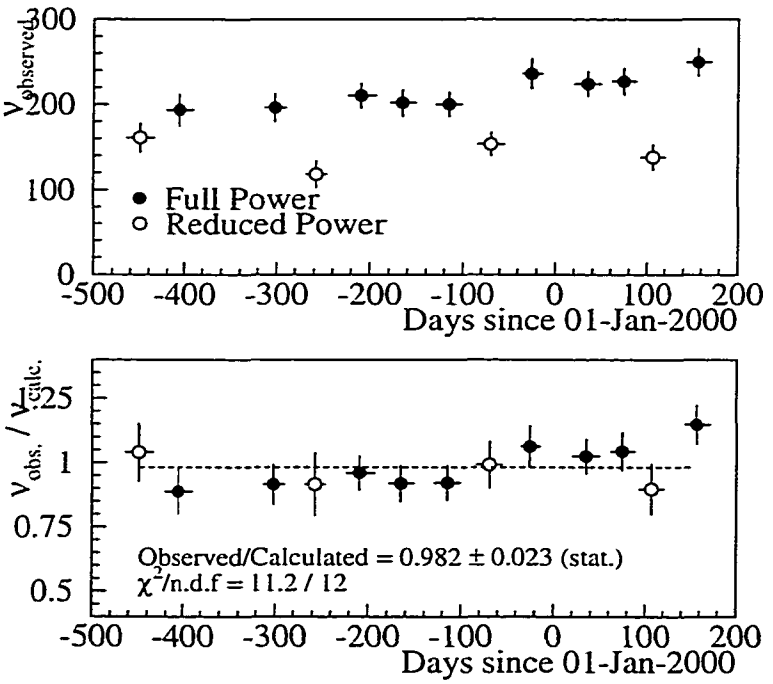

Figure 5.5: The total $\bar{\nu}_e$ rate and ratio of observed/calculated $\bar{\nu}_e$ interaction rates from the *swap* analysis for the 13 data periods.

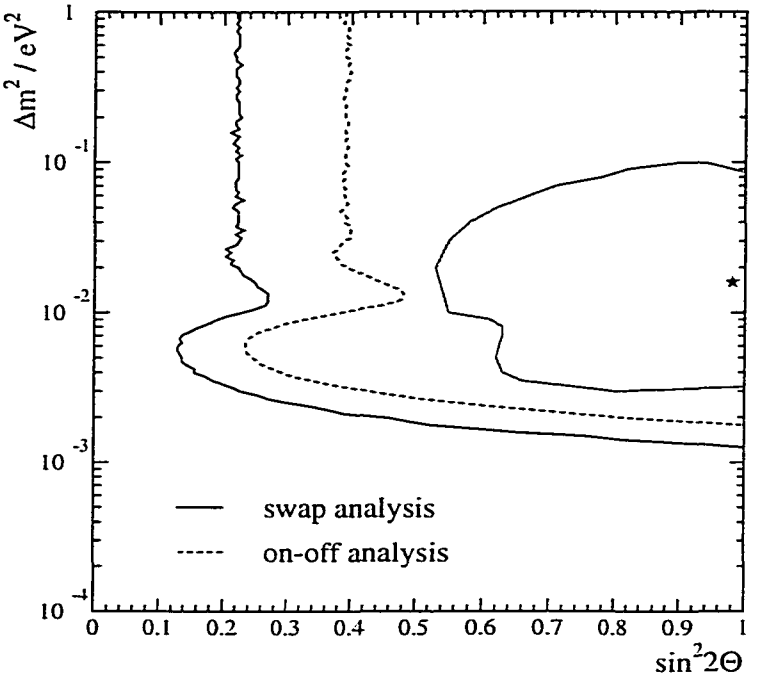

Figure 5.6: 90% CL limits of oscillation parameters for (a) the *on-off* analysis and (b) the *swap* analysis. The Kamiokande ν_e–ν_μ atmospheric neutrino 90% CL (dashed line) and best fit (star) are also shown.

are plotted in Figure 5.5. Overall $\frac{R_{obs.}}{R_{calc}} = 0.982 \pm 0.023(\text{stat.}) \pm 0.053(\text{sys.})$. The signal to noise ratio in the data is 0.83 ± 0.02. The 90% CL region for this analysis follows the same χ^2 formula, Equation 5.1, as for the *on-off* analysis. We find for the non-oscillation hypothesis $\chi^2/\text{n.d.f.} = 11.3/12$. A best fit of the data to oscillation hypotheses was found at $\sin^2 2\theta = 0.04, \Delta m^2 = 10^{-2}$, with $\alpha_{\text{best}} = 1.001$ and $\chi^2_{\text{best}}/\text{n.d.f.} = 11.2/10$. Again, the non-oscillation hypothesis is acceptable.

The excluded region for this analysis is shown in Figure 5.6, curve (b).

5.3 Systematic uncertainties

The systematic uncertainties have three sources: the prediction of expected $\bar{\nu}_e$ interactions, the efficiency estimate, and uncertainties introduced by the analyses' treatment of backgrounds. The uncertainty in the calculation of expected $\bar{\nu}_e$ interaction rate,

2.1%, is detailed in Appendix A.

The efficiency uncertainty can be further subdivided into that arising from direct comparisons of Monte Carlo e^+ and neutron efficiency from calibration measurements and that arising from the selection cuts themselves. As seen above, the Monte Carlo is able very well to reproduce, on average and stably over time, the observed neutron and positron efficiencies. The calibration runs taken with the ^{22}Na source, when compared with Monte Carlo simulations, show overall agreement across all locations of better than 1% in the efficiency predictions. However, the run-by-run agreement was at a level of 11% for these calibrations. Since the ^{22}Na source is similar to the inverse beta decay signal with the e^+ close to detector threshold, the positron efficiency uncertainty over the entire $\bar{\nu}_e$ spectrum is actually much smaller since the mean energy of the positrons from inverse beta decay is well above our detector threshold. The Am-Be calibrations consistently had better than 2% average agreement of neutron detection efficiency between simulation and data irrespective of the details of how one selected the events within the *inter-event* time window, in spite of some unexplained anomalies in the spectrum of measured capture times. We therefore assign a 2% systematic each for e^+ and n efficiency estimates.

The uncertainty due to the selection cuts was estimated by varying the cuts to see the effect on the analysis $\frac{R_{obs}}{R_{calc}}$ result. Each cut was varied individually to see the contribution of this cut to the final uncertainty, and then all cuts were varied simultaneously to account for possible correlations among the contributions of the individual cuts. These random variations were constrained to be within a sphere in "cut space" whose boundary is defined by the limits of the variations of each of the nine individual cuts. The individual cut contributions are summarized along with the results of the random variation in Table 5.5.

The $\mathrm{Max}(E_{tot,e^+}, E_{tot,n}) < 3.5$ MeV cut was changed from an original nominal value of 3 MeV as a result of discovering that variations in this cut had a significant effect on the *on-off* result below 3.5 MeV. This effect is shown in Figure 5.7. The ratio of neutrinos observed to calculated drops precipitously when the cut is extended below 3.5 MeV. This is due to fluctuations in the uncorrelated backgrounds below

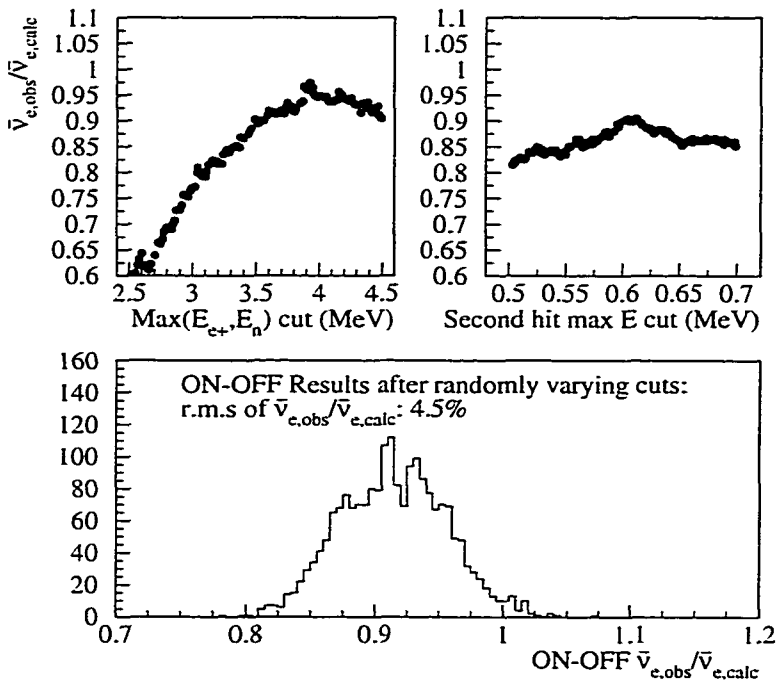

Figure 5.7: An example of the effect of varying selection cuts on the $\frac{R_{obs}}{R_{calc}}$ ratio in the *on-off* analysis. Upper left: sensitivity to the Max($E_{tot,e+}$, $E_{tot,n}$) cut, which was found to cause large changes in the result below a cut value of 3.5 MeV. Upper right: sensitivity to the positron cut, $E_{2,e+} < X$ MeV. Bottom: The $\frac{R_{obs}}{R_{calc}}$ of 1000 random variations on the cuts.

that energy. By comparison, the *swap* analysis is relatively immune to these fluctuations, as seen in Figure 5.8. The *on-off* systematic uncertainties quoted above are for the analysis without subtracting uncorrelated background. While subtracting uncorrelated background may reduce the systematics of some of the selection cuts, notably those rejecting uncorrelated background such as the Max($E_{tot,e+}$, $E_{tot,n}$) cut, uncertainties in how to vary the *inter-event* time cut and still coherently subtract the uncorrelated background prevent the investigation of the systematics of the background subtracted analysis. The *on-off* analysis is constrained by statistics, so this does not cause a great loss of discriminating power.

The *swap* method analysis has a somewhat smaller uncertainty due to selection cuts as the subtraction tends to cancel out systematics by cancelling background

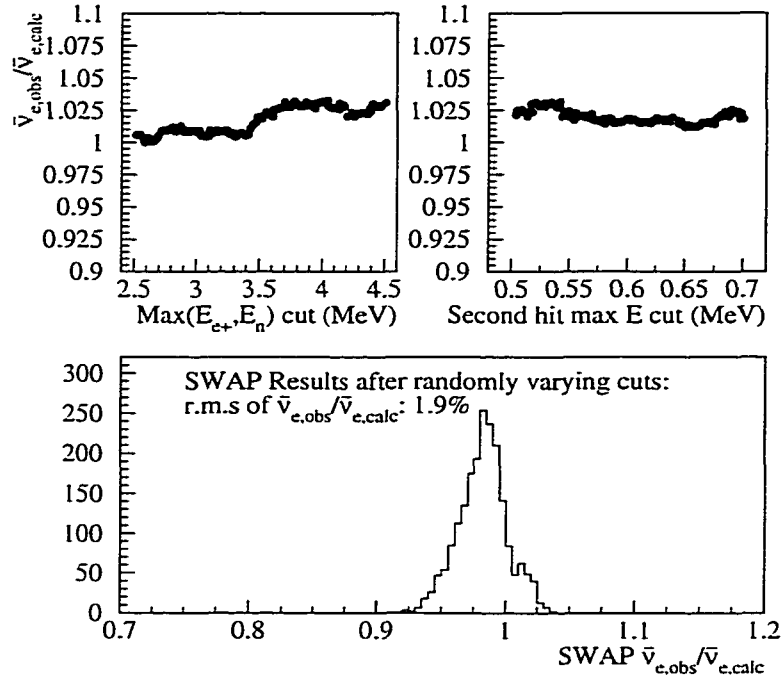

Figure 5.8: An example of the effect of varying selection cuts on the $\frac{R_{obs}}{R_{calc}}$ ratio in the *swap* analysis. Upper left: sensitivity to the $\mathrm{Max}(E_{\mathrm{tot,e+}}, E_{\mathrm{tot,n}})$ cut. Upper right: sensitivity to the positron cut, $E_{2,\mathrm{e+}} < X$ MeV. Bottom: The $\frac{R_{obs}}{R_{calc}}$ of 1000 random variations on the cuts. Note the relative immunity of this analysis to changes in these cuts compared to the *on-off* analysis.

fluctuations. However, the *swap* analysis uses a Monte Carlo estimate of the proton recoil background. The $(1-\epsilon_{\mathrm{l}})B_{\mathrm{pn}}$ corrections used have a 250% and 50% uncertainty for spallation and muon capture respectively (see Tables 5.2 and 5.3). These introduce a 3.2 and 1.3% systematic error respectively for a total of 3.5% in the final *swap* result.

The *on-off* analysis uses the assumption that all backgrounds are stable across the full power and reduced power periods being subtracted. The uncorrelated background can be measured and accounted for. We showed that the detector efficiency for correlated backgrounds was stable at a level of 1%. Using the spallation fast neutron background, $rB_{\mathrm{pn}}^{\mathrm{data}}(\mathrm{E_{l,e+}} > 10\mathrm{MeV}) = 8.6$ events/day, from the *swap* result, this leads to a small 0.1 event/day uncertainty in the *on-off* subtraction whose average signal is 8.0 events/day, introducing a 1.3% potential error in the result. The veto

inefficiency as monitored via the Michel electron rate varied by 5%. The *swap* analysis predicts 5 background events/day from stopping muons, implying a 0.3 events/day variation. This introduces a 3.7% uncertainty in the result. Together the systematics of the varying backgrounds cause a 3.9% error in the *on-off* analysis.

All of the systematic uncertainties are summarized in Table 5.6. The total systematic uncertainty is obtained by adding the individual errors in quadrature.

