

MEASUREMENT OF NEUTRINO OSCILLATION
WITH KAMLAND

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
Jason A. Detwiler
January 2005

UMI Number: 3162394

INFORMATION TO USERS

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 bleed-through, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send 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.

UMI[®]

UMI Microform 3162394

Copyright 2005 by ProQuest Information and Learning Company.

All rights reserved. This microform edition is protected against unauthorized copying under Title 17, United States Code.

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

© Copyright 2005 by Jason A. Detwiler
All Rights Reserved

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.



Blas Cabrera

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:

Abstract

A measurement of neutrino oscillation based on a 599 ton-year exposure of the Kamioka Liquid-scintillator Anti-Neutrino Detector (KamLAND) is presented. In the absence of oscillation, we expect 310.0 ± 23.1 reactor electron-antineutrinos ($\bar{\nu}_e$'s) with energies $E > 3.4$ MeV, and an additional 13.2 ± 5.9 background events. Only 226 events were detected, corresponding to $\bar{\nu}_e$ disappearance at the 99.96% CL. Interpreted in terms of two flavor neutrino oscillations $\bar{\nu}_e \leftrightarrow \bar{\nu}_x$, the best fit to the KamLAND data gives a mass-squared difference $\Delta m^2 = 8.05^{+0.95}_{-0.72} \times 10^{-5}$ eV² and a mixing angle $\sin^2 2\theta = 0.68^{+0.20}_{-0.22}$. These values are in excellent agreement with the Large Mixing Angle solution to the “solar neutrino problem”. The distortion of the measured spectral shape is consistent with neutrino oscillation, and disagrees with the un-oscillated spectral shape at 98.2% significance. Assuming CPT invariance, a combined analysis of solar neutrino results and the KamLAND spectrum yields the precision measurement $\Delta m^2 = 7.86^{+0.75}_{-0.59} \times 10^{-5}$ eV², $\tan^2 \theta = 0.391^{+0.036}_{-0.034}$.

Acknowledgments

The work presented here resulted from the efforts of many KamLAND collaborators, whom I would like to recognize and commend for their hard work and diligence in performing this beautiful experiment. I would like to give special recognition to Suzuki Atsuto, Giorgio Gratta, Stuart Freedman, and the KamLAND Collaboration Council for their leadership and vision. KamLAND is supported by the Center of Excellence program of the Japanese Ministry of Education, Culture, Sports, Science, and Technology under grant 09CE2003, and by the United States Department of Energy under grant DEFG03-00ER41138. The reactor data are provided courtesy of the following electric associations in Japan: Hokkaido, Tohoku, Tokyo, Hokuriku, Chubu, Kansai, Chugoku, Shikoku, and Kyushu Electric Power Companies, Japan Atomic Power Company, and the Japan Nuclear Cycle Development Institute. The Kamioka Mining and Smelting Company provided service for activities in the mine. Part of my funding was provided by a generous grant from the ARCS Foundation.

I am sincerely indebted to my adviser, Giorgio Gratta, for his consistent guidance over the years, for the insights he gave me into the methods of scientific exploration, and for his patience. My deepest thanks go to Nikolai Tolich, with whom I worked closely throughout my years at Stanford, for exemplifying a clear-headed approach to research, for our many discussions, and for his direction, assistance, and friendship. Thank you also to Patrick Decowski and Glenn Horton-Smith for their tireless efforts and for answering all of my questions, and to Kazumi Ishii, Bruce Berger, Dan Dwyer, Brian Fujikawa, and Lindley Winslow for their direct contributions and comments. Thank you also to Zelimir Djurcic, Sanshiro Enomoto, Evgueni Yakushev, and all of my KamLAND friends whose collaborative efforts and companionship made a difficult

task highly enjoyable. Special thanks go to Masayuki Koga and Kengo Nakamura for their guidance and camaraderie while I was in Japan, and to Fumiko Arao for her tutelage. お世話になりました。どうもありがとうございました。 Many thanks to Sam Waldman, Jesse Wodin, Yoshi Uchida, Thomas Koffas, Yifang Wang, the Gratta research group, and my fellow graduate students at Stanford University for their stimulating discussions and ideas. I would also like to thank my family and all of my friends, especially my parents Daryl and Elaine Detwiler, Kristi Detwiler, Kandace Hollenbach, Thomas Detwiler, and Sumiko and Akira Sebe for their love and support. Finally I would like to thank Joy Sebe, whose kindness, generosity, and endless support inspires and empowers me to pursue my dreams.

Contents

| | |
|---|-----------|
| Abstract | iv |
| Acknowledgments | v |
| 1 Introduction | 1 |
| 1.1 The Solar Neutrino Problem | 1 |
| 1.2 Neutrino Mass | 2 |
| 1.3 Neutrino Oscillation | 4 |
| 1.4 Matter Effects | 6 |
| 1.5 Neutrino Oscillation Experiments | 7 |
| 1.5.1 Solar Neutrinos | 7 |
| 1.5.2 Atmospheric Neutrinos | 12 |
| 1.5.3 Reactor Antineutrinos | 14 |
| 1.5.4 LSND | 15 |
| 1.6 Three-Flavor Oscillations | 17 |
| 2 KamLAND | 20 |
| 3 The Detector | 27 |
| 4 Event Reconstruction | 36 |
| 4.1 Event Building | 36 |
| 4.2 Waveform Analysis | 36 |
| 4.2.1 Pedestal and Baseline Subtraction | 37 |

| | | |
|----------|--|-----------|
| 4.2.2 | Pulse Finding | 39 |
| 4.2.3 | Sampling Rate Measurement | 40 |
| 4.2.4 | Q_0 , Yield, T_0 , and T - Q Corrections | 41 |
| 4.3 | Vertex Fitting | 43 |
| 4.3.1 | Charge-Based Pre-fit | 44 |
| 4.3.2 | Time-Based Fit | 44 |
| 4.3.3 | Calibration and Performance | 46 |
| 4.4 | Energy Estimation | 48 |
| 4.4.1 | Calculation of E_a | 49 |
| 4.4.2 | Calibration and Performance | 50 |
| 4.4.3 | Energy Resolution | 54 |
| 4.4.4 | Conversion from E_a to E | 55 |
| 4.5 | Muon Track Fitting | 58 |
| 4.5.1 | Muon Tagging | 58 |
| 4.5.2 | Track Reconstruction | 59 |
| 4.5.3 | Performance | 63 |
| 4.6 | Total Reconstruction Efficiency | 65 |
| 5 | Data Reduction | 67 |
| 5.1 | Data Cleaning and Livetime Calculation | 67 |
| 5.2 | Event Pairing | 69 |
| 5.3 | Time Correlation | 69 |
| 5.4 | Space Correlation | 71 |
| 5.5 | Delayed Energy Cut | 75 |
| 5.6 | Fiducial Volume Cut | 78 |
| 5.7 | Spallation Cuts | 84 |
| 5.8 | Prompt Energy Analysis Threshold | 86 |
| 5.9 | Total $\bar{\nu}_e$ Detection Efficiency | 87 |
| 6 | Backgrounds | 89 |
| 6.1 | $\bar{\nu}_e$ Backgrounds | 89 |
| 6.1.1 | Man-made $\bar{\nu}_e$ Sources | 89 |

| | | |
|----------|--|------------|
| 6.1.2 | Geological $\bar{\nu}_e$ Sources | 90 |
| 6.1.3 | Cosmic $\bar{\nu}_e$ Sources | 91 |
| 6.2 | Accidental Backgrounds | 92 |
| 6.3 | Correlated Non- $\bar{\nu}_e$ Backgrounds | 93 |
| 6.3.1 | Backgrounds from $^{13}\text{C}(\alpha, n)^{16}\text{O}$ | 94 |
| 6.3.2 | $^9\text{Li}/^8\text{He}$ Events | 99 |
| 6.3.3 | Fast Neutrons | 102 |
| 6.4 | Total Backgrounds | 103 |
| 7 | Reactor Signal Estimation | 104 |
| 7.1 | Fission Yields | 104 |
| 7.2 | $\bar{\nu}_e$ Spectra | 108 |
| 7.3 | $\bar{\nu}_e$ Flux and Baseline | 109 |
| 7.4 | Total Reactor $\bar{\nu}_e$ Signal | 109 |
| 8 | Analysis | 113 |
| 8.1 | $\bar{\nu}_e$ Disappearance and Average Survival Probability | 113 |
| 8.2 | Time Variation | 114 |
| 8.3 | Spectral Analysis | 116 |
| 8.4 | Spectral Distortion | 124 |
| 8.5 | Solar + KamLAND Analysis | 128 |
| 9 | Conclusions | 130 |
| 9.1 | Results and Implications | 130 |
| 9.2 | KamLAND Prospects | 131 |
| | Bibliography | 134 |

List of Tables

| | | |
|-----|--|-----|
| 4.1 | Muon Classification and Rates | 60 |
| 4.2 | Cuts for Evaluating the Total Reconstruction Efficiency | 66 |
| 5.1 | Absorption Lengths and Vertex Resolutions of Calibration Sources | 73 |
| 5.2 | $\bar{\nu}_e$ Detection Efficiency | 88 |
| 6.1 | Total Backgrounds | 103 |
| 7.1 | Systematic Uncertainties | 111 |

List of Figures

| | | |
|-----|---|----|
| 1.1 | Solar Neutrino Spectrum | 8 |
| 1.2 | Comparison of Measured and Theoretical Solar Neutrino Rates | 9 |
| 1.3 | Neutrino Oscillation Solutions for Solar Neutrino Experiments | 11 |
| 1.4 | Global Solar Neutrino Oscillation Solution (LMA) | 12 |
| 1.5 | Zenith Angle Distribution of Atmospheric ν 's in Super-Kamiokande | 13 |
| 1.6 | Atmospheric Neutrino Oscillation Solution | 14 |
| 1.7 | LSND Detected Energy Spectrum | 16 |
| 1.8 | LSND Neutrino Oscillation Solution | 17 |
| | | |
| 2.1 | Sensitivity of Reactor $\bar{\nu}_e$ Experiments | 21 |
| 2.2 | Reactor Locations Around Japan | 23 |
| 2.3 | Baseline Distribution of $\phi_{P_{th}}$ at KamLAND | 24 |
| | | |
| 3.1 | Detector Schematic | 28 |
| 3.2 | PMT Installation, Photograph | 29 |
| 3.3 | Inner Detector Construction, Photograph | 30 |
| 3.4 | Distribution of $nsmx_{ID}$ | 34 |
| | | |
| 4.1 | Example of a Raw Waveform | 38 |
| 4.2 | Pulse Finding Example | 40 |
| 4.3 | ^{60}Co Pulse-Time Distribution | 46 |
| 4.4 | Vertex Reconstruction Biases for Various Calibration Sources | 47 |
| 4.5 | Position Distribution of Physics Events | 48 |
| 4.6 | z -Variation of ^{60}Co Energy Reconstruction | 51 |