Period	Reactor Status	N	N'	$(1-\epsilon)B_{pn}$ spallation	$(1-\epsilon)B_{pn}$ μ capture	S_ν	BG ($B_{unc+nn+pn}$)	R_{obs}	R_{calc}
1998 1	890 m Off	32.0 ± 1.0	21.4 ± 0.8	-0.2	0.3	12.5 ± 1.4	19.5 ± 0.9	161 ± 18	155
1998 2	Full Power	37.1 ± 1.1	24.2 ± 0.9	-0.2	0.3	15.1 ± 1.5	22.0 ± 1.0	194 ± 19	218
1999 1	Full Power	52.5 ± 1.5	32.8 ± 1.2	-0.3	0.4	23.2 ± 2.0	29.3 ± 1.3	196 ± 17	214
1999 2	750 m Off	43.4 ± 1.4	31.5 ± 1.2	-0.2	0.5	13.9 ± 1.9	29.4 ± 1.3	118 ± 16	130
1999 3	Full Power	53.3 ± 1.3	32.4 ± 1.0	-0.2	0.5	24.6 ± 1.8	28.7 ± 1.2	210 ± 15	219
1999 4	Full Power	53.7 ± 1.4	33.4 ± 1.1	-0.3	0.4	23.8 ± 1.9	29.9 ± 1.2	202 ± 16	220
1999 5	Full Power	51.4 ± 1.4	31.2 ± 1.1	-0.3	0.4	23.7 ± 1.8	27.6 ± 1.1	200 ± 15	217
1999 6	890 m Off	48.0 ± 1.3	32.5 ± 1.0	-0.2	0.4	18.2 ± 1.7	29.8 ± 1.1	154 ± 14	155
1999 7	Full Power	54.7 ± 1.6	31.5 ± 1.2	-0.2	0.4	27.4 ± 2.1	27.3 ± 1.4	236 ± 18	222
2000 1	Full Power	51.5 ± 1.3	30.2 ± 1.0	-0.2	0.5	25.1 ± 1.7	26.4 ± 1.1	224 ± 15	219
2000 2	Full Power	52.0 ± 1.4	30.3 ± 1.1	-0.2	0.5	25.4 ± 1.8	26.5 ± 1.2	227 ± 16	218
2000 3	890 m Off	42.5 ± 1.2	29.6 ± 1.0	-0.2	0.5	15.0 ± 1.7	27.5 ± 1.1	138 ± 15	154
2000 4	Full Power	53.7 ± 1.4	30.6 ± 1.0	-0.3	0.7	26.9 ± 1.8	26.8 ± 1.2	250 ± 17	217

Table 5.4: The results of the *swap* analysis for the 13 periods. Errors are statistical only. Systematic errors are estimated to be 5.1%

Cut	Nominal Value	Variation	$\sigma_{Obs/Calc}$(%)	
			on-off	*swap*
e^+–n *inter-event* time	200	±50	3.5	1.0
e^+–n ΔColumn	3.5	±1.5	1.1	1.6
e^+–n ΔRow	2.5	±1.5	0.8	1.8
e^+–n Δz position	100	±20	2.0	0.6
Max(E_{tot,e^+}, $E_{tot,n}$) minimum	3.5	+1.0	1.9	0.3
E_1 minimum	1.0	±0.5	3.3	1.1
E_{2,e^+} maximum	0.6	±0.1	2.4	0.5
E_3 minimum	0.03	+0.02	2.9	1.0
$E_{tot,e^+} - E_{1,e^+}$ maximum	1.2	±0.2	1.7	1.2
Sum over cuts (quadrature):			7.0	3.3
With correlations:			4.5	1.9

Table 5.5: The systematic errors of the neutrino selection cuts. E_{tot}, E_1, E_2 and E_3 refer to the total energy and first, second, and third most energetic hits, while e^+ and n refer to the prompt and delayed banks, respectively. Each cut was varied within the range specified to see the effect on the result of the analysis. In order to take into account correlations among the cuts, all selections were simultaneously varied to get the result of the last line.

Error Source	*on-off* (%)	*swap* (%)
e^+ efficiency	2.0	2.0
n efficiency	2.0	2.0
$\bar{\nu}_e$ flux prediction	2.1	2.1
$\bar{\nu}_e$ selection cuts	4.5	1.9
B_{pn} estimate	N/A	3.5
Background variation	3.9%	N/A
Total	6.9	5.3

Table 5.6: Summary of the systematic uncertainties.

Chapter 6

Conclusion

In conclusion, the data taken from the Palo Verde experiment show no evidence for $\bar{\nu}_e \rightarrow \bar{\nu}_x$ oscillation. This result, along with the results reported by Chooz[44] and Super-Kamiokande[72], excludes two family $\nu_\mu \longleftrightarrow \nu_e$ mixing as being responsible for the atmospheric neutrino anomaly as originally reported by Kamiokande[73]. Later results of Super-Kamiokande, in particular data on the zenith angle distribution of muons and electrons, suggest that muon neutrinos ν_μ strongly mix with either ν_τ or with a fourth flavor of neutrino sterile to weak interaction.

One can combine this result with the Superkamiokande result, however, to gain extra insight into a more general three family mixing model. The most general approach would involve six unknown parameters, three mixing angles and a complex phase for the mixing matrix, and two independent mass differences. Analyses this general are premature; however, an intermediate approach consists of a simple generalization of the two flavor scenario, assuming that $m_3^2 \gg m_1^2, m_2^2$ (i.e. $\Delta m_{13}^2 = \Delta m_{23}^2 = \Delta m^2$, while $\Delta m_{12}^2 \simeq 0$). As mentioned in Section 1.3 above, these assumptions may come naturally from analyzing the solar and atmospheric neutrino results from a three family mixing perspective. In such a case the mixing angle θ_{12} becomes irrelevant and one is left with only three unknown quantities: $\Delta m^2, \theta_{13}$, and θ_{23}. With this parameterization the $\bar{\nu}_e$ disappearance is governed by

$$P(\bar{\nu}_e \rightarrow \bar{\nu}_x) = \sin^2 2\theta_{13} \sin^2 \frac{\Delta m^2 L}{4E_\nu} , \tag{6.1}$$

while the $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ oscillation in this scenario responsible for the atmospheric neutrino results, are described by

$$P(\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau) = \cos^4\theta_{13} \sin^2 2\theta_{23} \sin^2 \frac{\Delta m^2 L}{4E_\nu} \,. \tag{6.2}$$

A preliminary analysis of the atmospheric neutrino data based on these assumptions has been performed [74] and its results are shown in Figure 6.1 for the $\bar{\nu}_e$ disappearance channel. One can see that while the relevant region of the mass difference Δm^2 is determined by the atmospheric neutrino data, the mixing angle θ_{13} is not constrained very much. Here the reactor neutrino results play a decisive role.

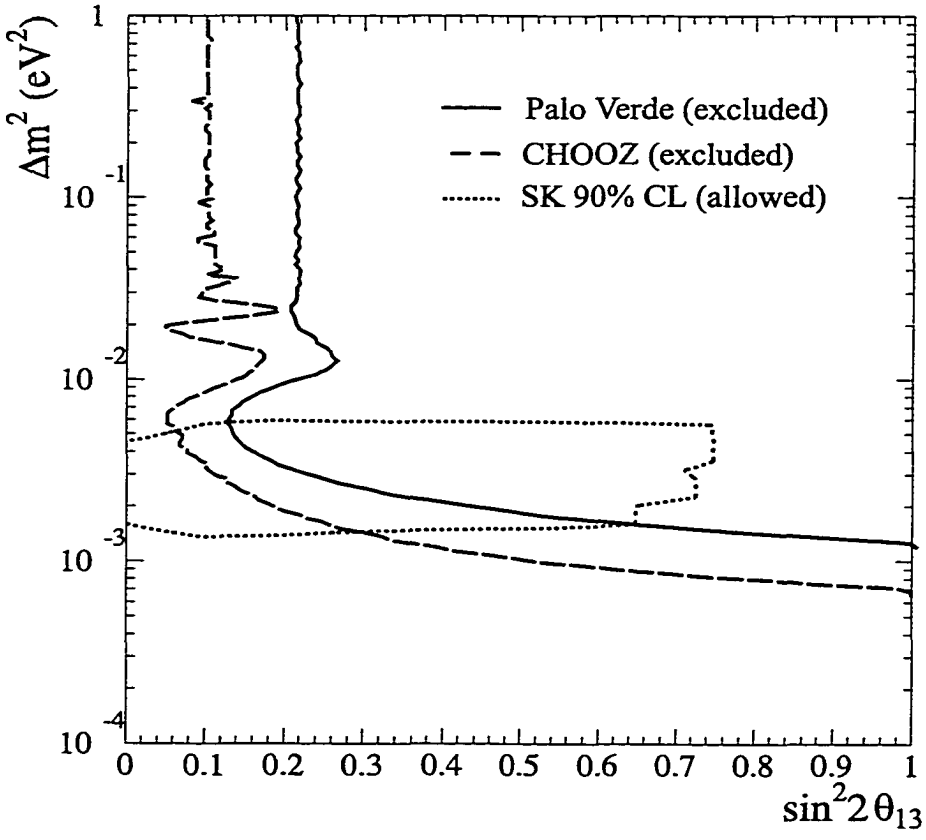

Figure 6.1: Exclusion plot showing the *allowed* region of θ_{13} and Δm^2 based on the Super-Kamiokande preliminary analysis (the region *inside* the curve). The region *excluded* by the neutrino reactor experiments are to the right of the corresponding curves.

The next generation of neutrino oscillation experiments, including the KamLAND reactor neutrino experiment, should provide enough information to answer many of the open questions in the nature of neutrino mass. The Palo Verde experiment stopped taking data on July 30, 2000 and is now decommissioned.

Appendix A

Expected Signal Estimation

In addition to understanding the absolute efficiency of the detector, neutrino disappearance experiments require understanding how many neutrinos from the source are expected to interact with the detector. Since it is unfeasible to place an identical detector close to the reactor, the neutrino flux emerging is estimated by modeling the core and extracting the total fission rates of the relevant Uranium and Plutonium isotopes. The resulting absolute fission rates are then converted to a neutrino spectral density. This spectrum, with cross sections for inverse beta decay and knowledge of the distances to the reactors, gives an estimation of the interactions per day expected in the target mass of the detector and the expected positron spectrum. The methods of these three steps above and the uncertainties in each are discussed below.

In every case an uncertainty is defined as the uncertainty introduced in the absolute number of neutrino interactions per day expected in the detector target, averaged over a full fuel cycle. During a fuel cycle the expected neutrino spectrum will evolve as the fuel is burnt. The uncertainties below could also affect not just the absolute number of interactions but the spectral shape observed as well, and is explored below.

A.1 Fission Rates in the Reactors

The three reactors at Palo Verde are identical 3.876 GW cores built by Combustion Engineering. Each core will generally be run at 100% power for around 500 days

and then be shut down for a 30–40 day refueling cycle. During a refueling cycle approximately one third of the fuel elements are replaced and the remaining elements are shuffled to optimize the heat production of the core. The state of each core is recorded hourly in a plant history file, among other places, which is how the experiment monitored the state of the plant and extracted the fission rates of the three units.

During a power cycle, the fuel's composition changes as Plutonium isotopes are bred and Uranium is depleted. The total fission rates of each isotope in the core is therefore a function of both thermal power of the reactor and the total burnup of the fuel in the core. The four isotopes whose fission and decay chains are the source of virtually all the energy and neutrinos produced by the reactor are ^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu. Other Pu isotopes, ^{240}Pu and ^{242}Pu, are bred in the core; but their fission contributes only a few parts in ten thousand to the total. Figure A.1 shows the evolution of the isotopes' total fission rates in a core under a normal full power fuel cycle. Each isotope produces a unique neutrino spectrum through its decay chain, so changing the relative fission rates of the isotopes will produce a significant change in the emitted neutrino spectrum and the number of events seen in the detector.

A proprietary core simulation code called ROCS (Reactor Operation and Control Simulator) was used to convert the state of a core into an estimation of the total fission rates for ^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu. The simulation requires five inputs from the plant: secondary calorimetric power (the most accurate of various power measurements); and the pressure, inlet temperature, flow rate, and ppm of Boron (used for thermal neutron absorption) in the primary cooling system. These inputs, in addition to a physical model of the core structure and initial fuel composition, are used to iteratively find a neutron flux solution for the core, which is then used to determine the fission rates. For simplicity ROCS models only a quarter of the core and assumes fourfold symmetry for the other three quadrants, since fuel is loaded symmetrically. A single flux solution will burn, or deplete, the fuel for a given time step before a new flux solution is sought.

Uncertainties in the model output can be separated into those resulting from inaccuracy of the inputs and those intrinsic to the algorithm used by the code itself. In

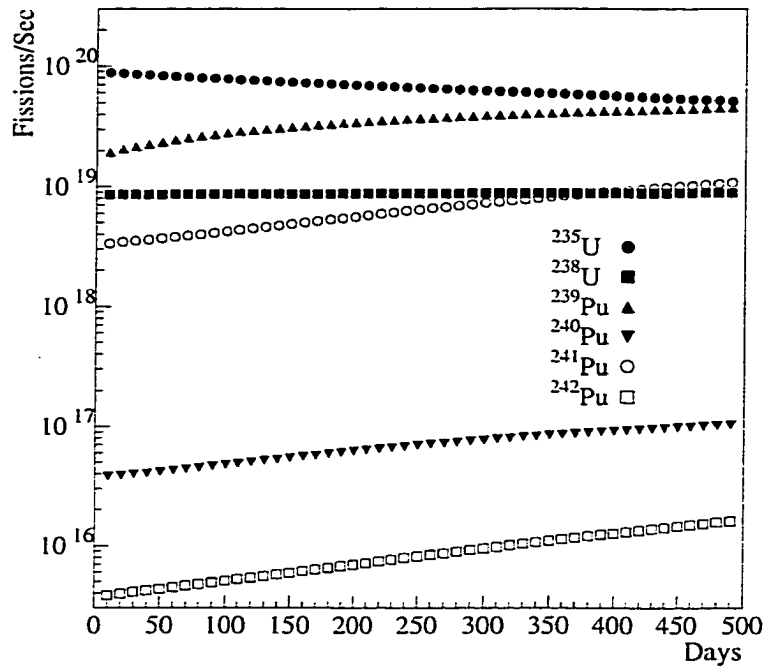

Figure A.1: Fission Rates for various isotopes during a typical 500 day fuel cycle for one of the Palo Verde reactors. The relative changes of rates as the fuel is burnt leads to significant changes in the neutrino spectrum.

order to estimate the uncertainty in the expected neutrino signal due to instrumentation errors of the plant, the ROCS simulation was run for a set of full 500 day cycles with each of the inputs varied independently. The change in expected neutrino signal was plotted versus input deviation, shown in Figure A.2. Correlations among inputs were not not explored. Uncertainties for each input quantity were then acquired [75, 76] from plant instrumentation analyses. Table A.1 summarizes the resulting contributions from plant instrumentation errors on the expected signal. Of the inputs only one, power, had significant effect on the core modeling. When cross-checked with Combustion Engineering [77] it was found that the ROCS algorithm relies less on the Boron input and relies more heavily on an effective criticality estimate given by the constraints of fuel composition and given power. The model uses the other three inputs mainly for the thermodynamics and density of the water in the primary cooling loop as they pertain to neutron transport.