| | | |
|------|--|----|
| 4.7 | Energy of Spallation Neutrons Vs. r^3 | 52 |
| 4.8 | Time Variation of ^{60}Co Energy Reconstruction | 53 |
| 4.9 | Energy Distribution of Physics Events | 53 |
| 4.10 | Energy Resolution Function | 54 |
| 4.11 | Energy Scale Fit to γ Calibration Source Energies | 57 |
| 4.12 | Identifying Muons: $nsmx_{OD}$ Vs. n_{pe} | 59 |
| 4.13 | Muon Event Display | 61 |
| 4.14 | Definitions of Geometrical Parameters for Muon Tracks | 62 |
| 4.15 | Muon Track Angular Distribution | 64 |
| 4.16 | Reconstructed Muon Impact Parameter Distribution | 65 |
| | | |
| 5.1 | Neutron Capture Time | 70 |
| 5.2 | Δt Distribution of $\bar{\nu}_e$ Candidates | 71 |
| 5.3 | Δr Distribution of $\bar{\nu}_e$ Candidates | 72 |
| 5.4 | Δr Distributions of ^{68}Ge and $^{241}\text{Am}/^9\text{Be}$ Source Events | 74 |
| 5.5 | E_d Distribution of $\bar{\nu}_e$ Candidates | 75 |
| 5.6 | E_d Vs. Δt for $\bar{\nu}_e$ Candidates | 76 |
| 5.7 | $^{241}\text{Am}/^9\text{Be}$ Prompt Energy Distribution | 77 |
| 5.8 | $^{241}\text{Am}/^9\text{Be}$ np -Capture γ Energy Distribution | 77 |
| 5.9 | Distribution of $\bar{\rho}^2$ Vs. \bar{z} for $\bar{\nu}_e$ Candidates | 79 |
| 5.10 | r^3 Distribution of $^{12}\text{B}/^{12}\text{N}$ Events | 81 |
| 5.11 | $^9\text{Li}/^8\text{He}$ Radial Bias Vs. \bar{r}^3 | 82 |
| 5.12 | $^9\text{Li}/^8\text{He}$ Radial Bias Near r_{fid} | 83 |
| 5.13 | r^3 Distribution of Spallation Neutrons | 84 |
| 5.14 | Live-Fraction Distribution of Data Set | 86 |
| 5.15 | Prompt and Delayed Energies of $\bar{\nu}_e$ Candidates | 87 |
| 5.16 | Triggering Efficiency | 88 |
| | | |
| 6.1 | Prompt Energy Spectrum of Geo-Neutrino Backgrounds | 91 |
| 6.2 | Energy Distribution of Accidental Background | 93 |
| 6.3 | $^{210}\text{Po}/^{222}\text{Rn}$ Peak in $nsmx_{ID}$ Distribution | 95 |
| 6.4 | Time Variation of $^{210}\text{Po}/^{222}\text{Rn}$ Peak Height | 96 |

| | | |
|------|---|-----|
| 6.5 | Radial Distribution of ^{210}Po Events. | 96 |
| 6.6 | $^{13}\text{C}(\alpha, n)^{16}\text{O}$ Prompt Spectrum | 98 |
| 6.7 | Δt_μ Distribution of $^9\text{Li}/^8\text{He}$ and $\bar{\nu}_e$ Events | 99 |
| 6.8 | Energy Distribution of $^9\text{Li}/^8\text{He}$ Events | 100 |
| 6.9 | L_μ Distribution of Spallation Neutrons | 101 |
| 7.1 | ^{235}U Fission Fragment Mass Number Distribution | 105 |
| 7.2 | Calculated Fission Rates of a BWR | 107 |
| 7.3 | $\bar{\nu}_e$ Energy Spectra of ^{235}U , ^{238}U , ^{239}Pu , and ^{241}Pu | 108 |
| 7.4 | $\bar{\nu}_e$ Flux at KamLAND | 110 |
| 7.5 | Effective Baseline Variation | 110 |
| 7.6 | Expected Reactor $\bar{\nu}_e$ Spectrum (No Oscillation) | 112 |
| 8.1 | Correlation of Observed and Expected $\bar{\nu}_e$ Rates | 115 |
| 8.2 | Measured Prompt Energy Spectrum and Oscillation Best-Fit | 121 |
| 8.3 | Oscillation Parameter Confidence Levels | 122 |
| 8.4 | Best-Fit Oscillation Parameters to MC Spectra | 124 |
| 8.5 | χ^2_P Distribution of MC Spectra | 125 |
| 8.6 | Goodness of Fit Variation with Binning | 125 |
| 8.7 | Scaled No-Oscillation Best-Fit Spectrum | 126 |
| 8.8 | Oscillation Parameter Confidence Levels, Shape Only | 127 |
| 8.9 | $L_{\text{eff}}/E_{\bar{\nu}_e}$ Distribution of $\bar{\nu}_e$ Candidates | 128 |
| 8.10 | Oscillation Parameter Confidence Levels, Solar + KamLAND | 129 |

Chapter 1

Introduction

1.1 The Solar Neutrino Problem

In 1962 Raymond Davis Jr. of Brookhaven National Laboratory set out to build a device capable of detecting neutrinos from the sun. His aim was to investigate the Standard Solar Model (SSM) [1], the theory of the complex chain of nuclear reactions that fuel the sun and keep it shining. In 1968 when Davis' first results were released [2], the SSM was a relatively flexible theory with poorly constrained parameters due to sparse solar experimental data with large uncertainties [3]. However, the SSM gave a testable prediction for the number of neutrinos that Davis' experiment would detect above the background.

Davis' experiment was located 1.5 km underground in the Homestake mine in South Dakota, and consisted of a large tank filled with 100,000 gallons of a chlorine-based cleaning fluid, C_2Cl_4 . Neutrinos interacting with a ^{37}Cl atom in the fluid would create radioactive ^{37}Ar (threshold = 814 keV), the production rates of which were measured monthly by monitoring its activity. The incident neutrino flux inferred from the measurements marked the first detection of neutrinos from the sun, or any other non-terrestrial source, an achievement for which Davis shared the 2002 Nobel Prize in Physics. However, it was the exact value of the neutrino flux detected that set off a flurry of activity in astro-particle physics that has since captured many of the best minds and hands in the field, pushing technological limits and generating

new ideas and constraints for theories beyond the Standard Model of particle physics [4]: Homestake only detected about one third of the neutrinos the SSM predicted it would.

In these early years, there was a great deal of skepticism about the discrepancy between Homestake and the SSM, which came to be known as the Solar Neutrino Problem [5, 6]. In particular, large uncertainties in the SSM prevented many scientists from taking these initial results seriously [7]. But as the years progressed, the SSM was dramatically refined, and the statistical and systematic uncertainties in the Homestake measurement were reduced, and to this day the deficit still remains [8].

Gradually it became apparent that the solution to the solar neutrino problem might lie in the properties of the neutrino itself. Among other things, it was realized that the neutrino may exist in a mixed state, analogous to the K^0 system, if it has a small but finite mass [10]. Over astronomical distances or through interactions with dense matter, the mixing could result in the “disappearance” of the original neutrino flavor produced by the sun and the “appearance” of other flavors [11]. Such disappearance of electron neutrinos could explain the flux deficit observed in the Homestake experiment. Although the Standard Model assumes that the neutrino is massless, grand unified theories naturally accommodate small neutrino masses through the “seesaw mechanism” [9], described in Section 1.2. With the success of big-bang nucleosynthesis [12], the discovery of the cosmic microwave background [13], and later the growing evidence for dark matter [4], it became apparent that neutrino mass also has implications for early cosmology and the matter balance of the universe. This prompted an all-out experimental assault on neutrino mass, which in the last few years has crescendoed into a golden age for experimental neutrino physics, in which the subject of this thesis, the KamLAND experiment, has played a seminal role.

1.2 Neutrino Mass

The Standard Model contains 3 left-handed neutrinos ν_{lL} ($l = e, \mu, \tau$), with no mass term. To generalize the theory, we add new fields ν_{iR} corresponding to possibly heavy

right-handed neutrinos. In a CPT-invariant gauge field theory containing such fields, the mass term in the Lagrangian can appear in one of two forms, Dirac or Majorana. The former may be written as

$$-L^D = \bar{\nu}_{iR} m_{ij}^D \nu_{jL} + h.c., \quad (1.1)$$

analogous to the mass terms of the other Standard Model (Dirac) leptons. L^D conserves total lepton number but not necessarily individual lepton number, depending on whether m^D has finite off-diagonal terms. The Majorana mass term has the form

$$-L^M = \frac{1}{2}(\bar{\nu}_{kL}^c m_{kjl}^M \nu_{jL} + h.c.) + \frac{1}{2}(\bar{\nu}_{iR}^c m_{ijR}^M \nu_{jR} + h.c.), \quad (1.2)$$

where ν^c denotes the charge-conjugate of ν . L^M allows transitions $\nu \leftrightarrow \bar{\nu}$, as well as ν - ν (or $\bar{\nu}$ - $\bar{\nu}$) creation and annihilation. Total lepton number changes by two units in these interactions. L^M is allowed only if neutrinos have no additive conserved charges, as is the case for electrically charged leptons and quarks. The possibility of Majorana neutrinos has motivated searches for total lepton number violating processes, such as neutrinoless double- β decay [14].

The most general Lagrangian includes both Dirac and Majorana mass terms, which may be written as

$$-L^M - L^D = \frac{1}{2}\bar{\nu}^c M \nu + h.c., \quad (1.3)$$

where $\nu = \begin{pmatrix} \nu_L \\ \nu_R^c \end{pmatrix}$, and $M = \begin{pmatrix} m_L^M & (m^D)^T \\ m^D & m_R^M \end{pmatrix}$ (the flavor indices will be suppressed for the remainder of this section). The relative scales of m^D , m_L^M , and m_R^M are unknown. However, the most phenomenologically tantalizing situation is the case where the scale of m^D is roughly the same order as the other Standard Model lepton masses, m_L^M is zero or is negligibly small, and m_R^M is at a scale much larger than m^D , corresponding to the scale of new physics. In this case, diagonalization of M leads to

two classes of eigenvectors, three of the form

$$\nu_1 \approx (\nu_L - \nu_L^c) + \frac{m^D}{m_R^M}(\nu_R - \nu_R^c) \quad (1.4)$$

with eigenmasses $m_1 \approx \frac{(m^D)^2}{m_R^M}$, and the rest with eigenmasses $m_2 \approx m_R^M$, with the form

$$\nu_2 \approx (\nu_R + \nu_R^c) + \frac{m^D}{m_R^M}(\nu_L + \nu_L^c). \quad (1.5)$$

The ν_1 correspond to the three Standard Model neutrinos. Their left-handedness derives from the suppression of the right-handed components by the term $\frac{m^D}{m_R^M}$. Their lightness falls naturally from the approximate relation $m_1 m_R^M \approx (m^D)^2$: ratcheting up m_R^M naturally drives m_1 down, a behavior for which this relationship has come to be referred to as the “seesaw mechanism”. The effective mass term for the ν_1 has the form of L^M .

Another phenomenologically viable situation occurs when m_R^M is similar to or below the electroweak scale, in which case there may be more than three light neutrinos with Majorana mass terms. However, measurements of the invisible width of the Z boson by the LEP experiments have determined that exactly three light neutrinos participate in the weak interactions [4], so any additional light neutrinos must be “sterile”, i.e. have no weak couplings. Finally, if $m_R^M = 0$, there are only three neutrinos, all with Dirac mass terms. For the cases of both sterile and Dirac neutrinos, the lightness of the Standard Model neutrinos, some 5 or more orders of magnitude smaller than all other known leptons, has no natural explanation in terms of new, high energy physical scales.

1.3 Neutrino Oscillation

Neutrinos come in three flavors, electron-, mu-, and tau-neutrinos (ν_e , ν_μ , and ν_τ). Neutrino oscillation is a flavor transformation, in which a ν_e , for example, may be converted into a ν_μ or a ν_τ . Thus an experiment measuring a flux of ν_e 's may find that some of them have “disappeared” on their journey to the detector, having oscillated

into undetected flavors. Note that neutrino oscillation does not conserve individual lepton number, while total lepton number is still conserved.

Neutrino oscillation can only occur if neutrinos have mass. Neutrinos are produced and detected in weak interactions, which couple to the weak eigenstates ν_l , where $l = e, \mu, \tau$. For massive neutrinos, the weak eigenstates may be expressed as a linear combination of three mass eigenstates ν_i , $i = 1, 2, 3$, with mass m_i :

$$\nu_l = \sum_i U_{li} \nu_i. \quad (1.6)$$

U is a 3×3 unitary mixing matrix and is analogous to the CKM matrix in the quark sector.

As an ultra-relativistic neutrino propagates through vacuum with energy E , the phase of each mass eigenstate will change at different rates according to the Schrodinger equation,

$$\nu(L) = \sum_i U_{li} e^{-i \frac{m_i^2 L}{2E}} \nu_i, \quad (1.7)$$

where $L = 0$ corresponds to the emission of the neutrino in weak eigenstate ν_l . At distances $L > 0$ the varying phases of the different mass eigenstates will rotate the neutrino into an admixture of weak eigenstates. The probability of detecting flavor $\nu_{l'}$ is found to vary with L as

$$P_{\nu_l \rightarrow \nu_{l'}} = \sum_{i,j} U_{li} U_{l'i}^* U_{lj}^* U_{l'j} e^{-i \frac{\Delta m_{ij}^2 L}{2E}}, \quad (1.8)$$

where $\Delta m_{ij}^2 = m_i^2 - m_j^2$. Note that $P_{\nu_l \rightarrow \nu_{l' \neq l}} = 0$ if all of the m_i are zero or equal, or if $U = 1$.

There are several subtleties in the derivation of Equation 1.8, particularly in the assumptions of the ultra-relativistic limit and the coherence of Equation 1.6. However, these subtleties are irrelevant for the experiment presented in this thesis, as well as for the experiments discussed in Section 1.5. In-depth formulations of the oscillation probability function starting from a wave-packet description of the neutrino state may be found in [15].

In order to get a feel for Equation 1.8, consider for the moment only two neutrino flavors, ν_e and, say, ν_x . In this case, \mathbf{U} is a 2×2 unitary matrix and is parameterized by a single mixing angle θ , the angle in eigenspace between ν_e and ν_1 :

$$\mathbf{U} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}. \quad (1.9)$$

Examining Equation 1.7, we see that the relative phase of the two mass eigenstates changes by 2π over a distance

$$L_{osc} = \frac{4\pi E}{\Delta m^2}. \quad (1.10)$$

In terms of the two parameters θ and L_{osc} , the probability for a neutrino emitted as ν_e to be detected as ν_e is then found from Equation 1.8 to be

$$P_{\nu_e \rightarrow \nu_e} = 1 - \sin^2 2\theta \sin^2 \left(\frac{\pi L}{L_{osc}} \right). \quad (1.11)$$

This is an oscillating function of L with period $2L_{osc}$ and amplitude determined by the size of the mixing angle θ . As will be shown later, this two-flavor approximation adequately describes the disappearance of electron-neutrinos and antineutrinos when the baseline L is very large.

1.4 Matter Effects

Equation 1.11 describes two-flavor neutrino oscillations in vacuum. The situation becomes slightly more complicated for neutrinos propagating through matter. As first recognized by Wolfenstein[16], Mikheyev, and Smirnov[17], while all three weak states participate in neutral-current interactions with normal matter, only electron neutrinos have additional charge-current interactions with the electrons in the material being traversed. This modifies the phase of the ν_e -component relative to the other components by 2π over a distance L_0 given by

$$L_0 = \frac{2\pi}{\sqrt{2}G_F N_e}, \quad (1.12)$$

where G_F is the Fermi constant, and N_e is the number density of electrons in the material. This phenomenon, called the MSW effect after the three authors listed above, alters the two flavor oscillation probability (Equation 1.11) as follows:

$$P_{\nu_e \rightarrow \nu_e} = 1 - \sin^2(2\theta_m) \sin^2\left(\frac{\pi L}{L_m}\right) \quad (1.13)$$

$$\tan 2\theta_m \equiv \tan 2\theta \left(1 + \frac{L_{osc}}{L_0} \sec 2\theta\right) \quad (1.14)$$

$$L_m \equiv L_{osc} \left[1 + 2 \cos 2\theta \frac{L_{osc}}{L_0} + \left(\frac{L_{osc}}{L_0}\right)^2\right]^{-\frac{1}{2}} \quad (1.15)$$

Note that L_0 appears in the equations defining θ_m and L_m only in the fraction L_{osc}/L_0 . In the low-density limit, $L_0 \rightarrow \infty$, and θ_m and L_m approach the vacuum oscillation values.

1.5 Neutrino Oscillation Experiments ¹

1.5.1 Solar Neutrinos

The largest and oldest body of evidence for neutrino oscillation comes from solar neutrino experiments like Davis' Homestake experiment described in Section 1.1. Three experiments, GALLEX [19], GNO [20], and SAGE [21], detected solar neutrinos using gallium via the interaction ${}^{71}\text{Ga} + \nu_e \rightarrow {}^{71}\text{Ge} + e^-$ (threshold = 233 keV). With their lower threshold, the gallium experiments sample more of the solar neutrino spectrum, illustrated in Figure 1.1 [22], and in particular the lower energy pp neutrinos emitted by the reaction $4p \rightarrow {}^4\text{He} + 2e^+ + 2\nu_e$ that is responsible for the majority of the energy production of the sun. As in the case of Homestake, the gallium experiments determined the number of incident neutrinos by measuring the radioactivity of the unstable ${}^{71}\text{Ge}$ daughter nuclei. Once again a deficit relative to the SSM was detected, although the gallium experiments saw slightly more than half of the expected number

¹For a recent review of experimental progress in neutrino oscillation physics and massive neutrino physics in general, see [18] and [14].




Figure 1.1: Solar neutrino energy spectrum predicted by the SSM. The uncertainties in the fluxes from the various reactions are noted. The thresholds for the different classes of solar neutrino experiments are shown at the top of the figure. [22]

of neutrinos, while Homestake saw roughly one third. A comparison of the measured ratios to SSM predictions is shown in Figure 1.2 [23].

The other three experiments with published results on solar neutrinos, Kamiokande [24], its successor Super-Kamiokande [25], and SNO [26], detect solar neutrinos with a different technique, utilizing Cherenkov emission in large detectors filled with water. This technique gives a directional, real-time measurement of the differential neutrino flux. Kamiokande and Super-Kamiokande, detectors near Kamioka, Japan, consist of large caverns lined with photomultiplier tubes and filled with 5 and 50 kton, respectively, of ultra-pure water. These experiments detect the Cherenkov emission of the scattered electron in the interaction $\nu + e^- \rightarrow \nu + e^-$. The analysis thresholds for the experiments were 7 and 5 MeV, respectively. While all three neutrino flavors participate in this interaction, the cross-section for ν_e is roughly 6 times larger than that for ν_μ and ν_τ . The detected correlation between the direction of the ν signal and that of




Figure 1.2: Comparison of measured solar neutrino rates to SSM predictions. [23]

the sun by Kamiokande-II [27] marked the first direct evidence that the sun indeed emits the neutrinos detected in solar neutrino experiments. Both of the Kamioka experiments detect roughly half of the SSM prediction, as shown in Figure 1.2. Neither experiment observes a significant day-night asymmetry in the flux, implying that matter effects for neutrinos passing through the earth on their way to the detector at night do not alter significantly the oscillation probability. The Kamioka experiments also do not observe a significant energy dependence of the suppression of the solar neutrino flux relative to SSM predictions.

SNO, located deep underground in Sudbury, Canada, detects the same electron-scattering (ES) interaction as the Kamioka experiments using 1 kton of heavy water suspended in a spherical acrylic vessel lined with photomultiplier tubes. The presence of deuteron (d) allows for the detection of neutrinos using two additional interactions: the charge-current (CC) interaction $\nu_e + d \rightarrow p + p + e^-$, and the neutral-current (NC) interaction $\nu + d \rightarrow p + n + \nu$. The neutron in the final state of the NC interaction is detected via the 6.25 MeV γ released when the neutron is captured by another

deuterium atom. During a second phase of the experiment, two tons of salt (NaCl) were mixed in with the heavy water. Cl has a larger neutron capture cross-section, and emits a higher-energy γ , at 8.6 MeV; the combination of these two factors greatly enhances the neutron detection efficiency. Since all three neutrino flavors participate in the NC interaction equally, its measurement by SNO represented the first flavor-independent detection of solar neutrinos, and for the first time experiment and theory were in beautiful agreement (see the right-most bar in Figure 1.2). Meanwhile, the CC and ES interactions still showed the suppression observed by all ν_e -dominant detection channels. As in the Kamioka experiments, no significant day-night asymmetry or spectral distortion was observed.

To interpret these experimental results in terms of neutrino oscillations, it must be kept in mind that the density at the center of the sun where most of the neutrinos are created is very high, with $L_0 \approx 200$ km. Except for very small values of Δm^2 , matter effects are so significant that, after the neutrinos emerge from the sun, subsequent vacuum oscillations do not alter the distribution between the three flavors appreciably. The complicated relationships between the matter and vacuum oscillation parameters expressed in Equations 1.14 and 1.15 result in non-trivial regions of oscillation parameter values that accommodate a particular experiment. Hence solar experiments individually allow values of $\tan^2 \theta$ and Δm^2 over many orders of magnitude, generating the complicated allowed regions shown in Figure 1.3 [28]. This figure does not include the results from the salt phase of SNO. The overlap of the allowed regions traditionally yielded 4 distinct oscillation “solutions”. The solution at the lowest values of Δm^2 corresponds to vacuum oscillation, and does not require any significant flavor transformation due to the MSW effect. The other three solutions do, one at $\Delta m^2 \approx 10^{-5}$ eV² with a small mixing angle (SMA), one at about the same value of Δm^2 but with a large mixing angle (LMA), and one at lower values of $\Delta m^2 \approx 10^{-7}$ called simply “LOW”. Adding the SNO salt result gives the overlap shown in Figure 1.4 [26], leaving only the LMA solution, centered at $\Delta m_{sol}^2 = 6.5 \times 10^{-5}$ eV² and $\tan^2 \theta_{sol} = 0.40$.

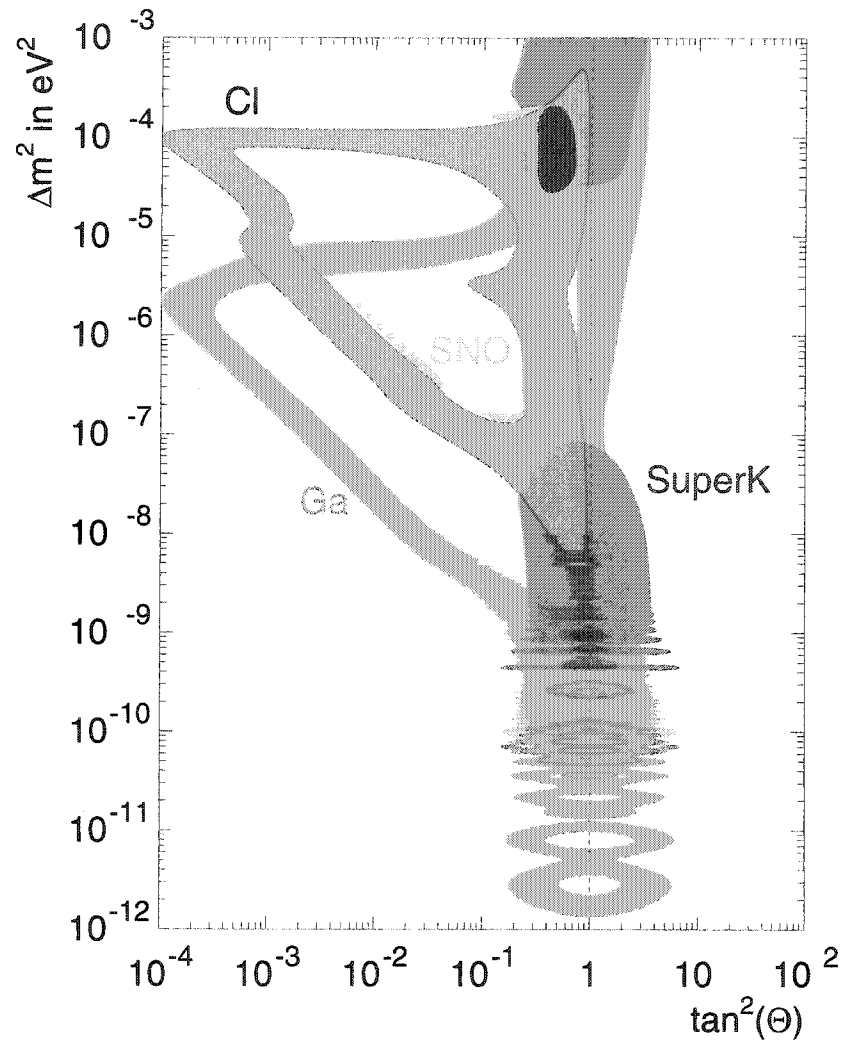


Figure 1.3: Allowed regions for solar neutrino mixing parameters prior to the SNO salt-phase results, divided according to experimental category. The red patch corresponds to the LMA MSW solution. [28]

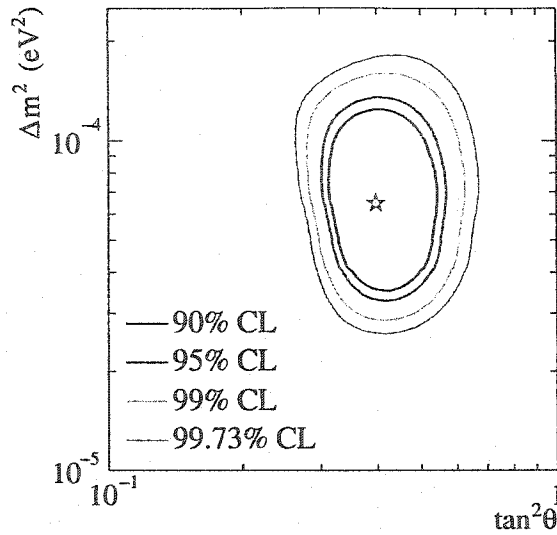


Figure 1.4: Allowed region for neutrino mixing parameters including all experimental solar neutrino results to date. This region corresponds to the LMA MSW solution shown as the red patch in Figure 1.3. The best-fit point is at $\Delta m_{sol}^2 = 6.5 \times 10^{-5}$ eV², $\tan^2 \theta_{sol} = 0.40$. [26].