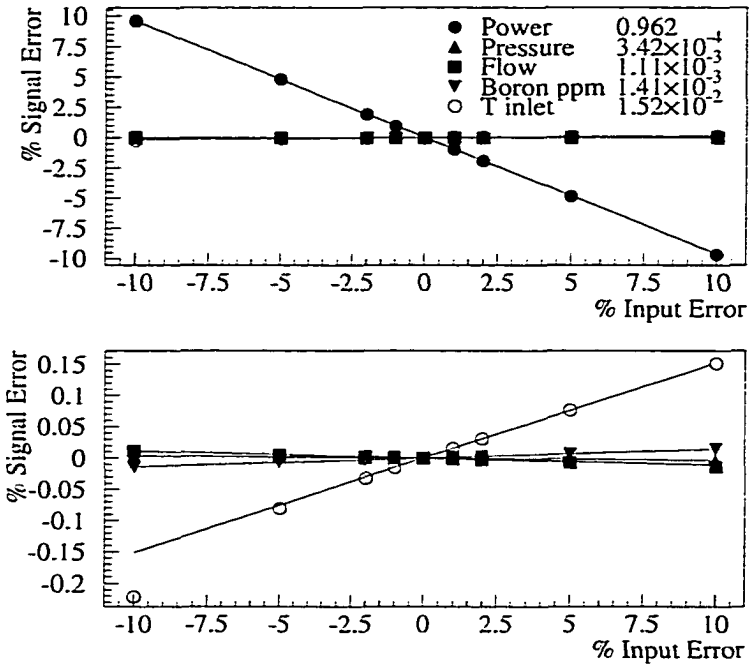

Figure A.2: The change in expected neutrino signal as a function of varying the five input to the core simulation (ROCS). The numbers in the key are the slopes of the fitted lines. Note that even varying the non-power inputs 10% has little effect on the output of the simulation.

The error intrinsic to the ROCS model was estimated using data supplied by Combustion Engineering [78]. Five fuel pellets from the Calvert Cliffs nuclear plant were analyzed for isotopic contents after being burned through several fuel cycles which had been modeled by ROCS. The five samples were taken from representative areas of the core to look at the accuracy of the model versus proximity to water, hot spots, or boundaries. The measured and predicted isotopic abundances were compared. All samples were measured to have a total exposure (burnup) within 1-2% of what the model predicted. Assuming that the discrepancy between an isotope's predicted and measured abundance is linearly related to the inaccuracy of the fission rate estimates, the fission rates' intrinsic error and therefore the error in expected

Input	ΔSignal/Δinput	Input Error (%)	Signal Error (%)
Power	.962	0.76	0.73
Pressure	3.42×10^{-4}	.4	$1.4 \times 10^{-}4$
Flow	1.11×10^{-3}	.21	$2.3 \times 10^{-}4$
PPM Boron	1.41×10^{-3}	2.0	$2.8 \times 10^{-}4$
Inlet Temp.	1.52×10^{-2}	.4	$6.1 \times 10^{-}3$

Table A.1: A summary of the errors introduced by uncertainty in the plant instrumentation inputs to the ROCS simulation. Note that power is the only input with a strong effect on the output. Correlations among inputs were not explored.

signal can be estimated. This assumption is valid as long as the neutron spectrum within the core is not strongly affected by the excess or lack of each isotope modeled. With the errors in composition found, of at most 2%, this is a valid assumption. The results, averaged over the five samples, are tabulated in Table A.1. This study was

Isotope	Pred.-Meas. Abund. (%)	ΔSignal/ΔFission	Signal Error (%)
^{235}U	-1.190	0.610	-.727
^{238}U	0.003	0.112	0.337×10^{-03}
^{239}Pu	1.915	0.209	0.399
^{241}Pu	1.305	0.068	0.089
Total: Linear Sum			-0.238
Total: Quadrature Sum			0.834

Table A.2: A summary of the ROCS intrinsic error estimate: The third column is the contribution of each isotope to the expected neutrino signal. The totals are for adding the isotopes' contribution either linearly, i.e. allowing cancellations, or in quadrature.

used by Combustion Engineering to improve the neutronics of the model, so these results can be viewed as a worst case of the current version of ROCS. Underburning one isotope is generally correlated with overburning another isotope since the model normalizes to the power input, so summing the results of the individual isotopes linearly is justified; however, summed in quadrature the resulting error is still not a dominant error.

Since the fission rates are being estimated in three units rather than one unit,

whether the above errors sum in quadrature or linearly, i.e. whether they are random or systematic, was explored. The intrinsic errors are considered systematic, since all three units use identical algorithms. The calculation [79] for the secondary calorimetry show that most of the uncertainty in this measurement comes from the secondary flow estimate, notably the diameter of the venturi pipe used. This is probably not correlated across units and was considered a random error.

To monitor the plant, we updated a database daily with each unit's total fission rate averaged throughout the day. Ideally one would use the state of the reactor from the previous day and model that with the current day's inputs for one 24 hour time step. The final state modeled would then be used the next day as the initial state, etc. Modeling of the core was only possible on plant computer systems. Due to plant computer security restrictions, a firewall, access was not granted to the plant's computers from outside and scripted automation of the above proved impossible. In practice, therefore, we modeled an entire full power cycle beginning with the initial state after refueling. The output of this gave a template for the cycle of full power fission rates versus EFPD (effective full power days). On any particular day, in order to extract fission rates, the total cycle EFPD to date was integrated from the cycle history, and the full power fission rates for that EFPD were modulated by the average power for that day. Since the units are run constantly at 100% power unless accidentally tripped or ramping to a refuelling, this method gives an accurate estimate of the fuel composition and therefore of fission rates. The other inputs also remain constant during normal cycle operation, with the exception of Boron which is constrained by the criticality of the reactor to follow a given concentration as fuel is burnt. If any significant deviation from 100% power occurs, the fission rates can be retroactively updated by re-modeling the cycle using the history files as input when access to the plant computers is available.

A.2 From Fission Rate to Neutrino Spectrum

The total fission rates of the four isotopes ^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu derived from modeling the reactors must be converted to a neutrino spectrum. For all but ^{238}U

direct measurements of the beta spectrum from fissioning of the isotopes by thermal neutrons exist [52, 53]. These were converted to neutrino spectra by fitting the observed beta spectra to a set of 30 hypothetical beta-branches. The neutrino spectrum of ^{238}U has not been measured due to its requiring more energetic neutrons to fission, and will therefore be discussed separately.

To more conveniently use the measured data, a fitted version of the data discussed in [80] was used. In this fit, the data for each isotope were fit to a three parameter function of energy:

$$\frac{dN_\nu}{dE_\nu} = \exp(a_0 + a_1 E_\nu + a_2 E_\nu^2)$$

where dN_ν/dE_ν is in units of $\bar{\nu_e}$ per MeV per fission. Figure A.2 shows the residuals of these fits. The error bars in the figure are the measurement errors at 90% CL. The measurement error and fit residuals were treated as separate sources of uncertainty in the expected signal. The results are summarized in Table A.2. The 90% errors of the data were used instead of converting to 68CL; however, the contributions from different isotopes were added in quadrature in obtaining the total. These uncertainties are the dominant ones in the final signal error. The errors tabulated may seem surprisingly small seeing the inaccuracy of the measurements above 5 MeV. Figure A.4, which shows the neutrino spectrum expected to interact in the detector, shows that the majority of the signal is in the range 2-5 MeV.

As mentioned above, no measured spectrum is available for ^{238}U, and here one must rely on calculation. While the isotope contributes only 7.7% of the thermal power to the reactor on average, it contributes 11.2% to the final neutrino signal. This is due to a slightly harder neutrino spectrum. The calculated spectrum was taken from ref [54]. The results there for ^{235}U and ^{241}Pu were compared to the ILL measurements [53, 52] and found to agree within approximately 10%. This was therefore used as an uncertainty in the spectrum of ^{238}U.

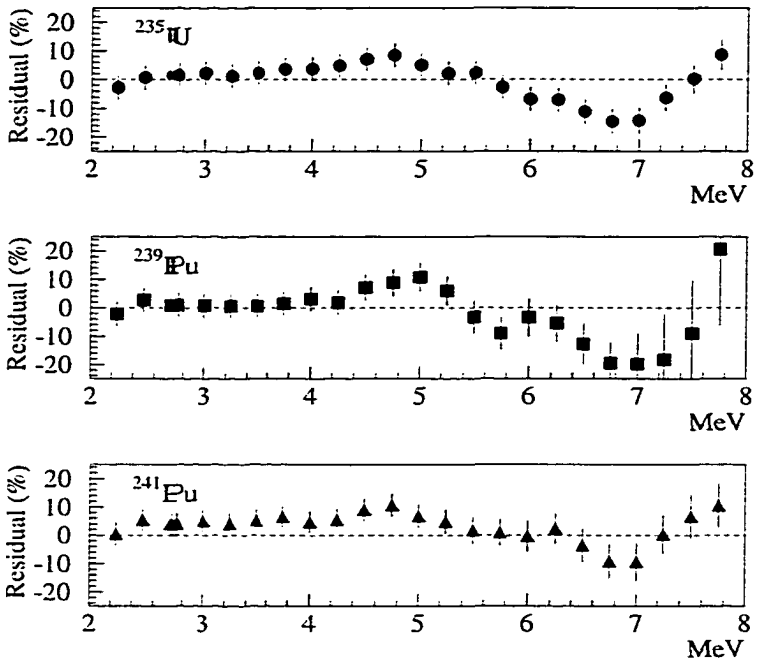

Figure A.3: The residuals (Fit-Data/Data) of the fit in [80] to the measured neutrino ILL spectrum [53, 52]. The error bars are the measurement errors.

A.3 From $\bar{\nu}_e$ Spectrum to Expected Signal

The neutrino spectrum emitted by the reactors must next be converted to an expected positron spectrum in the experiment. There are three contributions to this calculation: the distance to the core, the cross section for the reaction $p+\bar{\nu}_e \longrightarrow n+e^+$, and the number of protons in the target.

The distance to the reactors were surveyed by plant personnel and are known to within a meter, and therefore contribute nothing significant to the expected signal uncertainty. The cross section for inverse beta decay has been calculated in [81]. The lowest order cross section was modified for several effects such as neutron recoil, weak magnetism, and radiative corrections. The estimated accuracy of the resulting modified cross section is 1%.

Isotope	Contribution	ILL Err. (%)	Fit Error (%)	Total Error (%)
^{235}U	0.610	2.59	1.19	2.85
^{238}U	0.112	N/A	N/A	1.12
^{239}Pu	0.209	1.03	0.30	1.07
^{241}Pu	0.068	0.30	0.32	0.43
Total:				3.27

Table A.3: A summary of the errors from converting fission rates to a neutrino spectrum. Contribution is the portion of the expected signal due to the isotope. The two error types are from the empirical data and the fit to the data.

The number of protons in the target requires knowledge of the chemical composition and mass of both the scintillator which comprises its bulk and the acrylic tanks which hold the scintillator. During mixing and filling the mass of the scintillator was tracked to much better than 1%. Another source of uncertainty is the composition of the scintillator. Knowing the density of the various components in the mixture, one can try to reconstruct the expected density of the mixture and cross check with the measured scintillator density. These agreed within 0.7%. Finally, the number of protons in all components of the mixture are known except for the mineral oil, which accounts for 67% of the protons. Mineral oil is composed of hydrocarbons (C_nH_{2n+2}) of varying lengths; a particular oil is characterized by its average chain length. The percentage of the mass in protons is then $\frac{n+1}{7n+1}$ For the oil used, $< n >$ was given as 22.4 (14.8% protons). Varying this number by 10% results in a 0.5% uncertainty in the number of protons. Accounting for all these uncertainties gives the scintillator protons count at $8.79\times10^{29}\pm$ 1%.

The acrylic cells housing the scintillator add a significant number of protons to the target in that there is some efficiency for a neutrino interaction there to be detected. This efficiency is estimated at 26% and therefore one fourth of the uncertainty in the number of protons in the acrylic should be included in the total uncertainty of the calculation. Estimating from the dimensions of the cells the volume of the acrylic, and measuring the density of a sample of acrylic, there are found to be $1.66\times10^{29}\pm$ 5%. The uncertainty comes mainly from the density measurement. The above results in the total number of protons in the target to be $10.45\times10^{29}\pm$ 1.3%. Table A.3

summarizes the uncertainties of this section.

Source	Value	Error (%)
Unit Distances	887.5,750.4,888.5 ± 1.0 [m]	0.26
Target Effective Protons	$9.22\times10^{29}\pm 1.20\times10^{28}$	1.30
$\sigma(p+\nu \longrightarrow n+e^+)$	N/A	1.00
Total		1.66

Table A.4: The total uncertainty in the expected signal.

A.4 Summary of Signal Uncertainties and Spectral Uncertainties

Table A.4 lists all the uncertainties in deriving the expected neutrino signal in the detector based on the calculation presented above. The total expected rate of neutrino interactions is known to 3.8%.