1.5.2 Atmospheric Neutrinos

While the results of the seven solar neutrino experiments discussed in Section 1.5.1 are described elegantly by neutrino oscillation, the first clear evidence of the phenomenon in a single experiment came from measurements of atmospheric neutrinos. Atmospheric neutrinos are produced in the decay chain of charged pions generated in collisions of cosmic rays with the earth's upper atmosphere. The pions decay via $\pi^\pm \rightarrow \mu^\pm + \nu_\mu$; the daughter muons decay via $\mu^\pm \rightarrow e^\pm + \nu_e + \nu_\mu$. Hence one expects roughly two ν_μ 's for every ν_e , a phenomenology that has been verified by detailed cosmic ray Monte Carlo simulations.

Measurements of the ratio of ν_μ to ν_e in the atmospheric flux were made by the water Cherenkov detectors Kamiokande [29] and IMB [30], for which atmospheric neutrinos created backgrounds to nucleon decay searches. These two experiments found roughly the same number of ν_μ as ν_e , a result that was initially considered to be an anomaly [31]. The Soudan experiment later also reported a ν_μ -to- ν_e ratio of

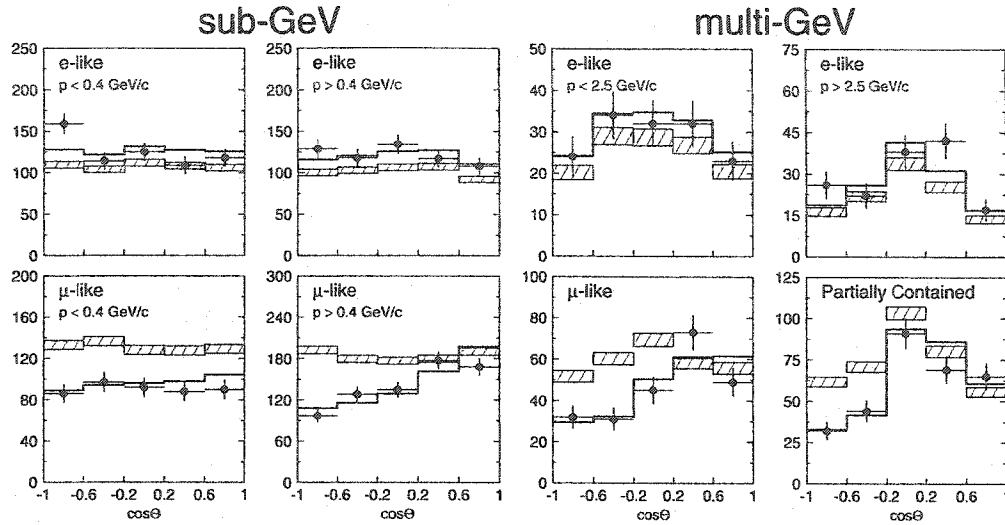


Figure 1.5: Zenith angle distribution of atmospheric neutrino candidates in Super-Kamiokande (data points). The Monte Carlo no-oscillation expectation is shown as the hatched boxes; the solid line shows the best-fit oscillation solution. [33]

about 1 using a different experimental technique [32].

The drastically larger exposure of the Super-Kamiokande experiment allowed for a high-statistics measurement of the zenith angle dependence of the atmospheric ν_μ and ν_e fluxes [33]. The distribution for ν_e agrees reasonably well with Monte Carlo calculations, but the upward ν_μ flux, which travels an extra distance of the earth's diameter relative to the downward flux, exhibits the pronounced deficit shown in Figure 1.5. The deficit and its variation with zenith angle is explained beautifully by ν_μ - ν_τ oscillations, the case for which has only strengthened with updated results [34]. The confidence levels in Δm^2 - $\sin^2 2\theta$ parameter space for the neutrino oscillation fit to the Super-Kamiokande data is shown in Figure 1.6. The best fit point is at $\Delta m_{atm}^2 = 2.4 \times 10^{-3} \text{ eV}^2$, $\sin^2 2\theta_{atm} = 1.0$.




Figure 1.6: Inclusion region for atmospheric neutrino oscillation. The best fit point is at $\Delta m_{atm}^2 = 2.4 \times 10^{-3} \text{ eV}^2$, $\sin^2 2\theta_{atm} = 1.0$. [34]

1.5.3 Reactor Antineutrinos ¹

Reactor $\bar{\nu}_e$ disappearance experiments provide a complementary experimental program to the solar and atmospheric neutrino experiments described in Sections 1.5.1 and 1.5.2. First, these are “laboratory-style” experiments, in which one is in control of both the source and the detector: reactors can be turned on and off, they have been simulated in detail, and calculations of their $\bar{\nu}_e$ emission can be accurately cross-checked with high-statistics measurements (see Chapter 7). Second, for Δm_{sol}^2 , matter effects in rock are negligible for the few-MeV neutrinos emitted by reactors ($L_0 \approx 10^4 \text{ km}$), simplifying the physics involved to mere vacuum oscillations. In addition, while solar experiments detect neutrinos, reactor-based experiments detect *antineutrinos*, typically via observation in scintillator of the inverse β interaction. The different detection techniques required for these particles decouples the systematics between the two experimental classes entirely.

Reactor $\bar{\nu}_e$ experiments with source-detector distances of up to $\sim 1 \text{ km}$ have been performed prior to KamLAND with no significant disappearance detected. In this

¹For an in-depth review of reactor neutrino physics prior to KamLAND, see [35].

light, early experiments at distances ≤ 100 m performed at Grenoble [36], Goesgen [37], Rovno [38], Krasnoyarsk [39], Bugey [40], and Savannah River [41] may be viewed as consistency checks that the reactor $\bar{\nu}_e$ signal can indeed be calculated accurately. Reactor spectrum models have been compared in detail in the high-statistics multiple baseline experiment Bugey 3, in which the overall spectral uncertainty was limited to 1.4% [42]. The latest generation of experiments at distances of 1 km, Palo Verde [43, 44] and CHOOZ [45], searched for $\bar{\nu}_e$ disappearance for $\Delta m^2 \approx \Delta m_{atm}^2$. These experiments achieved systematic errors on the order of a few percent. The lack of oscillation signal in these experiments indicate that atmospheric oscillations are primarily due to $\nu_\mu \leftrightarrow \nu_\tau$ mixing.

Testing solar oscillation parameters with a reactor-based experiment requires baselines at or exceeding ~ 100 km. However, since reactors emit neutrinos isotropically, the signal decreases inversely with the square of the distance between the reactor and the detector. Maintaining a measurable flux of neutrinos in such a long-baseline experiment requires either a very large detector, a very large reactor, or both.

1.5.4 LSND

One last experiment must be mentioned in any discussion of neutrino oscillation, the Los Alamos Liquid Scintillator Neutrino Detector (LSND) [46]. This was an accelerator-based experiment, in which the neutrino source was the beam dump of an 800 MeV proton beam. Such a beam creates copious amount of charged pions, of which there are 8 times as many π^+ as π^- . Only about 5% of the π^- decay in flight; they are highly absorbed due to their large capture probability on nuclei. Moreover, of the μ^- that are produced by π^- decays, roughly 90% get captured via reactions of the form $N + \mu^- \rightarrow^{12} N'$, and no neutrinos are emitted. Hence the predominant source of neutrinos is the decay of the π^+ , which yields ν_μ , $\bar{\nu}_\mu$, and ν_e (the latter two are emitted in the decay of the daughter μ^+). 30 m from the beam stop, a 160 ton liquid scintillator detector searched for the appearance of $\bar{\nu}_e$ via inverse β -decay on protons. A significant excess of $\bar{\nu}_e$'s above backgrounds was indeed detected. The spectrum of the events is plotted in Figure 1.7 [46], along with the contributions from




Figure 1.7: Energy spectrum of $\bar{\nu}_e$ candidates in LSND. The lower histogram shows the expected background distribution. The middle histogram shows the contributions from $\bar{\nu}_e$ contamination in the beam. The upper histogram gives the best-fit oscillation expectation. [46]

backgrounds (lower histogram) and $\bar{\nu}_e$ contamination in the beam (middle histogram). The best-fit oscillation expectation is shown in the upper histogram. Figure 1.8 [48] shows the confidence levels for the mixing parameters derived from the oscillation fit to the LSND data, allowing values of Δm_{LSND}^2 in the range 0.1 - 10 eV².

A similar experiment at a shorter baseline, KARMEN [47], detected no $\bar{\nu}_e$ appearance; its contours exclude much of the LSND solution, but do not have the sensitivity to fully contradict the LSND results. The Bugey reactor experiment also limits the LSND allowed region to smaller values of $\sin^2 2\theta$. The LSND result has yet to be confirmed by an independent experiment. A new accelerator-based $\bar{\nu}_e$ -appearance experiment, MiniBooNE [48], is currently underway and taking data. With projected sensitivity covering the entire LSND allowed region (see Figure 1.8), MiniBooNE is poised to either verify or reject, once and for all, the LSND result.

If the LSND result is indeed verified by MiniBooNE, the interpretation in terms of neutrino mixing becomes a little awkward. The three known neutrino flavors can

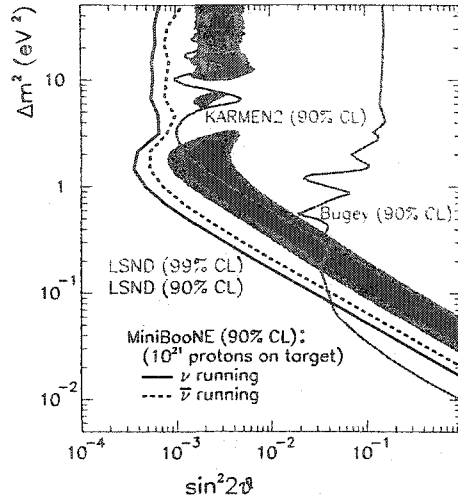


Figure 1.8: Inclusion region for neutrino oscillations consistent with the LSND $\bar{\nu}_e$ appearance signal, and exclusion regions for the KARMEN and Bugey experiments. The black lines show MiniBooNE's expected sensitivity. [48]

give rise to at most two independent mass-differences, so this third region, incompatible with both Δm_{sol}^2 and Δm_{atm}^2 , poses somewhat of a theoretical conundrum. The resolution would require either an additional sterile light neutrino flavor, or the breaking of CPT invariance, so that ν - and $\bar{\nu}$ masses are not equivalent. Additionally, Δm_{LSND}^2 implies the existence of at least one mass eigenstate with $m > 0.4$ eV, at odds with recent studies of the cosmic microwave background [49]. Since the LSND result has yet to be independently verified, it is common to ignore Δm_{LSND}^2 and its complicated theoretical implications when analyzing the larger picture of neutrino oscillation.

1.6 Three-Flavor Oscillations

Until now we have discussed neutrino oscillation results primarily in terms of the two-flavor description leading to Equation 1.11. With experimental parameters in hand, we conclude this introduction with a brief argument for the appropriateness of

this simplification, particularly with regard to the analysis presented in Chapter 8.

When all three known active neutrino flavors are included, \mathbf{U} may be parameterized as follows:

$$\mathbf{U} = \mathbf{U}_{sol} \times \mathbf{U}_{atm} \times \mathbf{U}_{13} \times \mathbf{U}_M, \quad (1.16)$$

$$\mathbf{U}_{sol} = \begin{pmatrix} \cos \theta_{12} & \sin \theta_{12} & 0 \\ -\sin \theta_{12} & \cos \theta_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \mathbf{U}_{atm} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_{23} & \sin \theta_{23} \\ 0 & -\sin \theta_{23} & \cos \theta_{23} \end{pmatrix}$$

$$\mathbf{U}_{13} = \begin{pmatrix} \cos \theta_{13} & 0 & e^{-i\delta} \sin \theta_{13} \\ 0 & 1 & 0 \\ -e^{-i\delta} \sin \theta_{13} & 0 & \cos \theta_{13} \end{pmatrix} \quad \mathbf{U}_M = \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{-i\alpha/2} & 0 \\ 0 & 0 & e^{-i(\alpha/2+\beta)} \end{pmatrix}. \quad (1.17)$$

The six free parameters comprise three mixing angles, θ_{12} , θ_{23} , and θ_{13} , one CP-violating phase δ , and two Majorana phases, α and β . With three masses in the theory there are two independent values of Δm_{ij}^2 . The values of $\tan^2 \theta_{sol}$ and Δm_{sol}^2 measured in solar neutrino experiments correspond to $\tan^2 \theta_{12}$ and Δm_{12}^2 , respectively. Likewise, the parameters measured in atmospheric neutrino experiments correspond to $\tan^2 \theta_{23}$ and Δm_{23}^2 .

The experimental result that $\Delta m_{sol}^2 \ll \Delta m_{atm}^2$ implies that $L_{12} \gg L_{23}$ (recall Equation 1.10). Under this approximation, the oscillation probability (Equation 1.8) simplifies greatly. In particular, for distances much larger than L_{23} , the probability for ν_e to be detected as ν_e is

$$P_{\nu_e \rightarrow \nu_e} \approx \sin^4 \theta_{13} + \cos^4 \theta_{13} \left[1 - \sin^2 2\theta_{12} \sin^2 \left(\frac{\pi L}{L_{12}} \right) \right]. \quad (1.18)$$

Assuming CPT invariance so that $P_{\bar{\nu}_e \rightarrow \bar{\nu}_e} = P_{\nu_e \rightarrow \nu_e}$, the failure to observe $\bar{\nu}_e$ disappearance at previous reactor antineutrino experiments with baselines up to 1 km limits $\sin^2 \theta_{13}$ to be quite small, less than 0.03 [45], in which case one may further approximate

$$P_{\nu_e \rightarrow \nu_e} \approx 1 - \sin^2 2\theta_{12} \sin^2 \left(\frac{\pi L}{L_{12}} \right) \quad (1.19)$$

Note that this is identical to the oscillation probability equation obtained when only two neutrino species are considered, Equation 1.11. This implies that oscillations of ν_e on length scales much larger than L_{23} are governed by the interference of only two of the mass eigenstates, validating the analysis of long-baseline ν_e and $\bar{\nu}_e$ disappearance experiments in terms of two flavors only.

Chapter 2

KamLAND

As explained in Section 1.5.3, detecting reactor $\bar{\nu}_e$'s at distances long enough to be sensitive to the LMA MSW solution to the solar neutrino problem requires either a very large detector or a very powerful reactor. This is illustrated in Figure 2.1, which plots the product of reactor power and detector size versus baseline and the corresponding Δm^2 sensitivity for reactor $\bar{\nu}_e$ experiments [35]. The Kamioka Liquid-scintillator Anti-Neutrino Detector (KamLAND) [50] reaches unprecedented Δm^2 sensitivity by satisfying both requirements simultaneously. First, it is the largest liquid scintillator detector ever built, with 1 kton of active volume. Second, its source consists of not one reactor complex but the entire Japanese nuclear power industry, with a net thermal power output of 200 GW.

Figure 2.2 [51] shows the locations of reactors around Japan and the position of KamLAND underneath Mt. Ikenoyama in Gifu Prefecture. Using the rated thermal powers of each reactor, the distribution of the power flux at KamLAND versus reactor distance is plotted in Figure 2.3. In real-life operations, reactor down-time results in a slightly lower power flux at the KamLAND site. Almost 80% of the power flux comes from reactors between 135 and 215 km, ensuring that the oscillation pattern for $L_{osc} \approx 100$ km is not washed out. Since reactors emit $\bar{\nu}_e$'s with energies of typically a few MeV, this configuration makes KamLAND sensitive to $\Delta m^2 \approx \Delta m_{sol}^2$, placing KamLAND in position to directly observe vacuum oscillations associated with the LMA MSW solution to the solar neutrino problem. However, at such low energies,

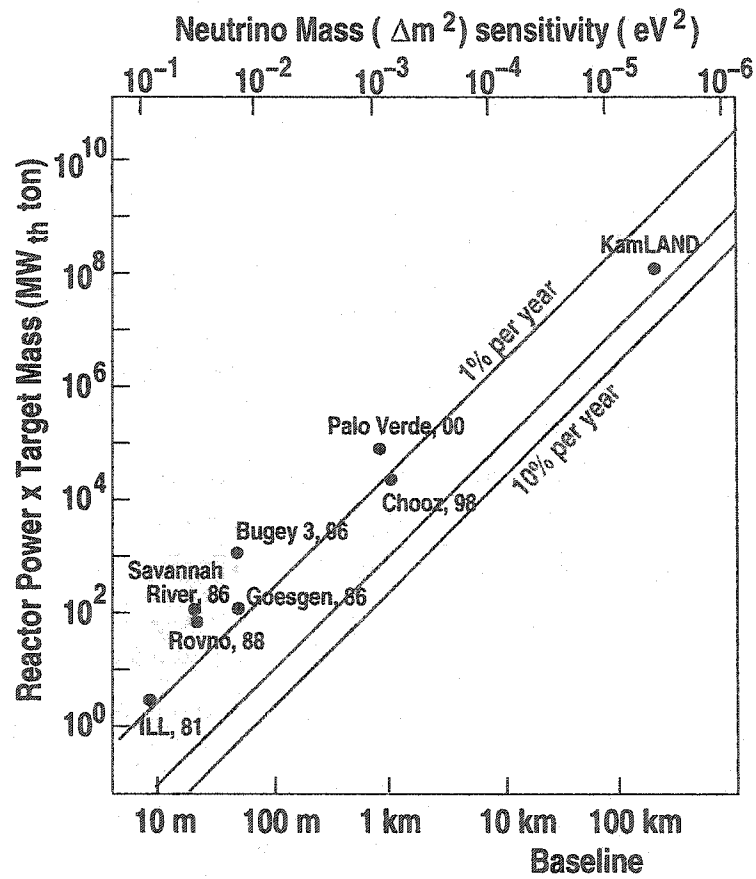


Figure 2.1: Product of reactor power and detector size versus baseline and the corresponding Δm^2 sensitivity for reactor $\bar{\nu}_e$ experiments. Published experiments are labeled by reactor site and year of publication. [35]

one must take care to avoid backgrounds from natural radiation that could potentially drown out the signal. Fortunately the $\bar{\nu}_e$ detection method itself serves as a powerful background rejection, as will be discussed shortly.

Electron antineutrinos are detected at KamLAND via inverse β -decay, in which a $\bar{\nu}_e$ annihilates with a proton (i.e. hydrogen atom) in a mineral oil based liquid scintillator, producing a positron and a neutron: $\bar{\nu}_e + p \rightarrow e^+ + n$. The positron is ejected from the interaction vertex with kinetic energy given by

$$E_{e^+} = E_{\bar{\nu}_e} - E_n - \Delta m_{n-p} - m_e, \quad (2.1)$$

where $E_{\bar{\nu}_e}$ is the $\bar{\nu}_e$ energy, E_n is energy of the neutron, Δm_{n-p} is the mass difference between the neutron and the proton, and m_e is the mass of the positron. Since the proton and neutron are so massive compared to the e^+ and $\bar{\nu}_e$, E_n is small compared to the other terms on the right hand side of the equation. The positron quickly slows down as it deposits its energy in the scintillator by ionizing molecules in the liquid, and then annihilates with a nearby electron, producing two 511 keV γ 's.

Scintillators emit light (i.e. “scintillation”) when ionized by charged particles and penetrating radiation. In KamLAND’s organic liquid scintillator, ionized valence electrons relax quickly to the first excited molecular orbital via internal degradation and collisions with other molecules, processes that do not emit radiation. From the first excited state, there is a high probability for radiative decay to a vibrational excitation of the ground state. The energy of the emitted photons, which is on the order of a few eV, is insufficient to re-excite transitions from the ground state to the first excited state, explaining why scintillators are transparent to their own radiation. For inverse β -decay, ionization from the positron’s energy loss and the two annihilation γ 's occurs so quickly relative to the time scale of scintillation emission that they appear as a single flash promptly following the moment at which the interaction occurs. The scintillation light is emitted isotropically, and its brightness depends on the total amount of ionization, which is roughly proportional to the initial kinetic energy of the positron plus the energy of the two annihilation γ 's. Hence by collecting the scintillation light an estimate may be made of the energy of this “prompt” event,

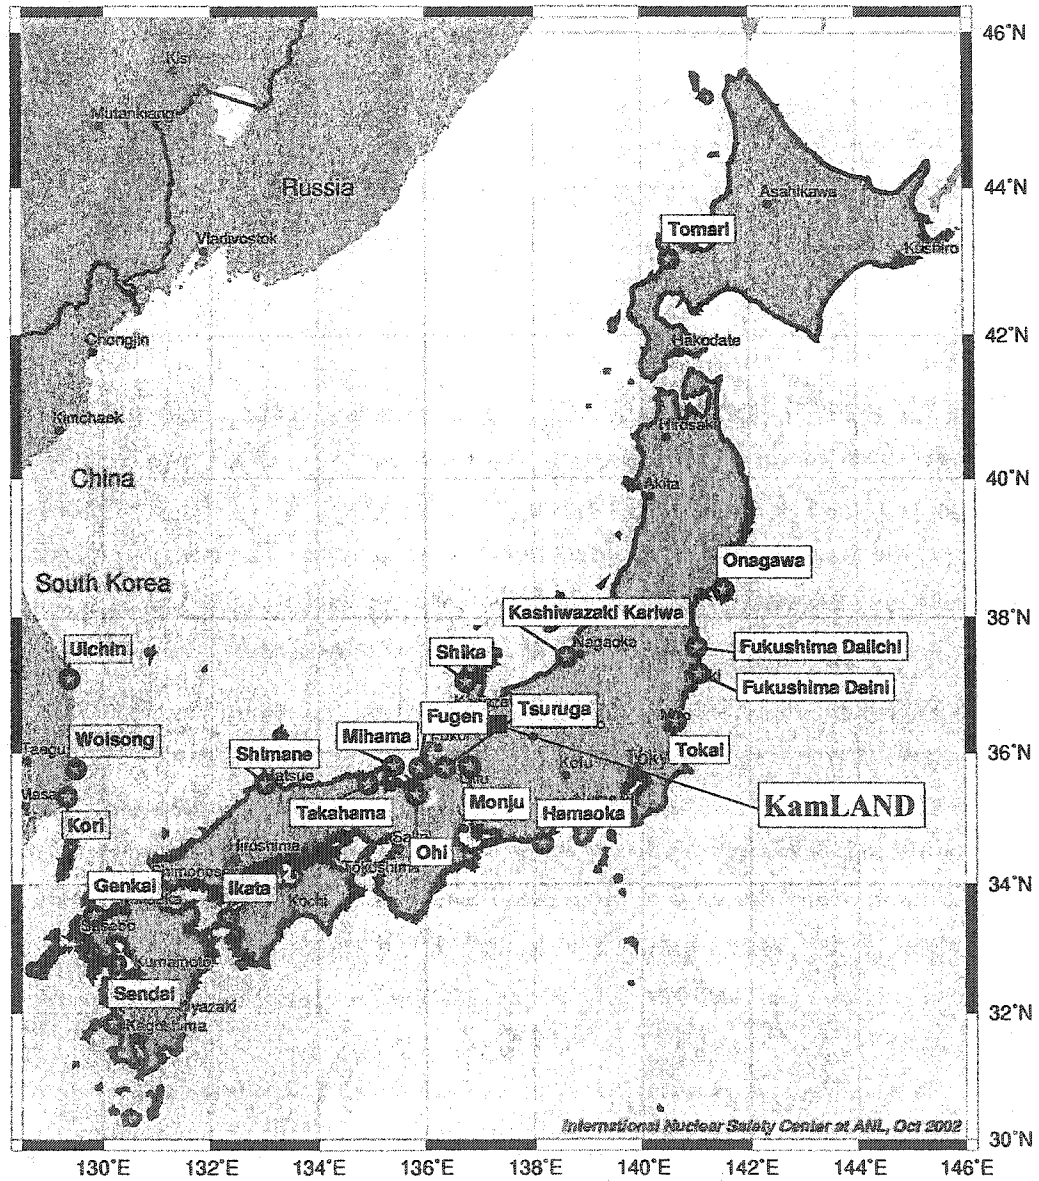


Figure 2.2: The stars mark reactor locations around Japan. The location of KamLAND is marked by the blue square. [51]




Figure 2.3: Baseline distribution of the rated thermal power flux ($\phi_{P_{th}}$) at KamLAND from Japanese and Korean reactors. Nearly 80% of the total flux comes from the peak centered at 180 km. Korean reactors comprise roughly 3% of the total flux, while reactors from the rest of the world contribute at the 1% level.