Previous short baseline reactor neutrino experiments have experimentally measured the rate of neutrinos coming from reactor cores at Gösgen [82] and Bugey [83]. Both found agreement between modeled and observed spectra. Bugey, in particular, found excellent agreement both in shape (χ^2 9.23/11) and in absolute rate (0.99 observed/expected). The experiment used the same ILL measured neutrino spectra and a similar calculation of the ^{238}U spectrum. The final Bugey result is quoted as having a 1.4% total uncertainty. Since they use a comparable calculation to the one presented here, we can use this measurement to reduce the uncertainty in parts of the Palo Verde calculation.

Some possible differences remain between the Palo Verde and Bugey calculations, such as whether the ILL measured anti-neutrino spectra were fit or used as raw empirical data. However, these differences are minor. If we separate all of the Palo-Verde-specific uncertainties in the expected signal estimate, we can add these to the Bugey uncertainty. This is summarized in the right column of Table A.4. We find that the Palo Verde experiment has a 1.6% uncertainty associated with contributions

Error Source	Calculation (%)	With Bugey (%)
^{235}U	2.9	..
^{238}U	1.1	..
^{239}Pu	1.1	..
^{241}Pu	0.4	..
ν/MeV/Fission (total):	3.3	..
$\sigma(p + \nu \longrightarrow n + e^+)$	1.0	..
ROCS intrinsic error	0.3	..
Total in Common:	3.4	1.4
Reactor Power	0.7	0.7
Target protons	1.3	1.3
Distance to cores	0.3	0.3
Total error:	3.8	2.1

Table A.5: The systematic uncertainty in the expected neutrino signal. The left column does not use the Bugey result, whereas the right attributes the uncertainties common to both experiments to the quoted Bugey uncertainty of 1.4%.

unique to itself, and that the other contributions to the uncertainty are in common with Bugey's.

Finally, the uncertainty in the spectrum of neutrinos expected to be seen in the detector was calculated, including all sources listed in Table A.4. The result, shown in Figure A.4, is that there is no significant deformation of the spectrum to be expected due to the uncertainties. Both the actual spectrum and the bin by bin percentage uncertainties are plotted. The coarse granularity of the binning of empirical fission spectra errors are responsible for the apparent discontinuities in the bin errors. The effect of the uncertainties is to change the overall normalization of the spectrum seen. From this spectrum the expected positron spectrum can be calculated and compared with the data.

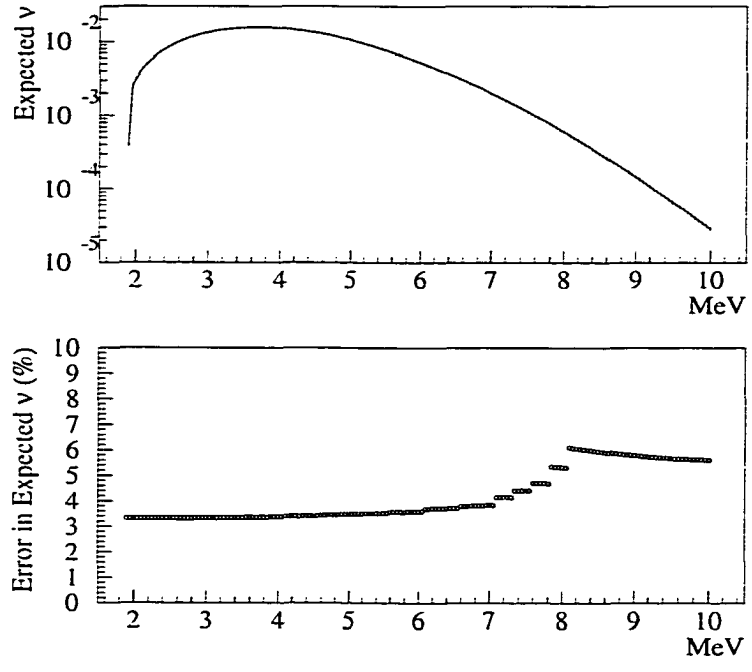

Figure A.4: Top: The expected spectrum of neutrinos which interact in the detector. The actual positron spectrum measured in the detector is slightly modified by the kinetic energy carried away by the inverse beta decay neutron, ~50 keV. Error bars (barely visible) are for all the uncertainties discussed in the text. Bottom: The percentage error in the above spectrum by bin. Jumps are due to the graininess of the empirical error bars in energy.

Appendix B

The Central Detector PMTs

B.1 Introduction

The low event rate and low background nature of the Palo Verde measurement placed many constraints on the characteristics of the signal region phototubes. The signals of interest ranged from 40 keV to upwards of 10 MeV deposited in the cells, which represented a dynamic range at the PMTs photocathode of 1 to 400 photoelectrons. The PMTs therefore needed to resolve single photoelectron signals and be reasonably linear over this range. The time of arrival of signals from the PMTs are used to spatially locate the events in the cells with a resolution of approximately 0.3 m. In order not to degrade this resolution the PMTs should introduce less than 1 ns uncertainty into signal arrival times. The front end electronics were designed around a tube gain of 5×10^7 and which gain the PMTs must achieve and maintain stable to within 1%. Finally, the PMTs dark current, afterpulsing and intrinsic radioactivity must not produce high false trigger rates or greatly increase the expected background rates.

Presented below are results of measurements performed on a pair of prototype PMTs from EMI Electron Tubes Inc. to verify that they met these requirements of the experiment. The measurements focused on the following properties: gain, including the effects of magnetic fields and temperature; single photoelectron peak to valley ratio; linearity; signal transit time; rise time; dark current and afterpulsing

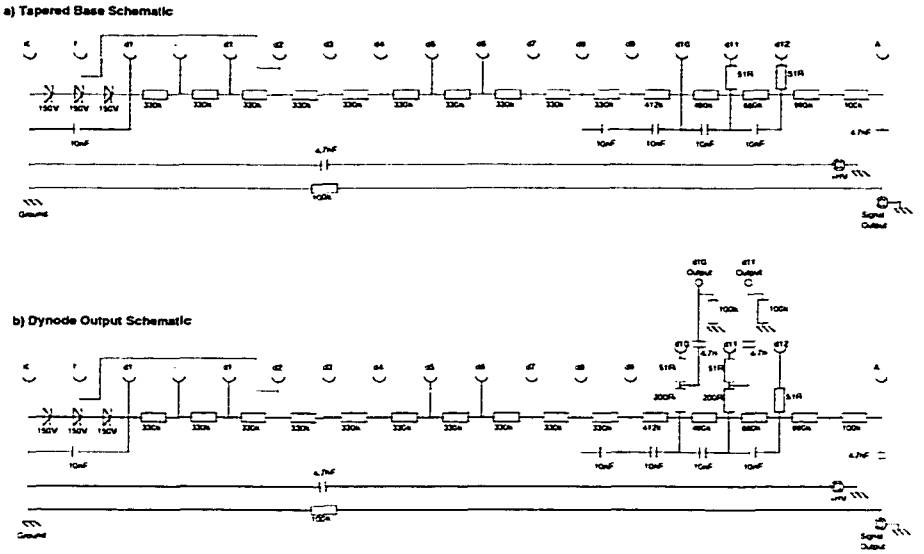

Figure B.1: a) The schematic of the tapered base. b) Modified version designed to allow a dynode signal output.

rates; dynode output characteristics; and radioactivity. From these tests acceptance criteria were designed for the central detector PMTs.

The tubes are 5 inch hemispherical front window tubes with an operating voltage in the range of 1500 to 2000 Volts. The tubes were tested with a tapered positive high voltage base supplied by EMI; the schematic is shown in Figure B.1. This same scheme was used for the central detector tubes in data taking with a slight modification to the dynode output; the 200Ω resistors were observed to cause ringing and were replaced with 100kΩ resistors. The high range signal was taken off of the 10th dynode.

B.2 Gain and single photoelectron properties

Since part of the regular calibration scheme for the detector was PMT gain monitoring using single photoelectrons, the separation of the one photoelectron peak from the lower amplitude noise was investigated. The absolute gain was then measured from the position of the single photoelectron peak. A small amplitude signal was produced by pulsing an LED and the resulting spectrum was measured. Due to the

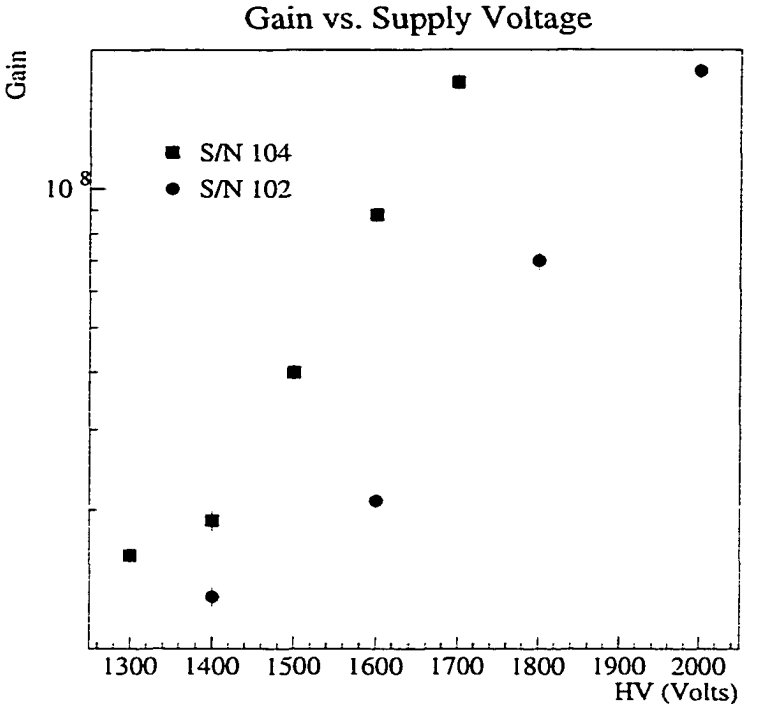

Figure B.2: The gain of the two prototypes for various supply voltages. The tubes were tested under identical conditions.

unknown frequency response of the an ORTEC shaping amplifier we estimate that this calibration introduced approximately 20% systematic error in the absolute gain, while relative gains could be measured to within 5% percent. Since relative rather than absolute gain measurements were required for the rest of our investigations, we did not pursue refining the calibration.

The spectra, a typical example of which is shown in Figure B.3, had discernible one and two photoelectron peaks with a single photoelectron peak to valley ratio of approximately 3 at a gain of 5×10^7. All of the tubes in the central detector achieved this gain at supply voltages between 1500 and 2000 V. The detector PMT nominal gain was later reduced by 20% for afterpulsing instabilities as discussed below. The two prototypes had significantly different gains, but both achieved the needed gain.

We then investigated the stability of the gain of the tubes. In particular we investigated two possible sources of drift in gain: power supply drift and the temperature dependence of the Zener diode voltages in the first stage. In order to stabilize the

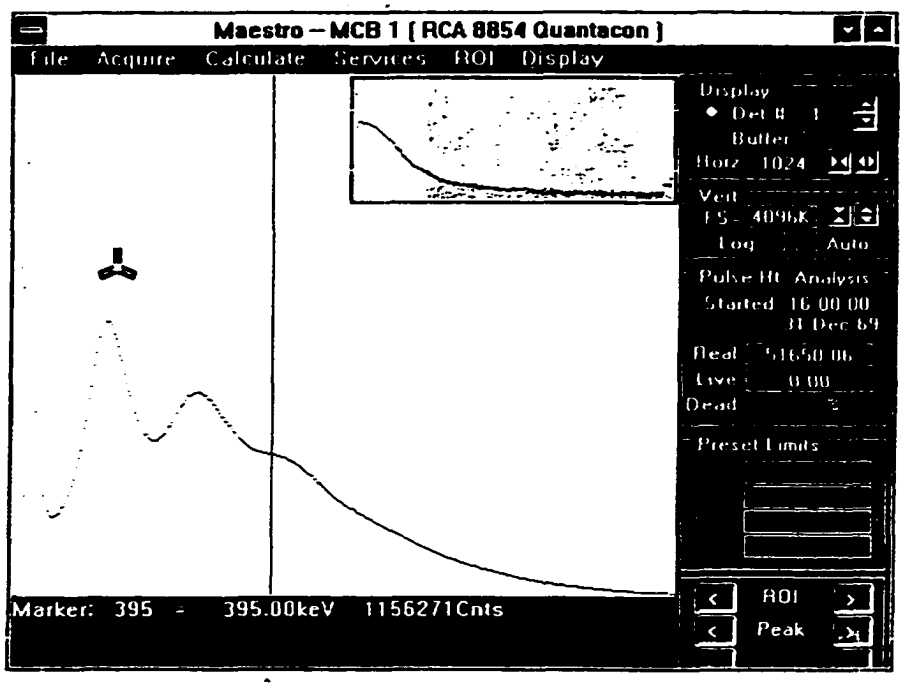

Figure B.3: A typical spectrum from a gain measurement. The first three photoelectron peaks can be seen.

gain to within 1% of 5×10^7 the slope of the gain curve demands the voltage be stable to within 3 Volts, or around 0.2%. This is within the specifications of the power supply used at Palo Verde. The Zeners' temperature dependence was measured to be $0.1\%C^{-1}$, or for the first stage of the base 0.45 VC^{-1}. The underground laboratory temperature fluctuates by no more than ± 3 C. To relate this fluctuation to gain drift, we measured the effect of first stage voltage on gain. The gain changed $0.6\%V^{-1}$, which, for the range of temperatures expected represents a gain drift of 0.8%.

We also investigated the effect of the Earth's magnetic field on the PMT gain, which can cause significant degradation of performance in larger tubes because of induced distortions in the path of electrons propagating from photocathode to first dynode. We measured the gain of the tube in several orientations at 1750 Volts. Two sets of orientations were considered, the first by varying the direction of the axis of symmetry of the tube, and the second by rotating about the axis of symmetry. These results are shown in Figure B.4. The gain of the tube varied by around 10% across

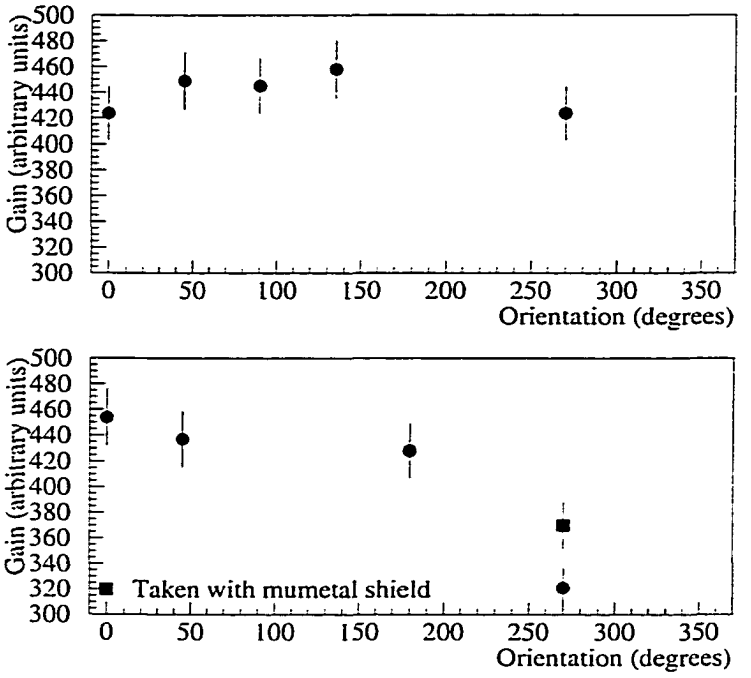

Figure B.4: Top) The gain across various orientations of the tube axis. Bottom) The gain as measured at several angles of rotation about the axis of the tube. One outlying point has been measured twice to show the effect of Mumetal shielding.

the range of orientations. One orientation in particular had a 30% drop in gain. This gain was remeasured with a 0.005 inch thick Mumetal shield around the tube which brought the gain back up to 15% below average. The experiment has substantially thicker magnetic shielding around the tubes. The detector also allows calibration of the tubes in situ, subtracting the effects of magnetic fields at the tubes as long as these effects are stable.

B.3 Linearity

The dynamic range of signals in the experiment is from 1.5 to 400 photoelectrons. While non-linearity is regularly calibrated and accounted for in reconstruction, the magnitude of the correction should ideally be minimal. The prototype PMTs showed no significant nonlinearity to 300 photoelectrons. The average PMT in the detector

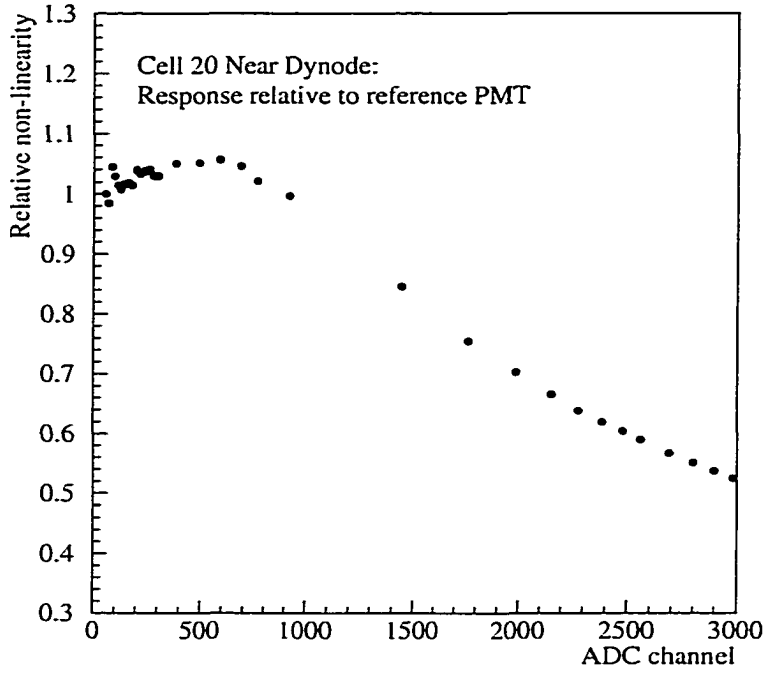

Figure B.5: The nonlinearity of a central detector PMT, measured as the response of the dynode output versus a reference phototube. These data from regular calibrations were used to correct for nonlinear effects offline. 3000 ADC channels represents approximately 1200 photoelectrons, or 60 MeV at the center of the detector.

has around 10% deviations over the range up to ~400, a typical example is shown in Figure B.5.

Note that the nonlinearity is positive for some signal sizes. This is due to a drop of voltage across the last dynodes as the current is shunted through the phototube rather than flowing through the resistor of the last stage divider. The voltage is redistributed among the other stages, increasing their voltage and consequently gain. This will lead to a net increase in overall gain of the tube, or positive non-linearity. The addition of capacitors for charge storage across the last set of dynodes reduces but does not eliminate this effect. As signal sizes get larger the competing effect of dynode saturation will cause negative nonlinearity. Bases can be optimized for linearity by but at the expense of worse timing resolution.

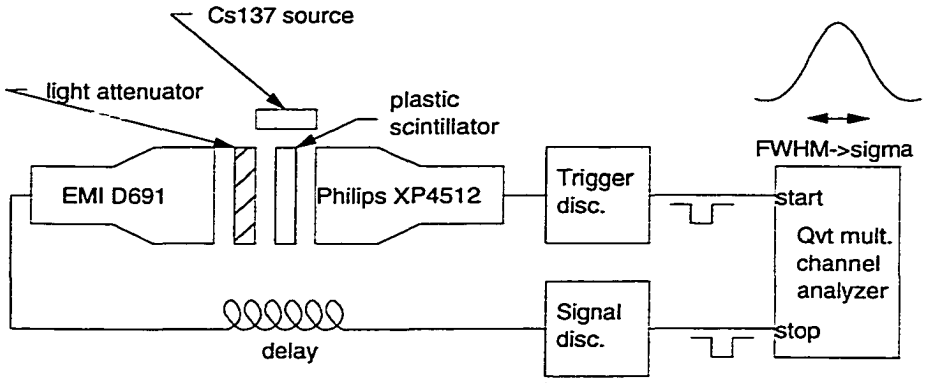

Figure B.6: The setup used for the PMT timing jitter measurement.

The commissioning set of energy scale calibrations showed a surprising 100% positive nonlinearity at the ^{228}Th 2.6 MeV point, which could not be explained by the effects described above. This was eventually attributed to the relatively low current used by the bases, 200 μamp. If the average current demand at the anode output becomes a significant fraction of this, greater than a few percent, the positive nonlinearities mechanism is exacerbated since the effectiveness of the capacitors is reduced. We studied this phenomenon and all subsequent calibrations used sources which produce average currents <0.3% of the base quiescent current.

B.4 Signal transit time and rise time

As mentioned above the PMTs should not introduce more than 1 ns uncertainty into the signal arrival time. We investigated the transit time jitter and the transit time's variation with supply voltage. For transit time jitter the worst case is small signals, specifically single photoelectron signals. Larger signals, composed of many photoelectrons, have smaller jitters.

To measure the single photoelectron transit time jitter, the setup of Figure B.6 was used. Two PMTs were used, one illuminated by a large number of photons as trigger and the other, the tube under test, illuminated by one photoelectron on average. A small block of scintillator and a ^{137}Cs source were sandwiched between the two PMTs. The trigger PMT, a Philips XP4512, received the scintillator signal directly.

The EMI tube opposite it received light attenuated to the single photoelectron level by an optical filter. The discriminated trigger signal gave a start signal, which due to the large signal size has little jitter. The discriminator of the delayed EMI signal gave a stop signal. The width of the resulting peak was measured to have sigma of 1.3 ns. The jitter intrinsic to the electronics was measured to be 80 ps, and the large signal jitter of the Philips, measured by the same setup with two Philips tubes receiving unattenuated signals, was found to be 180 ps. After subtraction of these factors in quadrature, then, the single photoelectron jitter of the tube is still 1.3 ns.

We measured the 10-90% rise time of the tube to be 4.8 ns.

B.5 Dynode output characteristics

A base was built and tested with outputs at the 10th and 11th dynodes per the schematic of Figure B.1b. After discovering problems with ringing in the dynode outputs we stabilized the output by replacing a resistor as mentioned above. We investigated the attenuation factor of the dynode outputs and their effect on the anode signal characteristics. By integrating the charge from the dynode output and dividing it by the charge from the anode output, we found attenuation factors of 40 for the dynode 10 output and 7 for the dynode 11 output. The PMTs installed used the tenth dynode for their high range output and had an average attenuation of 20. The shape, single photoelectron transit time jitter, and rise time of the anode pulse were remeasured with this base and are not significantly affected by the presence of the dynode outputs.

B.6 Dark current rates

Dark current, the rate of signals from a completely unlit tube, can contribute to false triggers which increases the deadtime of the experiment, adversely affecting efficiency. The second concern is the uncorrelated background contribution, degrading the signal to noise ratio. We measured the dark rate at various thresholds using a discriminator fed into a scaler. The tube was left powered at 1600 Volts and unlit for 12 hours

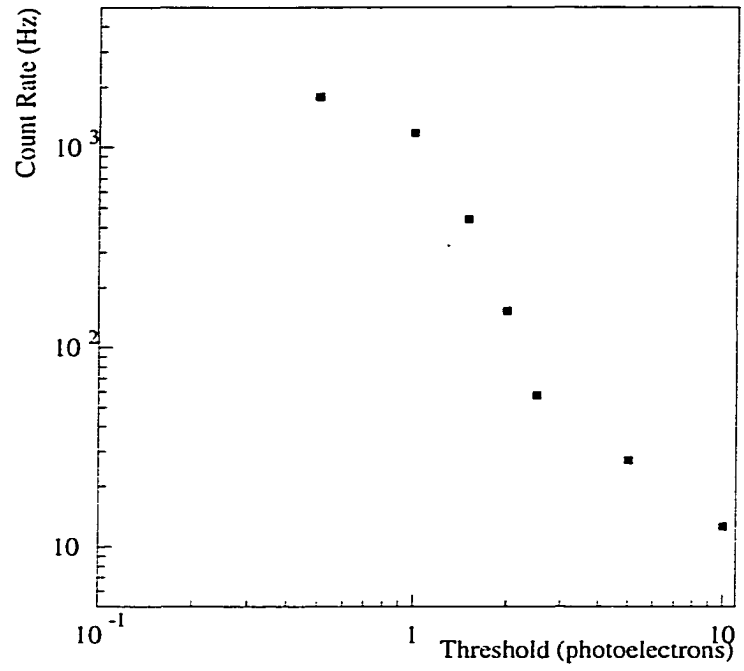

Figure B.7: Dark count rates for different thresholds. Thresholds are given in photoelectrons. Statistical errors are negligible.

before this measurement to avoid the effects of exposing the photocathode to light, which can lead to residual increased dark currents. The results in the form of rates for signals above a given threshold are shown in Figure B.7.

The Palo Verde trigger thresholds were 50 mV and 400 mV for the *low* and *high* thresholds, respectively. On average there were 750 kHz and 5 kHz rate of each. At the measured dark rates for these thresholds (1 and 8 photoelectrons), the dark rates of the 132 central detector PMTs were not the dominant contribution to the trigger rates.

B.7 Afterpulsing

Afterpulsing, an effect in which a real signal pulse is followed by ghost pulses, is caused by two mechanisms. On short (< 100 ns) time scales luminous reactions,

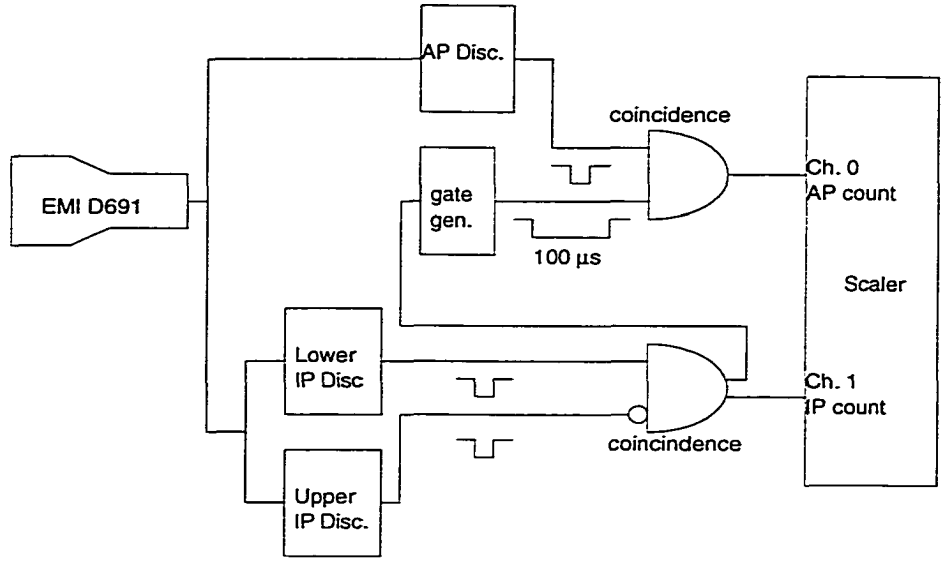

Figure B.8: Electronics setup for the afterpulse measurement. AP is afterpulse and IP is initial pulse, the two IP discriminators trigger on a range of amplitudes.

in which light is emitted by the dynode chain as they are bombarded by electrons. This light, if able to reach the photocathode, can initiate an afterpulse. This can cause positive nonlinearity effects, and in severe cases the positive feedback can become uncontrollable and the PMT cease to function. At high gains this effect was a problem for the PMTs used at Palo Verde and eventually necessitated the reduction of the central detector PMT gain by 20%. The second afterpulse mechanism comes from propagating electrons ionizing residual gases in the PMT vacuum, producing a delayed afterpulse up to several microseconds later when the positive ion hits the photocathode. Afterpulses represent noise with a time correlation, which can be triggered on relatively easily in our trigger scheme. Afterpulses large enough to mimic neutron capture are rare; the main concern is loss of livetime from digitizing these events.