E_p , which by Equation 2.1 gives a measure of the incoming $\bar{\nu}_e$ energy. On average, $E_{\bar{\nu}_e}$ is given by

$$\bar{E}_{\bar{\nu}_e} = E_p + E_{\Delta}, \quad (2.2)$$

where $E_{\Delta} = \bar{E}_n + \Delta m_{n-p} - m_e \approx 0.8$ MeV (the bar denotes an average). The spread of $E_{\bar{\nu}_e}$ is narrow due to the smallness of E_n . In addition to the brightness of the scintillation light, if the arrival time of the spherical scintillation wavefront is measured at various positions around the detector, the position of the event may be precisely determined.

The neutron emitted concurrently with the e^+ in the inverse β -decay quickly thermalizes via elastic interactions with protons in the scintillator, and any scintillation emitted by the recoil nuclei in the process contributes negligibly to the brightness of the prompt event. The neutron then undergoes a series of elastic interactions at thermal energies, which do not produce detectable amounts of scintillation light, and is eventually captured by a nucleus in the scintillator. The neutrons are primarily

captured by protons via the interaction $n + p \rightarrow d + \gamma$, where d is a deuteron and the 2.2 MeV γ carries away the mass difference between the initial and final states. The detection of this delayed monoenergetic gamma, typically within a few meters of and a few hundred microseconds following the prompt positron event, provides the extremely powerful background suppression alluded to earlier.

Detection of coincidence $\bar{\nu}_e$ events at KamLAND gives a real-time measurement of the $\bar{\nu}_e$ flux and its energy dependence. The instantaneous differential detection rate of the reactor $\bar{\nu}_e$ signal, $\frac{d^2 N_{\bar{\nu}_e}(E_p, t)}{dE_p dt}$, may be expressed by the following equation:

$$\frac{d^2 N_{\bar{\nu}_e}(E_p, t)}{dE_p dt} = \int_0^\infty dE'_{\bar{\nu}_e} R(E_p, E'_{\bar{\nu}_e}) n_p \sigma(E'_{\bar{\nu}_e}) \epsilon(E'_{\bar{\nu}_e}) \times \sum_i^{\text{reactors}} \frac{I_i(E'_{\bar{\nu}_e}, t)}{4\pi L_i^2} P_{\bar{\nu}_e \rightarrow \bar{\nu}_e}(E'_{\bar{\nu}_e}, L_i; \Delta m^2, \sin^2 2\theta). \quad (2.3)$$

$R(E_p, E'_{\bar{\nu}_e})$ is the detector response function, including both the energy resolution and the conversion to prompt energy, expressed by Equation 2.2. The parameter n_p is the number of target protons, $\sigma(E'_{\bar{\nu}_e})$ is the cross-section for inverse β -decay [52], and $\epsilon(E'_{\bar{\nu}_e})$ is the possibly energy-dependent detection efficiency. $I_i(E'_{\bar{\nu}_e}, t)$ is the instantaneous differential $\bar{\nu}_e$ intensity of the i^{th} reactor with units [$\bar{\nu}_e$'s per unit energy per unit time], and L_i is the reactor's distance from KamLAND. $P_{\bar{\nu}_e \rightarrow \bar{\nu}_e}$ is the $\bar{\nu}_e$ survival probability, given by Equation 1.11¹. In addition to reactor antineutrinos, KamLAND is sensitive to $\bar{\nu}_e$'s from the sun [53], heavy element deposits in the earth [54], and any other significant antineutrino source. However, for energies greater than 3.4 MeV, the background from known $\bar{\nu}_e$ sources other than nuclear reactors is negligible. For this reason the analysis reported here is restricted to $\bar{\nu}_e$ energies above this threshold.

Including the contributions $\frac{d^2 N_{\text{bg},j}(E_p, t)}{dE_p dt}$ from non- $\bar{\nu}_e$ backgrounds that mimic the

¹Technically, $P_{\bar{\nu}_e \rightarrow \bar{\nu}_e} = P_{\nu_e \rightarrow \nu_e}$ assumes CPT invariance; if CPT is broken, $P_{\bar{\nu}_e \rightarrow \bar{\nu}_e}$ has the same form as Equation 1.11, but the parameters Δm^2 and $\sin^2 2\theta$ are different for neutrinos and antineutrinos.

$\bar{\nu}_e$ signal, the total instantaneous differential event rate $\frac{d^2 N(E_p, t)}{dE_p dt}$ at KamLAND is

$$\frac{d^2 N(E_p, t)}{dE_p dt} = \frac{d^2 N_{\bar{\nu}_e}(E_p, t)}{dE_p dt} + \sum_j^{\text{backgrounds}} \frac{d^2 N_{bg,j}(E_p, t)}{dE_p dt} \quad (2.4)$$

Equations 2.3 and 2.4 will be used to extract the best fit values of Δm^2 and $\sin^2 2\theta$ from the KamLAND data. It should be noted that, while the details of the present analysis differ from those published elsewhere [55, 56], the results and conclusions are in good agreement.

Chapter 3

The Detector

KamLAND is located in the cavern that housed the original Kamiokande experiment. The 2700 m.w.e. overburden reduces the cosmic ray flux by a factor of roughly 10^{-5} with respect to the surface flux. A schematic of the detector is shown in Figure 3.1. It may be divided into two distinct regions, a 1 kton liquid scintillator inner detector (ID) and a water Cherenkov outer detector (OD) that serves as a cosmic ray veto. The ID and OD are physically separated by an 18 m diameter spherical stainless steel containment vessel erected in the center of the cavern.

The ID is composed of concentric volumes of liquid scintillator (LS) and non-scintillating mineral oil. The scintillator [57] consists of 80% dodecane and 20% pseudocumene (1,2,4-trimethylbenzene) (by volume), with 1.52 g/liter of PPO (2,5-diphenyloxazole) as a fluor. Measurements of its optical properties give an attenuation length of 20 m, a light yield of 70% anthracene, and a refractive index of 1.44 (at $\lambda = 589$ nm, 15° C). The scintillator density is 0.778 g/cm³ at 11.5° C; it is measured to 0.01% precision, and an additional 0.1% error is assigned for the uncertainty in the temperature. The calculated carbon-to-hydrogen ratio of 1:1.969 was verified by elemental analysis to 2% precision. Combining this ratio with the scintillator density and using the molar masses of carbon and hydrogen gives a proton density of $6.608 \pm 0.006 \times 10^{28}$ m⁻³. A volume of 1200 m³ of scintillator is suspended in a transparent, 13 m diameter spherical balloon made of 135 μ m thick nylon/EVOH (ethylene vinyl alcohol copolymer) composite film. The balloon is reinforced and




Figure 3.1: A schematic diagram of the detector.

anchored to the containment vessel with a network of Kevlar ropes whose tension is monitored by a set of load cells at the top of the detector. A buffer of 1800 m^3 of dodecane and isoparaffin oil fills the space between the balloon and the containment vessel. The density of the buffer oil (BO) is 0.04% lower than that of the LS to maintain the shape of the balloon. The LS and BO were purified during detector filling using water extraction and nitrogen stripping techniques [58].

The scintillator is viewed by 1325 fast 17-inch aperture Hamamatsu photomultiplier tubes (PMT's) custom-designed for KamLAND, and 554 20-inch PMT's inherited from Kamiokande [59]. As shown in Figures 3.2 and 3.3, the PMT's are secured to the inner surface of the containment vessel, facing the center of the detector. The total photocathode coverage of the 17-inch tubes is 22%; when the 20-inch tubes are included, the coverage increases to 34%. The quantum efficiency of the PMT's is about 20% for 340-400 nm light. The enveloping buffer oil shields the scintillator from radioactivity in the PMT glass. A 3.3 mm thick, 16.6 m diameter acrylic sphere



Figure 3.2: A photograph of PMT installation. The workers are, from left to right, K. McKinny (University of Alabama), the author, and J. Ritter (California Institute of Technology). Photo taken by T. Takayama (Tohoku University).

just in front of the PMT's acts as a radon barrier.

A small opening in the top of the tank leads to a chimney and calibration device, allowing an operator to deploy calibration sources along the vertical axis of the detector (z -axis) using a line-and-reel system. KamLAND uses four γ sources (^{203}Hg , ^{68}Ge , ^{65}Zn , and ^{60}Co), one neutron source ($^{241}\text{Am}/^9\text{Be}$). ^{203}Hg undergoes β^- decay to the first excited state of ^{203}Tl . The emitted e^- is captured within the source, while the decay of the $^{203}\text{Tl}^*$ emits a 279 keV γ . ^{68}Ge undergoes β^+ decay to ^{68}Ga ; the positron annihilates inside the source, but the resulting pair of 511 keV γ 's escape, giving a calibration point at the inverse β -decay threshold. ^{65}Zn likewise undergoes β^+ decay to ^{65}Cu , which 50% of the time produces a 1.115 MeV γ . The β^- decay of ^{60}Co to ^{60}Ni emits a pair of 1.333 MeV and 1.173 MeV γ 's. The bombardment of ^9Be by α 's emitted in decays of ^{241}Am can produce free neutrons through several pathways, such




Figure 3.3: A photograph of the inner detector under construction. The workers who can be identified are B.K. Fujikawa (Lawrence Berkeley National Laboratory) handling the PMT on the far left, the author to his right, and J. Ritter (California Institute of Technology) at the center of the photo. Taken by T. Takayama (Tohoku University).

as $\alpha + {}^9\text{Be} \rightarrow {}^{12}\text{C}^{(*)} + n$, $\alpha + {}^9\text{Be} \rightarrow \alpha + {}^8\text{Be} + n$, or $\alpha + {}^9\text{Be} \rightarrow 3\alpha + n$. Captures of the free neutrons give γ lines at 2.2 and 4.95 MeV. The mode involving ${}^{12}\text{C}^{*}$ gives a prompt 4.4 MeV γ in coincidence with the neutron thermalization scintillation.

Three thermometers at the bottom, center, and top of the detector are affixed to a vertical line displaced slightly off the z-axis, out of the way of the calibration device. The ID is hermetically sealed and the space in the chimney above the LS and BO levels is flushed with radon-purified nitrogen gas to prevent oxygenation of the LS.

The OD is divided into four sections, labeled top, upper, lower, and bottom, containing 50, 60, 60, and 55 20-inch PMT's, respectively. To make up for the obstruction by the containment vessel, each section is lined with and physically separated by sheets of highly reflective Tyvek plastic that enhance light collection. The OD is filled with 2600 m³ of ultra-purified water that is filtered and circulated continuously. Several coils of wire looped about the cavern cancel the earth's magnetic field to minimize the Lorentz force on photoelectron paths in the PMT's.

Each PMT or "channel" is connected to two Analog Transient Waveform Digitizers (ATWD's), custom application-specific integrated circuits developed at Lawrence Berkeley National Laboratory. These chips record and measure the PMT pulses, behaving similar to oscilloscopes tracing voltage changes on every channel. They produce 128-sample digital waveforms at a variable sampling frequency, which is set to ~ 0.65 GHz to give approximately 200 ns of data in each waveform. The ATWD's are self-launching, and their low-noise properties allow single-channel discriminator thresholds as low as $\frac{1}{3}$ of a photoelectron (p.e.), giving a single photoelectron detection efficiency of $> 95\%$. The ATWD's record waveforms on three different gains, providing enough dynamic range to measure amplitudes of up to thousands of photoelectrons. Digitization of a pulse requires 30 μs ; by using two chips per PMT, the effect of the dead time is small. It is not insignificant, however, as will be discussed further in Section 5.5. In early versions of the ATWD logic, the second ATWD would be launched if a channel was still over threshold at the end of the digitization window of the first ATWD. This led to the collection of many spurious waveforms during high-energy events creating very large pulses, such as muons. In March 2002, the condition was changed to require the channel to re-cross the discriminator threshold.