Practically we are interested in the probability that an initial pulse of a given size will create an afterpulse of another given size during the 5-200 μs interval of the signal's delayed coincidence. This measurement involved the setup shown in Figure B.8. The signal was split, one half going to the pair of discriminators triggering on the initial pulse, and the other half going to the afterpulse discriminator. A

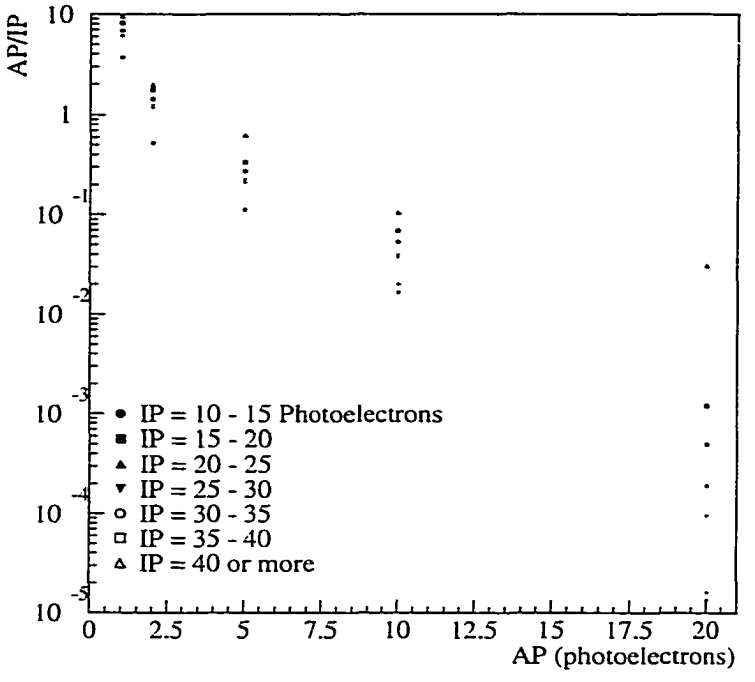

Figure B.9: Afterpulse probabilities by initial pulse range. The horizontal axis is the threshold of the afterpulse discriminator. Different symbols are for different initial pulse ranges. The statistical error is negligible above $10^{-}2$.

pair of discriminators defined the range of the initial pulse accepted by requiring a coincidence of the lower threshold discriminator output with the not of the upper threshold discriminator, where the lower threshold gate was timed to completely subtend the upper threshold gate. This coincidence triggered a initial pulse (IP) scaler count and generated 150 ns later a 100 μs long gate. The third discriminator was set to the afterpulse threshold, which in coincidence with the gate triggered an afterpulse (AP) scaler count. After subtracting the dark rate contribution to AP, we define the afterpulse probability as AP/IP. The results of this measurement are in Figure B.9. Note that the larger the initial pulse, the more likely an afterpulse, up to expecting around 10 single photoelectron afterpulses for particularly large signals.

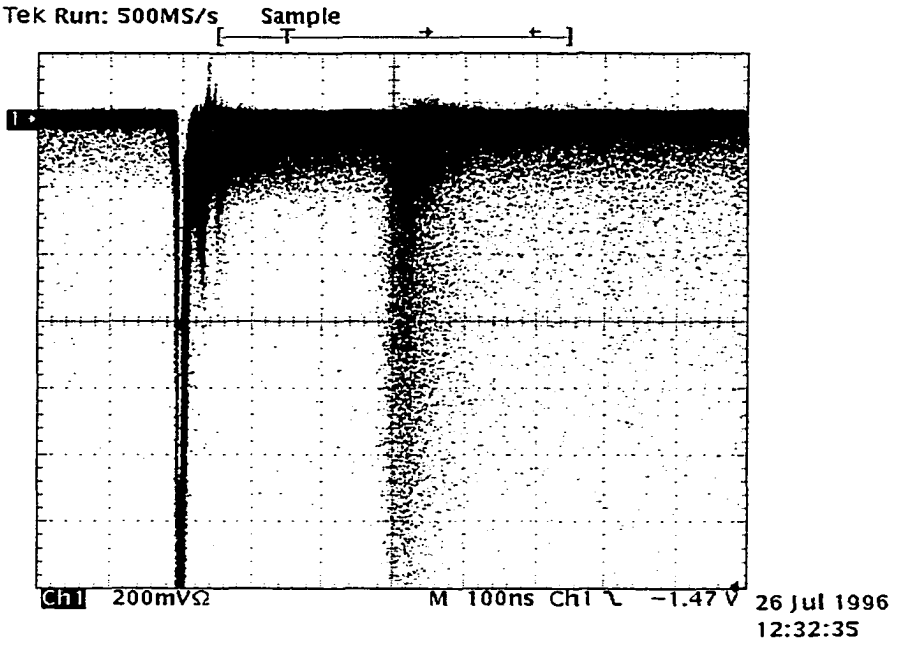

Figure B.10: An infinite persistence trace of the tube output. The horizontal scale is 100 ns/div. Note the afterpulses 300 ns after the main pulse.

There is a time structure for the afterpulsing. The afterpulsing is shown on two time scales in Figures B.10 and B.11. Specifically there is a tendency for an afterpulse to follow approximately 300 ns after the initial pulse. This is probably due to the presence of a specific contaminant in the phototube vacuum such as Helium. The afterpulse probabilities were remeasured with the 100 μs gate delayed for one μs to exclude this band; we found that the 300 ns band does not represent a significant portion of the total afterpulsing. The afterpulses also seem to occur almost entirely within the first five microseconds after the initial pulse, motivating the 5 μs trigger suppression after a *triple* trigger in data taking.

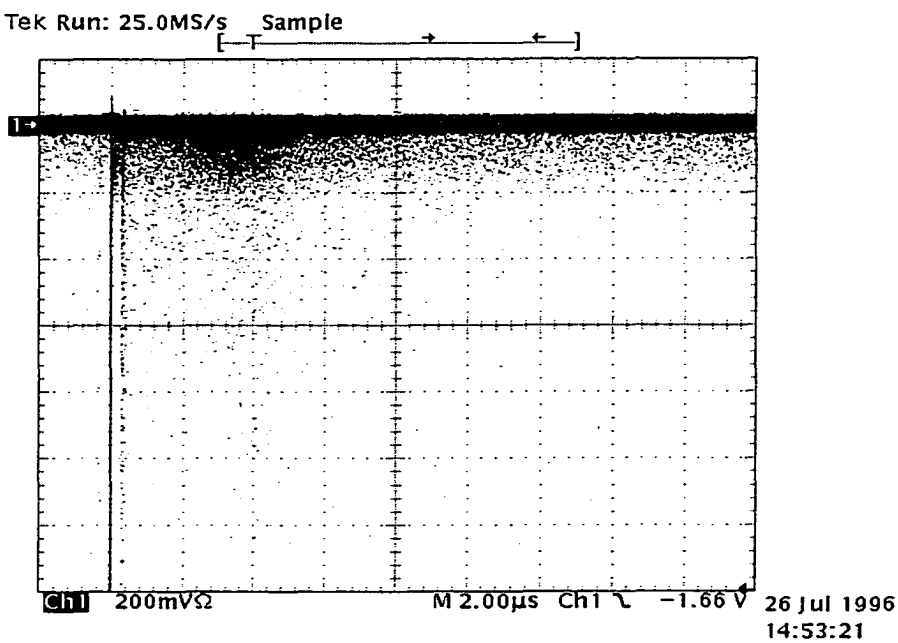

Figure B.11: An infinite persistence trace of the tube output. The horizontal scale is 2 μs/div.

Appendix C

In other words...
(A Non-technical Synopsis)

The elusive neutrino has gathered a lot of interest in experimental physics over the last fifty years. Wherefore the neutrino?

Neutrinos are surprisingly ubiquitous. They are generated in many nuclear interactions: in the fusion processes of the sun, in the natural radioactivity of the earth, in the constant shower of cosmic radiation in the upper atmosphere, and most copiously during the Big Bang of creation. Any given cubic centimeter of the universe has around 300 of them, making them a billion times more numerous than the proton or electron, the more mundane components of the universe. A natural question to ask about neutrinos, in consideration of their numbers, is whether they have mass. Even a small amount could spell a large difference in the composition of the universe. Until recently, nobody knew whether the neutrino had any mass. But lately there have been tantalizing hints.

Despite their numbers, neutrinos are the antisocial species in the particle zoo. They have no electric charge and don't interact with the strong nuclear forces binding protons and neutrons together in atomic nuclei. Neutrinos only interact through a force appropriately called the weak force. (They will interact through gravity as well, but gravity is far weaker even than the weak force and can be neglected.) The weak force is well named, for a neutrino could easily pass through a light year (10^{15} m) of

lead without noticing. Needless to say this presents the experimenter hoping to see neutrinos at a disadvantage.

There are three species of neutrinos named the electron, muon and tau neutrino in reference to the fundamental particle which they are paired with in their interactions. The muon and the tau particles are identical to the electron but successively heavier, but since they are not stable and tend to decay to the lighter electron, they are not common particles.

Direct searches for neutrino mass normally look for small changes in the energies of particles emitted in particle decays. If a neutrino is involved in the decay and has some mass, the energy tied up in the mass of the neutrino will be missing in the overall energy balance of the rest of the decay. Electron neutrinos are emitted in the decay of a special form of hydrogen called tritium, for example, and careful observations of the energies released in these decays have been performed to look for signs of neutrino mass. These experiments have seen no signs yet of a neutrino mass, and together the results imply that the mass of the electron neutrino is less than three electron Volts (eV). This impressively small mass limit is a few hundredth thousand times lighter than the electron, the lightest particle of known mass. Limits on the mass of the other kinds of neutrino are much worse because they are only produced in the decays of more exotic particles.

However, an even more sensitive probe of neutrino mass is available through the phenomenon of oscillation, a consequence of the quantum mechanical nature of neutrinos. According to quantum theory, particles can be modeled as traveling waves. In the macroscopic world we live in the characteristic sizes of these waves are far too small to be of practical relevance, but observable consequences of this wave-like nature are available in the behavior of elementary particles. The wavelength of a particle is related to its energy, higher energies having shorter wavelengths. This phenomenon occurs in light as well: for example, a very hot object which tends to emit more energetic light appears blue (blue has a short wavelength); when it cools appears to get redder and redder as the energy of the light it emits is reduced (red has longer wavelengths). Early theories explaining this temperature related radiation, called *black-body* radiation, were some of the earliest precursors of the modern form

of quantum mechanics.

The wave-like behavior is similar for neutrinos: the higher their energy, the shorter their wavelength. However, if they have mass, some of that energy is tied up not just in its momentum, but also in its mass according to the equation first come up with by Einstein, $E = mc^2$. We cannot easily measure the wavelength of a neutrino but nature may allow a special method for measurement via *mixing*. In *mixing*, a particular neutrino species, say the electron type, may actually be composed of not one traveling wave, but two or even more. When particles are actually composed of more than one wave (wavefunctions in the language of quantum mechanics) they are called superpositions. Each wavefunction represents the same momentum for the neutrino as was decided at the time of the neutrino's creation when the available energy of the decay was divided up. The two wavefunctions, however, may have different masses, and hence different wavelengths (since the total energy of each wave is different).

While it may sound bizarre to speak of a particle being a superposition of different masses, this phenomenon has been observed in other situations. Strange as it sounds, the particle state which interacts (called a weak interaction *eigenstate*, e.g. the electron neutrino) is not necessarily the same particle state which propagates through space (called a mass *eigenstate*). *Eigenstate* is another way of referring to a single wavefunction.

When a mixed state propagates, the component wavefunctions will evolve differently, each according to their wavelength. The superposition of the two, which forms the neutrino itself, therefore changes as it propagates. In fact it can, according to the relative point in a cycle of the two components, become an entirely different flavor, no longer an electron neutrino but rather a muon or tau neutrino. Oscillation, as this is called, is illustrated in the left side of Figure C.1. A close analogy to this phenomenon is seen (or rather heard) in sound waves (see the right side of Figure C.1). When a sound is composed of two close tones, a phenomenon called beating occurs, where the sound strengthens and weakens as the two tones reinforce or interfere with each other. Anyone who has tuned a guitar to a pitchfork may have noticed beating.

What governs the way the neutrino oscillates is the difference in mass of the

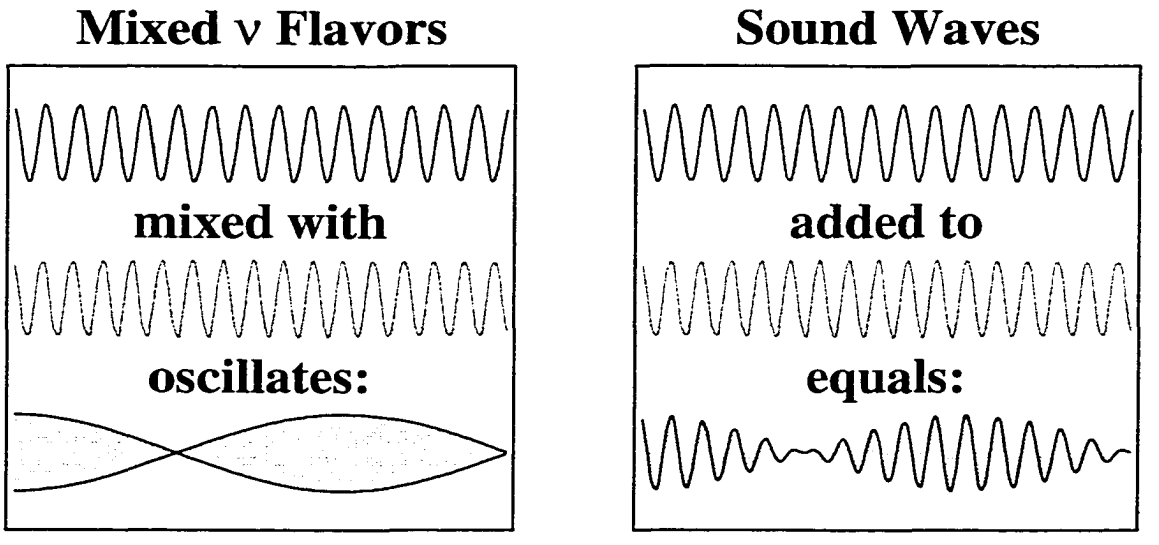

Figure C.1: Left: If neutrinos, when created as "flavor eigenstates", are actually mixed states of "mass eigenstates" then their flavor composition will oscillate as the different wavelengths reinforce or interfere with each other. Right: If a sound is composed of two slightly different pure tones, the sound will vary in intensity as the different wavelengths reinforce or interfere with each other. This phenomenon, called "beating", is analogous to what happens in neutrino oscillation.

superposed waves in its makeup. By looking for changes in the composition of a neutrino beam as it propagates we are in effect looking for these mass differences, hence the concept of a neutrino oscillation experiment.

Indications of oscillation were first seen 25 years ago at an experiment in the South Dakota Homestake mine designed to see neutrinos creating in the solar interior. They saw roughly a third of the expected number of neutrinos. Our understanding of the sun's energy production mechanism, in which huge quantities of neutrinos are produced in the fusion of hydrogen to form helium, has been confirmed to a great precision by helioseismology, in which the passage of sound waves through the solar interior is used to study the solar core. The solar model is therefore not in great doubt, yet other experiments have confirmed this neutrino deficit. The best solution for the solar neutrino "problem" is that the neutrinos are disappearing through oscillation into other species to which the experiments are insensitive.

Experiments designed to see neutrinos created in the upper atmosphere from cosmic radiation have also seen fewer than the expected number of neutrinos. One nice feature of these "atmospheric" neutrinos is that by measuring from which direction they came from, i.e. from below or above the experiment, one can deduce how far they traveled. Correlating the deficit with distance traveled gives the actual difference in wavelength of the superposition and hence the mass differences involved in the mixed neutrino state. The first neutrino oscillation solutions to the atmospheric neutrino anomaly were published in the early 1990's, and predicted significant mixing of neutrino states with mass differences on the order of $\sim$0.13 eV.

In view of the oscillation signatures from extraterrestrial sources, experiments have been focused lately on understanding neutrino oscillation via man-made sources. Coming either from accelerated particle beams or nuclear reactors, these man made neutrinos have the advantages of being relatively well understood and controllable.

The Palo Verde detector was one of two reactor neutrino experiments designed to look for the oscillations seen in atmospheric neutrino experiments. PVNGS, the Palo Verde Nuclear Generating Station, has three reactors, producing a total thermal output of 11.6 GW. Our neutrino detector was placed at a kilometer distance from the reactors in an vault 50 feet underground.

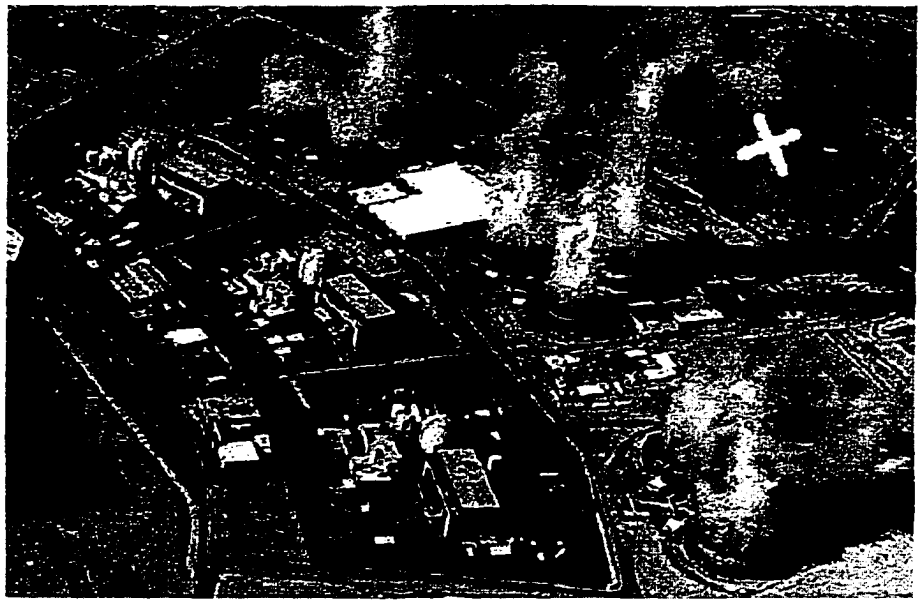

Figure C.2: The Palo Verde Nuclear Generating Station as seen from the air. The three reactors are housed in the dome shaped buildings on the left. The neutrino detector was located 50 feet underground at the white "X".

Reactors produce copious numbers of electron neutrinos (actually antineutrinos but the distinction is not important): 10^{22} neutrinos traversed the detector target every day when all three reactors were on. Of those, only 220 were expected to interact with the 15 tons of target material (assuming none of them oscillated to another neutrino type on their way!) We looked for oscillations by counting the number of electron neutrinos we saw at our one kilometer distance versus the known number of neutrinos emitted by the reactors. Any deficit seen would be a sign of oscillation. This form of neutrino experiment is sometimes referred to as a disappearance experiment. In order to accurately perform the measurement we needed to be careful of understanding exactly how efficient our detector was for seeing the interacting neutrinos, and in controlling sources of false signals from the background radioactivity around and in our detector.

A cut-away view of the detector is drawn in Figure C.3 to show the various layers. We employed several active and passive methods of suppressing backgrounds in our neutrino data. The innermost layer, an array of 66 acrylic cells filled with an liquid

scintillator, served as the neutrino target. The scintillator is a mixture of organic compounds designed to give off light when energy from particle interactions is deposited within it. The target was surrounded along the sides by a 1 meter buffer of water in steel tanks which, along with a buffer of pure mineral oil at the end of each cell, served as a passive shield against gamma rays and other sources of background noise from getting into the target. This buffer was surrounded on all sides by an active layer called the cosmic muon veto. Composed of 40 tanks filled with more liquid scintillator, this layer registered the passage of any cosmic muons which passed through the detector. These particles, formed by extraterrestrial protons hitting the Earth's atmosphere, can easily cause target interactions which mimic a neutrino event. We therefore did not allow any data to be valid during and immediately following their passage.

The reactor antineutrinos interacted in the target acrylic or scintillator via inverse beta decay,

$$\bar{\nu}_e + p \longrightarrow e + n,$$

which is essentially the reverse of the neutron decay process which creates them (see Figure C.4). The products of inverse beta decay, a positron and neutron, form a two part signal. The prompt part consists of the positron losing its kinetic energy and then annihilating with an electron, which forms a characteristic pair of 511 keV gamma rays. The neutron has to be captured by a nucleus before detection; in our case this capture occurs on the Gadolinium doped into the target scintillator, usually around 28 μs after its creation. Gadolinium has a large affinity for neutron capture and also releases a substantial burst of energy, 8 MeV, upon doing so.

When an energy deposit happens in the target, scintillation light is released and propagates towards the ends of the acrylic cells, where it is collected by photomultipliers, light sensitive detectors which convert the light into electrical impulses which can be recorded and analyzed. To reconstruct what happened in an event, we find the location of each energy deposit by the relative time of arrival of the light at each end of the cell, and the total energy deposited by noting the total light collected and correcting this light by a factor to account for extinction of light due to passage

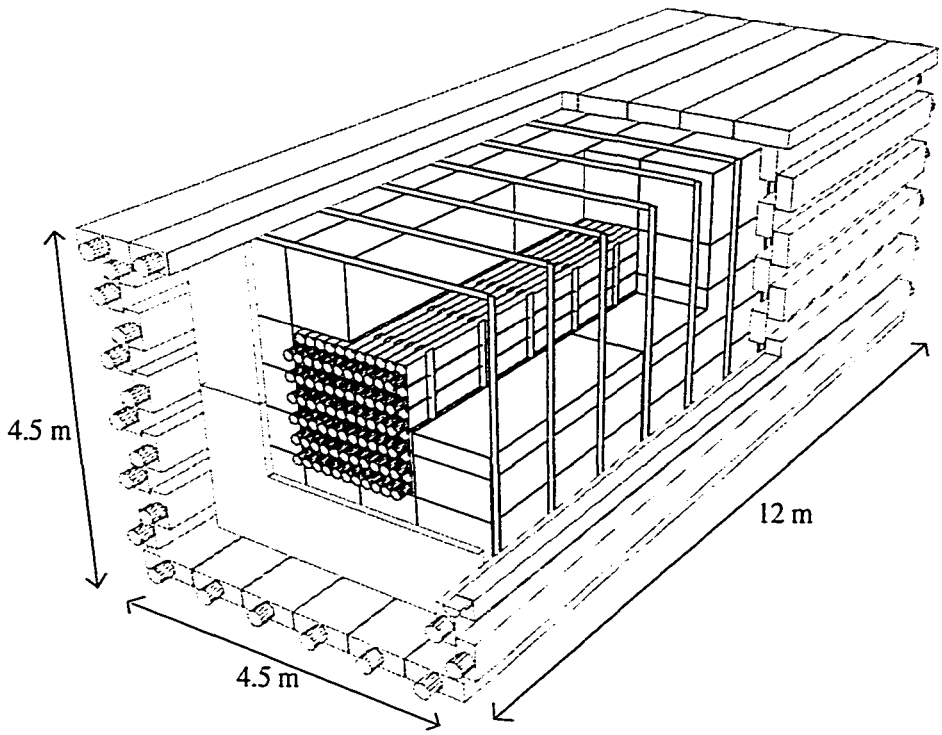

Figure C.3: The Palo Verde detector consists of three layers: a target array of 66 acrylic cells filled with an oil based liquid scintillator in the center, a one meter layer of passive water shielding, and an outer layer of active veto detectors to suppress cosmic muon induced activity.

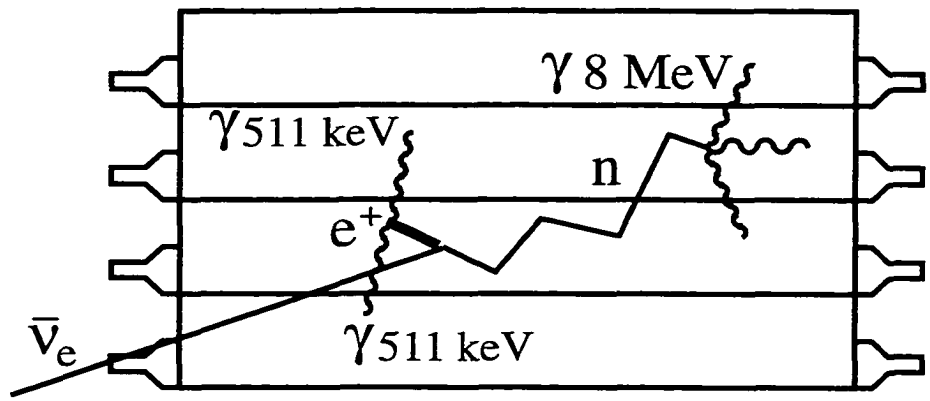

Figure C.4: A schematic of the inverse beta decay signal. The ionization and annihilation of the positron forms a prompt triple coincidence of energy deposits, followed by the thermalization and capture of the neutron forming a delayed triple coincidence.

through the scintillator.

In addition to the time correlation of the signal we used the detector segmentation to discriminate neutrinos from background events. Both the prompt and delayed parts of the inverse beta decay have a reasonable chance of depositing energy in three cells simultaneously. We therefore only recorded events which were triggered by finding two triply coincident energy deposits in close time succession.

After analyzing the reconstructed events and selecting the events which look like neutrino events, i.e. have appropriate hit combinations and energy characteristics, we ended up with around 50 events per day as a neutrino candidate set. We then needed to determine how many neutrinos actually interacted in the detector, and to account for the remaining backgrounds which were able to mimic the neutrino signal. The former required knowing the absolute efficiency for an inverse beta decay event to survive the trigger requirements and the further selections made.

The most straightforward method of determining the efficiency for neutrino detection would have been to calibrate the detector with a neutrino source. As these are not readily available in transportable form, we had to resort to a less direct method involving calibrating for each part of the event separately: use a positron source to measure the efficiency for the prompt portion of the events and use a neutron source to measure the efficiency for the delayed neutron capture portion of the events. Together

they make a reasonable analog for the whole inverse beta decay event.

However there were some differences between the source characteristics and the real signal events; for example, the neutrons from the decays of the Americium-Beryllium (Am-Be) neutron source used are much more energetic than those from neutrino interactions, and originated from where the source was inserted in the interstices of the target cell array, rather than within the target mass itself. To get around these differences, we used a detailed simulation of the detector and checked that it accurately predicted the response of the detector to these calibration sources. Once we convinced ourselves of the efficacy of this "Monte Carlo" simulation, we could rely on this to calculate our overall neutrino detection efficiency. We repeated these calibrations four times throughout the data taking period and at many locations throughout the target to ensure the simulation quality was consistent (see Figure C.5).

Based on these calibrations we found a 10.8% efficiency for neutrino detection. Noting that we had ~50 neutrino candidates per day in our data this implies over 450 inverse beta decays per day in our data. We expected only 220 interactions per day with no oscillations; the surplus is due to false neutrino candidates (background) still in our data set.

Two methods were used to account for the background in the data set. One method used the periods of reduced flux while one of the three reactors was off for refueling. By measuring the difference in neutrino candidate rates, we directly measured the refueled reactor's contribution to our signal. The second method of subtracting background used a more sophisticated understanding of the nature of the backgrounds to isolate their contribution to the neutrino candidate set. Briefly, the latter method used the unique neutrino signature of the prompt positron annihilation gamma rays, a low-energy feature lacking in the backgrounds.