The ATWD's and their accompanying electronics are mounted onto KamLAND Front-End Electronics (KAMFEE) boards that are read out by VME buses. There are 10 VME crates, with 20 boards per crate and 12 channels per board. One crate is dedicated to the PMT's in the OD. The data from each crate is read out by its own PC via fiber optic cables. The data acquisition (DAQ) software, KiNOKO, is a networked parallel processing system developed at Tohoku University. It is based on distributed object technology, and coordinates the data collection across the KamLAND DAQ network. In addition to DAQ performance monitoring and remedial online analysis tasks, KiNOKO folds the data streams from the different crates together and writes them to disk. Each ATWD is read independently, so their waveforms are written to disk out of time-order and are sorted off-line.

The ID PMT's have dark rates on the order of tens of kHz. Since each waveform requires 256 B of memory, it is impossible with current network and computer technology to record data from all channels all the time. Instead, data is only taken when a large number of 17-inch PMT's are hit within a very short time period. The trigger module that detects such coincidences and commands the ATWD's to write their waveforms to the data stream were developed at Stanford University. Every 25 ns, each KAMFEE board sends a digital signal to the trigger module specifying the number of channels on the board that exceeded their discriminator thresholds in the past 125 ns. The trigger module calculates the sum of all these signals, called "*nsum*", for the ID and for each OD section every 25 ns, compares it to a threshold, and issues trigger signals to the ATWD's if it deems necessary based on various criteria described below. Trigger information is read out to the data stream by a dedicated PC using KiNOKO.

In order to facilitate the association of waveforms with trigger records downstream, each trigger record and waveform is accompanied by a timestamp, the number of ticks counted by the trigger module's 40 MHz clock between the start of the run and the issuing of the associated trigger signal. The clock is fanned out from the trigger module to all KAMFEE boards so that the waveform timestamps can be applied at the ATWD level. The waveform records contain additionally the launch offset of the ATWD, the number of clock ticks between the launching of the ATWD and the

issuing of the trigger signal. This allows waveforms to be stored in the ATWD's for up to tens of clock ticks waiting for the trigger signal while still retaining the ability to reconstruct exact pulse timings downstream in the analysis. A GPS receiver delivers a 1PPS signal to the trigger module, periodically accompanied by an IRIG B time code. The GPS signals are used to monitor and verify the stability of the clock, to determine precisely the clock frequency, and to synchronize the timestamps with UTC to within 150 ns.

The trigger module design incorporates three Field-Programmable Gate Arrays (FPGA's), with which triggering algorithms of varying complexity can be implemented. KamLAND is typically operated in "prompt-delayed" triggering mode, in which a prompt trigger is issued when the ID $nsum$ ($nsum_{ID}$) exceeds 200, and a delayed trigger is issued when it goes above 120 within 1 ms of a prompt trigger. The prompt-delayed trigger rate is typically about 25 Hz. Trigger-only records, called "history" records, are recorded for up to 8 clock ticks (200 ns) whenever $nsum_{ID}$ exceeds 120. The maximum value of $nsum_{ID}$ in the train of history records accompanying an event, called $nsmx_{ID}$, is highly correlated with the total number of PMT hits, and hence the energy of the event. The distribution of $nsmx_{ID}$ for a typical run is shown in Figure 3.4. Features such as the triggering thresholds, several radio-peaks/shoulders, and the muon rate can be identified in this distribution, making it useful for preliminary analysis and online monitoring.

The top, upper, lower, and bottom OD sections are triggered when their individual $nsums$ exceed 6, 5, 6, and 7, respectively. Series of OD history records are also recorded when these thresholds are exceeded. Whenever the ID or any OD section is triggered, a trigger signal is also sent to the rest of the detector (this was not true for the first few months of data-taking due to a bug in an early version of the trigger module logic). The trigger module is also capable of issuing trigger signals from an external source, such as a signal generator, a laser pulse, or any TTL or NIM signal. Several other triggering schemes are implemented for test runs, calibrations, monitoring, etc., such as pre-scale triggers, and clock and pedestal captures described in Section 4.2.

One last triggering scheme that must be mentioned is "supernova mode". Should

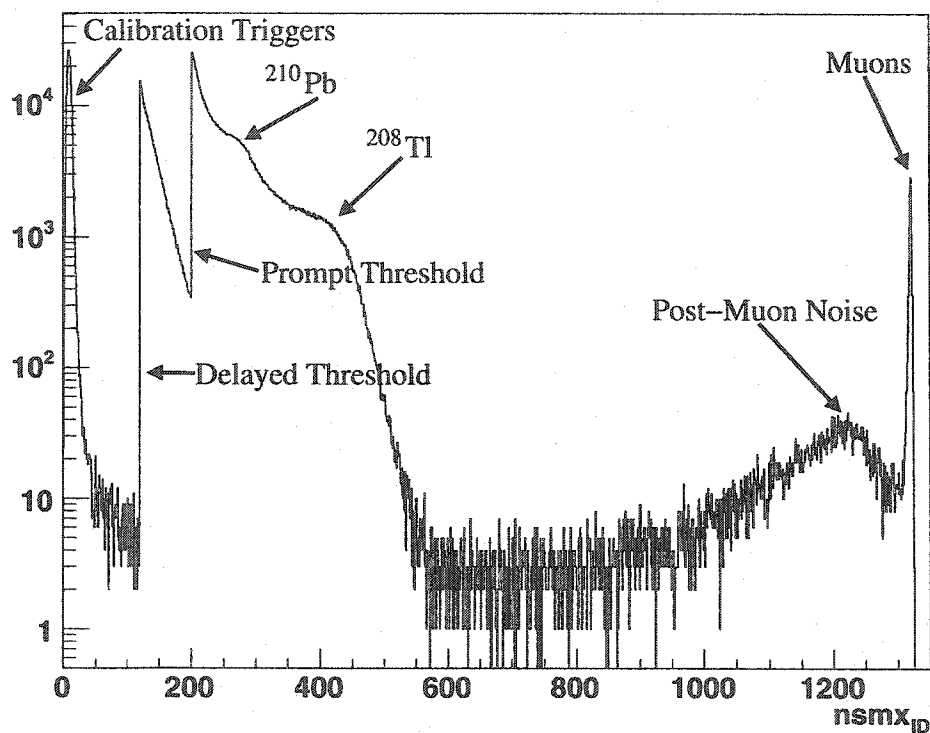


Figure 3.4: Distribution of $nsmx_{ID}$. The vertical lines at $nsmx_{ID} = 120$ and 200 represent the delayed and prompt trigger thresholds, respectively. The peak at very low $nsmx_{ID}$ is due to calibration triggers which are uncorrelated with physics events and therefore gives the minimum bias. The bump and the shoulder at $nsmx_{ID} = 250$ and 400 are due, respectively, to ^{210}Pb in the LS and ^{208}Tl concentrated near the thermometers and the balloon surface. The peak at $nsmx_{ID} = 1325$ is due to muons crossing the LS, while the peak just below 1325 is due to electronics noise following muons.

a supernova occur within our galaxy, KamLAND would likely detect around 300 $\bar{\nu}_e$'s with a mean energy of about 15 MeV. The initial observation rate would be on the order of 100 Hz, and would decay exponentially with a ~ 3 s decay constant [50]. The trigger module watches for supernovae by counting the number of events with $nsmx_{ID} > 772$ in the last 2^{25} clock-ticks (840 ms). If the count reaches 8 events, the trigger module switches to supernova mode for 1 minute, during which the threshold is lowered to 170 in order to increase the efficiency for detecting lower energy proton scatters of neutrinos emitted in the supernova. These events, if they are detectable above backgrounds, are useful for calculating the temperature and luminosity of the supernova [60]. Operators are not allowed to interrupt data taking while the detector is in supernova mode.

Before writing the data to disk, waveforms are first compressed by a factor of ~ 3 using an implementation of lossless Huffman encoding optimized for single-photoelectron waveforms. The compressed, asynchronous waveforms are written to disk along with the trigger data at a rate of about 180 GB per day. Further processing is done off-line. Data is taken 24 hours a day, stopping only for problems, calibration runs, upgrades, testing, or daily run-number changes. Data-taking is monitored by several automated programs and on-site collaboration members. The raw data is copied daily onto tapes, which are then mailed to data libraries in the US and Japan for analysis. The analysis reported here employed the High Performance Storage System and Parallel Distributed Systems Facility at the National Energy Research Scientific Computing Center at Lawrence Berkeley National Laboratory.

Chapter 4

Event Reconstruction

4.1 Event Building

The first step in the event reconstruction process is to sort the asynchronously collected waveforms and trigger records and group them into events. The sorting process is simple in principle since each waveform record contains the timestamp of the trigger signal issued to acquire it. However, during bursts of high activity, data from one ATWD may be recorded in the data stream up to 5 or 10 seconds after the data from another ATWD. Combined with data rates of around 2 MB/s, the memory handling for event building is rather involved. Waveforms arriving more than 5 minutes out-of-order cannot be sorted and are thrown out. The event building software was verified with other less complicated and less memory-intensive codes which simply count waveforms and trigger records for each time stamp. Comparing counts before and after the event building process established that the efficiency for building physics events is greater than 99.999%. The dropped events occur almost exclusively during noisy periods with trigger rates in the kHz range, which are vetoed for the analysis.

4.2 Waveform Analysis

After the waveforms are collected and time sorted into events they are processed by a series of algorithms which search for pulses in the waveforms, extract their arrival

times, and estimate the number of photoelectrons that arrived at each PMT. The times and photon counts are used by event reconstruction software to estimate the positions and energies of events. Described briefly below are the algorithms used in this analysis. Most of these algorithms assume that events create only enough light to produce at most a few photoelectrons per channel. For high energy events such as muon tracks, very large pulses are created in every channel. For some steps in the waveform analysis, a secondary algorithm is implemented that is optimized such large pulses. These algorithms are used only by the muon track fitting software.

4.2.1 Pedestal and Baseline Subtraction

The component of the ATWD that actually captures the PMT pulses is essentially an array of 128 capacitors, each of which stores the height of the incoming pulse at different times. The sampling frequency is tunable, and is set to ~ 0.65 GHz. A raw waveform is simply a listing of the voltages across the capacitors. An example of a raw waveform is shown in the upper trace in Figure 4.1. The appearance of this waveform is affected by two phenomena: a sample-to-sample variation called the “pedestal”, and a low-frequency fluctuation that results in an overall shift of the waveform’s baseline.

In the absence of an input signal, each capacitor rests at a constant base voltage that is different from capacitor to capacitor. These base voltages are the ATWD’s pedestal, and they must be subtracted from the waveform before it can be further processed. The pedestals are measured at the beginning of every run via a series of 50 forced-acquisition triggers in each gain of each ATWD. The waveforms obtained are averaged and subtracted from waveforms subsequently collected in the run. With PMT dark rates on the order of tens of kHz, the chances of a pulse accidentally occurring during a forced-acquisition trigger is small but non-zero. So the pedestal waveforms are scanned for samples deviating more than 4σ from their means; pedestal waveforms containing such deviations are flagged as noise and rejected, and the remaining pedestal waveforms are re-averaged. This method has been found to very efficiently remove pulses and other noise from the pedestals. An example of a pedestal

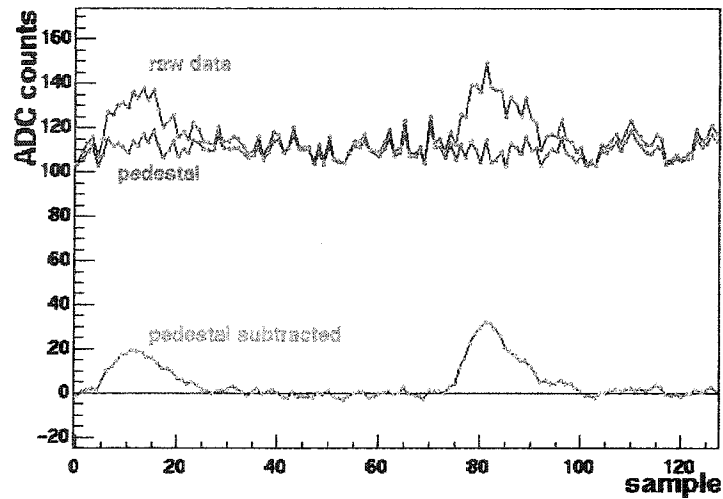


Figure 4.1: A raw waveform with two distinct pulses. The upper trace is the raw data; immediately below it is the pedestal for the ATWD that recorded the pulse. The bottom trace is the pedestal- and baseline-subtracted waveform.

waveform so obtained is shown immediately below the raw data trace in Figure 4.1.