After performing this subtraction, we can then directly retrieve the measured neutrino interaction rate in the detector and compare it with the calculated interaction rate. For both methods of subtraction, we measure a rate in line with expectation for no oscillation. The final ratio of observed to calculated neutrinos was $\frac{R_{\rm obs}}{R_{\rm calc}} = 0.982 \pm 0.023(\text{stat.}) \pm 0.053(\text{sys.})$ where *stat.* is the statistical component of the measurement error and *sys.* is the total systematic errors stemming from experimental

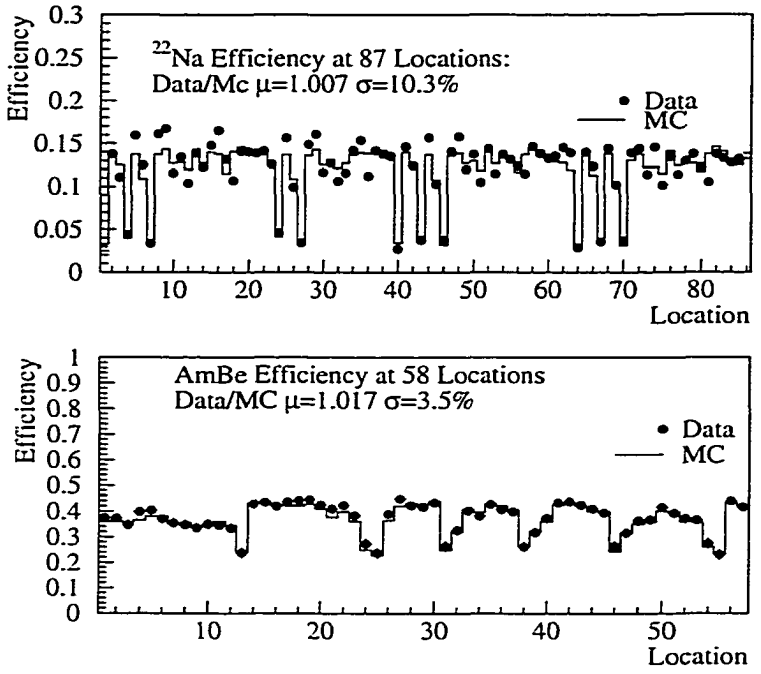

Figure C.5: To check that we are able to accurately estimate the detector efficiency for neutrino events, data were taken at many locations throughout the target with a positron (^{22}Na) and neutron (Am-Be) source. Plotted here are the Monte Carlo simulated and measured efficiencies for the two sources. On average they agree very well. Locations near the edge of the target tended to have lower efficiencies.

uncertainty.

In conclusion, the Palo Verde experiment saw no evidence for electron neutrinos oscillating to other species of neutrino. This result implies that the atmospheric neutrino anomaly mentioned above is not due to oscillation between muon and electron type neutrinos, but rather must stem from oscillations between the muon and tau type neutrinos. Future experiments are planned to test this scenario using very long baseline (700 km) beams of accelerator neutrinos. The results of Palo Verde rule out electron neutrinos being mixed to any significant degree with neutrinos whose mass difference is greater than ~0.03 eV.

What about future reactor neutrino experiments? One upcoming experiment in Japan called KamLAND will begin data taking soon with a baseline 100 times that

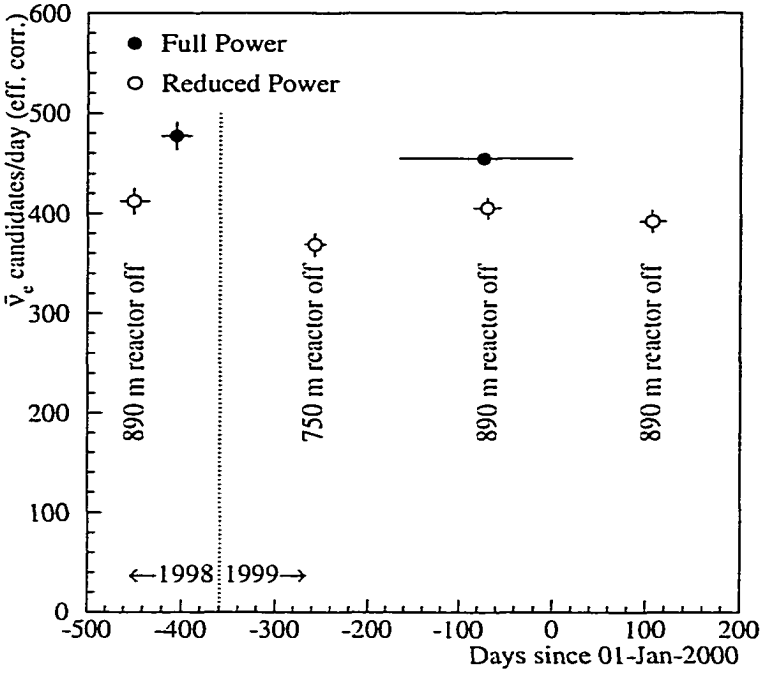

Figure C.6: The apparent target neutrino interaction rate plotted for full power (solid dots) and for the four refueling periods (hollow dots). The difference in candidate rates represents the refueled reactor's contribution to the data set. Data taken before 1999 were treated separately due to hardware changes in the detector.

of Palo Verde. It will probe significantly smaller neutrino masses than any previous neutrino experiment with a man-made source. The experience gained and analysis techniques used at Palo Verde should be readily applicable there.

Bibliography

[1] DONUT, M. Nakamura, Nucl. Phys. Proc. Suppl. **77**, 259 (1999), see also Neutrino 2000 proceedings.

[2] J. Bonn *et al.*, Nucl. Phys. Proc. Suppl. **87**, 271 (2000).

[3] V. M. Lobashev *et al.*, Phys. Lett. **B460**, 227 (1999).

[4] K. Assamagan *et al.*, Phys. Rev. **D53**, 6065 (1996).

[5] ALEPH, R. Barate *et al.*, Eur. Phys. J. **C2**, 395 (1998).

[6] M. Gunther *et al.*, Phys. Rev. **D55**, 54 (1997).

[7] D. O. Caldwell and R. N. Mohapatra, Phys. Rev. **D48**, 3259 (1993).

[8] J. N. Bahcall and S. L. Glashow, Nature **326**, 476 (1987).

[9] B. Pontecorvo, Sov. Phys. JETP **6**, 429 (1957).

[10] Z. Maki, M. Nakagawa, and S. Sakata, Prog. Theor. Phys. **28**, 870 (1962).

[11] S. M. Bilenkii and B. Pontecorvo, Phys. Rept. **41**, 225 (1978).

[12] L. B. Okun, M. B. Voloshin, and M. I. Vysotsky, Sov. Phys. JETP **64**, 446 (1986).

[13] J. N. Bahcall, S. Basu, and M. H. Pinsonneault, Phys. Lett. **B433**, 1 (1998), astro-ph/9805135.

[14] See http://www.sns.ias.edu/ jnb/ for a discussion of solar neutrinos with recent results.

[15] B. T. Cleveland *et al.*, Astrophys. J. **496**, 505 (1998).

[16] SAGE, J. N. Abdurashitov *et al.*, Phys. Rev. **C60**, 055801 (1999), astro-ph/9907113.

[17] GALLEX, W. Hampel *et al.*, Phys. Lett. **B447**, 127 (1999).

[18] Kamiokande, Y. Fukuda *et al.*, Phys. Rev. Lett. **77**, 1683 (1996).

[19] Super-Kamiokande, Y. Fukuda *et al.*, Phys. Rev. Lett. **82**, 2430 (1999), hep-ex/9812011.

[20] L. Wolfenstein, Phys. Rev. **D17**, 2369 (1978).

[21] S. P. Mikheev and A. Y. Smirnov, Sov. J. Nucl. Phys. **42**, 913 (1985).

[22] J. N. Bahcall and P. I. Krastev, Phys. Lett. **B436**, 243 (1998), hep-ph/9807525.

[23] Borexino, A. Ianni, Nucl. Phys. **A663-664**, 791 (2000).

[24] SUDBURY NEUTRINO OBSERVATORY, G. Aardsma *et al.*, Phys. Lett. **B194**, 321 (1987).

[25] P. Alivisatos *et al.*, STANFORD-HEP-98-03.

[26] Soudan-2, W. W. M. Allison *et al.*, Phys. Lett. **B449**, 137 (1999), hep-ex/9901024.

[27] R. Becker-Szendy *et al.*, Nucl. Phys. Proc. Suppl. **38**, 331 (1995).

[28] Kamiokande, S. Hatakeyama *et al.*, Phys. Rev. Lett. **81**, 2016 (1998), hep-ex/9806038.

[29] Super-Kamiokande, M. Nakahata, Nucl. Phys. Proc. Suppl. **87**, 125 (2000).

[30] V. Barger *et al.*, Phys. Lett. **B462**, 109 (1999), hep-ph/9907421.

[31] LSND, C. Athanassopoulos *et al.*, Phys. Rev. Lett. **81**, 1774 (1998), nucl-ex/9709006.

[32] Y. Declais *et al.*, Nucl. Phys. **B434**, 503 (1995).

[33] CCFR/NuTeV, A. Romosan *et al.*, Phys. Rev. Lett. **78**, 2912 (1997), hep-ex/9611013.

[34] L. Borodovsky *et al.*, Phys. Rev. Lett. **68**, 274 (1992).

[35] KARMEN, K. Eitel, (2000), hep-ex/0008002.

[36] KARMEN, T. E. Jannakos, Nucl. Phys. Proc. Suppl. **85**, 84 (2000).

[37] BooNE, E. Church *et al.*, (1997), FNAL P898.

[38] J. P. Meyer, Nucl. Phys. Proc. Suppl. **87**, 309 (2000).

[39] K2K, K. Nakamura, Nucl. Phys. **A663-664**, 795 (2000).

[40] ICARUS and NOE, F. Arneodo *et al.*, INFN-AE-99-17.

[41] OPERA, M. Guler *et al.*, CERN-SPSC-2000-028.

[42] F. Reines and C. L. Cowan, Phys. Rev. **90**, 492 (1953).

[43] G. Zacek, PhD thesis, Technical University Munich, 1984.

[44] M. Apollonio *et al.*, Phys. Lett. **B466**, 415 (1999), hep-ex/9907037.

[45] F. Boehm *et al.*, (2000), hep-ex/0003022.

[46] L3, M. Acciarri *et al.*, Phys. Lett. **B431**, 199 (1998).

[47] Model 9372 Electron Tubes Inc., Ruislip, UK.

[48] A. G. Piepke, S. W. Moser, and V. M. Novikov, Nucl. Instrum. Meth. **A432**, 392 (1999), nucl-ex/9904002.

[49] MACRO, S. P. Ahlen *et al.*, Nucl. Instrum. Meth. **A324**, 337 (1993).

[50] A. Marchioro, W. von Rueden, and G. McPherson, IEEE Trans. Nucl. Sci. **34**, 133 (1987).

[51] G. Gratta, L. Miller, C. Roat, D. Tracy, and Y. F. Wang, Nucl. Instrum. Meth. **A400**, 456 (1997).

[52] A. A. Hahn *et al.*, Phys. Lett. **B218**, 365 (1989).

[53] K. Schreckenbach, G. Colvin, W. Gelletly, and F. V. Feilitzsch, Phys. Lett. **B160**, 325 (1985).

[54] P. Vogel, G. K. Schenter, F. M. Mann, and R. E. Schenter, Phys. Rev. **C52** (2498).

[55] P. Vogel and J. F. Beacom, Phys. Rev. **D60**, 053003 (1999), hep-ph/9903554.

[56] L. Groshev *et al.*, Nucl. Data Tables **A5**, 1 (1968).

[57] R. Brun *et al.*, *GEANT 3*, , cern dd/ee/84-1 (revised) ed., 1987.

[58] P. Aarnio *et al.*, *FLUKA user's guide*, , cern tis-rp-190 ed., 1990.

[59] T. Gabriel *et al.*, , Ornl/tm-5619-mc ed., 1977.

[60] R. Craun and D. Smith, Nucl. Instrum. Meth. **60**, 239 (1970).

[61] M. Apollonio *et al.*, Phys. Rev. **D61**, 012001 (2000), hep-ex/9906011.

[62] R. Hertenberger, M. Chen, and B. L. Dougherty, Phys. Rev. **C52**, 3449 (1995).

[63] M. Aglietta *et al.*, Nuovo Cim. **C12**, 467 (1989).

[64] J. Delorme, M. Ericson, T. Ericson, and P. Vogel, Phys. Rev. **C52**, 2222 (1995), hep-ph/9504331.

[65] O. C. Allkofer and R. D. Andresen, Nucl. Phys. **B8**, 402 (1968).

[66] G. J. Feldman and R. D. Cousins, Phys. Rev. **D57**, 3873 (1998), physics/9711021.

[67] Y. F. Wang, L. Miller, and G. Gratta, Phys. Rev. **D62**, 013012 (2000), hep-ex/0002050.

[68] Karmen Collaboration, B. Armbruster (private communication).

[69] F. F. Khalchukov *et al.*, Nuovo Cimento **C6**, 320 (1983).

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

[71] K. V. Alanakian *et al.*, Sov. J. Nucl. Phys. **34**, 828 (1981).

[72] Super-Kamiokande, Y. Fukuda *et al.*, Phys. Rev. Lett. **81**, 1562 (1998), hep-ex/9807003.

[73] Kamiokande, Y. Fukuda *et al.*, Phys. Lett. **B335**, 237 (1994).

[74] Super-Kamiokande collaboration, preliminary result to be published.

[75] *Core Operating Limit Supervisory System (COLSS) and Core Protection Calculator (CPC) Meaurement Channel Uncertainties, 13-JC-RJ-205, Revision 2*, 1997.

[76] ABB Combustion Engineering Nuclear Operations, *PVNGS COLSS Secondary Calorimetric Power Uncertainty Analysis*, 1995, Proprietary.

[77] Thanks to Paul Gavin, Combustion Engineering, for help.

[78] Combustion Engineering, *Comparison of Calculated and Measured Transuranic Nuclide Content for Three- and Four-Cycle BG&E Unit 1 Fuel*, 1985.

[79] *Nuclear Administrative and Technical Manual: Secondary Calorimetric Power Verification.*

[80] P. Vogel and J. Engel, Phys. Rev. **D39**, 3378 (1989).

[81] P. Vogel, Phys. Rev. **D29**, 1918 (1984).

[82] CALTECH-SIN-TUM, G. Zacek *et al.*, Phys. Rev. **D34**, 2621 (1986).

[83] B. Achkar *et al.*, Phys. Lett. **B374**, 243 (1996).