Due to low-frequency variations in the bias on the input signals, the waveforms are found to float at a non-zero baseline, even after the subtraction of the pedestal. The baseline is different for every waveform, and must be determined from the waveform shape and subtracted away for every event. The procedure is to sort the ADC values from lowest to highest and compute their mean. Then samples with the largest deviations from the mean are successively removed from the sorted list, pushing the mean of the remaining samples closer to the flattest region of the waveform. The process is iterated until the deviations of the largest and smallest ADC values from the mean are approximately the same, to within an accuracy $\delta = 0.025$. The last calculated mean is the estimate of the baseline, which is then subtracted from the waveform. Unfortunately, this method fails for very large pulses which do not return to their baselines by the end of the sampling window. For such pulses, a secondary waveform analysis is performed in which the baseline subtraction consists simply of subtracting off the average of the first 10 samples in the waveform.

An example of a pedestal- and baseline-subtracted waveform is shown as the bottom trace in Figure 4.1. This particular waveform happened to require very little baseline subtraction. The apparent high-frequency noise in the raw waveform has been removed by the pedestal subtraction, giving a clean trace in which two single photoelectron pulses can be easily seen.

4.2.2 Pulse Finding

The first step in pulse finding algorithms is typically to smooth the pulses, removing high frequency noise. The waveforms in this analysis were smoothed using the Savitzky-Golay filtering algorithm [61], in which the smoothed value of each data point was calculated based on a k^{th} order polynomial fit to at least the $k + 1$ surrounding data points. Here a 4^{th} order polynomial was fit to the 15 surrounding data points. The Savitzky-Golay method is touted for its preservation of features such as peak height and width that other smoothing procedures tend to distort (see, for example, [62]).

The actual pulse finding algorithm proceeds as follows. First, the waveform is divided up into positive regions, contiguous sections within which all samples of the pedestal- and baseline-subtracted, Savitzky-Golay-smoothed waveform are greater than 0. Each section is considered a candidate pulse. Then the charge of each candidate pulse (the area beneath the waveform) is computed, and candidate pulses are rejected as noise if their charges are less than 15% of the summed charge of all candidate pulses. The summed charge is recomputed after each rejection. For the pulses that survive the 15% cut, the arrival times are determined from quadratic approximations of their peak positions. The times are accurate to within a fraction of a sample, much better than the ~ 2 ns transit time spread of the 17-inch PMT's.

One drawback of this pulse-finding algorithm is that multiple overlapping pulses are counted as a single pulse if the trace remains positive in the overlap region, as shown in Figure 4.2. However, what it lacks in complexity it makes up for in robustness due to its lack of dependence on pulse shape, and downstream analysis components are written to be as insensitive as possible to the algorithm's inability to

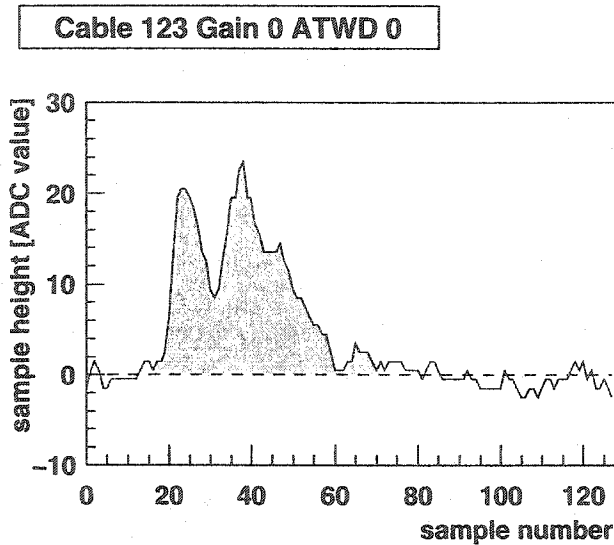


Figure 4.2: An example of a double pulse tagged by the pulse-finding algorithm. Because the trace does not drop below 0 ADC counts between the pulses, they are counted as a single pulse. The area of the pulse is the region shaded yellow.

distinguish multiple pulses

Another aspect in which this algorithm is found to be less-than-optimal is in the time estimation of very large pulses. Due to AC coupling, the peak position of such pulses depends essentially on their size, with larger pulses peaking later. It is for this reason that the pulse finding algorithm is insufficient for the needs of muon track reconstruction algorithms. The timing of a large pulse is better estimated by locating the time at which it begins to rise above its baseline. This is done by setting a simple threshold of 50 ADC counts and using the time at which the rising edge crosses it.

4.2.3 Sampling Rate Measurement

Converting the units of the pulse arrival times from [samples] to [nanoseconds] requires precise knowledge of the sampling frequency, which may vary from ATWD to ATWD. To determine the sampling frequency precisely, the 40 MHz clock signal sent to each KAMFEE board from the trigger module is digitized ~ 50 times by each ATWD at the beginning of every run. The waveforms are pedestal-subtracted and averaged,

and a Fourier transformation is performed. The largest peak in the power spectrum is identified as the 40 MHz line, and the weighted mean position of the the 5 bins to either side of the peak gives the corresponding frequency in [samples⁻¹]. The sampling frequency of the ATWD is then given by the inverse of the peak frequency times 40 MHz. The typical frequency is found to be ~ 0.65 GHz.

4.2.4 Q_0 , Yield, T_0 , and $T-Q$ Corrections

Each PMT and electronics channel exhibits slight differences in gains, cable lengths, quantum efficiency, noise etc., which may change from run to run due to, for example, electronics changes (board swaps) or temperature variations. Before pulses can be compared or combined in calculations downstream, their charges must be normalized relative to each other to account for such differences, and their times must be adjusted to account for any constant offsets. These adjustments are called Q_0 and T_0 corrections, respectively. Additionally, the reconstructed arrival time of a pulse may vary as the charge increases, necessitating an additional “ $T-Q$ ” correction for larger pulses.

The Q_0 's are determined from central ^{60}Co calibration data. Events are selected that have between 200 and 1000 waveforms, and whose position reconstructs to within a radius of 1 m from the detector center. A subtraction is made for dark noise appearing very early or very late in the event. The mean of the charge distribution \bar{Q} is calculated for each ATWD, as is the relative efficiency or “yield”, the fraction of events in which the ATWD registered a hit. With the \bar{Q} 's and yields, it is possible to convert the charge of an event into a quantity proportional to the number of photoelectrons incident on the tube. This is done by dividing by Q_0 , the charge of a single p.e. at the tube, corrected for yield. Denoting the yield by y , we write

$$\bar{Q} = \frac{Q_0}{y} \sum_{n=1}^{\infty} nP(n|\mu), \quad (4.1)$$

Where $P(n|\mu)$ is the poisson probability of observing n p.e. when $\mu = -\log(1 - y)$ were expected. For ^{60}Co calibration source events, $\mu \approx 0.09$. Since contributions to

\bar{Q} from $P(3|\mu)$ and above are at the 0.01% level, the series may be truncated at 2 p.e. Solving for Q_0 , one obtains

$$Q_0 = \frac{y\bar{Q}}{P(1|\mu) + 2P(2|\mu)}. \quad (4.2)$$

In this analysis, the quantity number-of-photoelectrons (n_{pe}) is used to refer to the Q_0 -corrected event charge. The Q_0 's for one ^{60}Co calibration are used for all data taken until the next calibration. Typical time between calibrations is ~ 1 week for early data, and ~ 1 month currently. The Q_0 's have a mean value of around 20 with a $\sim 10\%$ spread.

The T_0 's are also calculated from central ^{60}Co calibration runs with the same event selection as for the Q_0 's. For each event, the average time of events within 30 ns of the median event time is calculated. The difference of each pulse time from this average time is histogrammed for each ATWD over all events in the calibration run. The histograms are then smoothed with the 353QH algorithm [63], and preliminary values of the T_0 's are taken from their peak positions. Corrections to the T_0 's are calculated by repeating the entire process, but with pulse times corrected with the preliminary T_0 's. Additional iterations are not found to change the T_0 's significantly. The spread of the T_0 's over all tubes is about 20 ns. Like the Q_0 's, the T_0 's are updated for each calibration period.

As discussed in Section 4.2, the T - Q correction for very large pulses is so big that they require a completely different estimation of the pulse time. But even events with energies of a few MeV will have a few moderately large pulses. To minimize biases from these pulses in the position estimation of such events, a small T - Q correction is applied. The bias in the estimation of the pulse times were assumed to be a property of the peak finding algorithm used. Hence the T - Q correction was estimated and applied as an average correction, identical for all channels. The correction was obtained from 337 nm laser calibration data; light in this frequency range is absorbed by the scintillator and re-emitted. The laser is set to pulse at a constant rate, making event selection trivial, and the intensity of the light is precisely varied using a series of neutral-density filters. For each laser event, pulses were Q_0 - and T_0 -corrected, and

for similar charges the difference between the pulse times and the trigger times were histogrammed. The peak pulse time was then plotted as a function of charge q , and the variation with charge was fit with a 3rd order polynomial in q . The time correction $\delta t(q)$ was determined to be

$$\delta t(q) = 0.5175q - 0.0195q^2 + 0.0003q^3. \quad (4.3)$$

Note that the 0th order coefficient would shift all times by the same amount and thus affects only the reconstruction of the global event time. Since the T - Q correction is assumed to be a property of the waveform analysis rather than a time-varying property requiring periodic calibration, the same coefficients are used for all run periods.

4.3 Vertex Fitting

Scintillation light is emitted from the event vertex, the interaction point, in all directions. The closest PMT's will receive the earliest pulses, and further PMT's will receive later pulses with the time spread determined by the speed of light in the scintillator. The timing pattern of the pulses provides the most accurate event position information, as well as the best estimate of the exact time at which the event occurred. However, pulse time errors, re-emission, attenuation, and any differences between the indices-of-refraction of the scintillator and buffer oil complicate the event's pulse-time distribution, preventing a straightforward analytical calculation of the event vertex and time.

The algorithm developed for this analysis finds the event vertex iteratively. A pre-fitter based on the pulse-charge distribution of the event provides an initial guess, and the pulse-time distribution is used to push the vertex, step by step, to an optimal position. The algorithm makes use of the known detector geometry, and has two tunable parameters, the "effective" light speeds in the scintillator and the buffer oil. These parameters are different from the physical speeds of light in these media in that they incorporate dispersion and retardation effects such as re-emission. While a more

physical parameterization of the light transport may be desirable, as shown below the fitter performs so well with just these two parameters that the incorporation of further detail into the light transport model was deemed unnecessary.

4.3.1 Charge-Based Pre-fit

The charge-based pre-fit is essentially a Q_0 -corrected charge-weighted average of the positions of PMT's registering hits in an event. This estimate of the event vertex is inherently biased: for an ideal, point-like vertex in a spherical detector with uniform photocathode coverage, even ignoring absorption, reflections, etc., a simple integration will show that the radius of the event is reconstructed roughly a fraction $2/3$ too small for events not too close to the PMT's. Fits to calibration data show that in practice the fraction is about 0.62, and the pre-fit radius is accordingly expanded to counteract the bias. Even the angular event coordinates will be slightly misreconstructed by the pre-fit, due to the holes in the PMT coverage at the chimney and at the bottom of the detector. However, such small effects are ignored, since the purpose of the pre-fit is merely to place the vertex in roughly the right spot so that the time-based fitting procedure can quickly and easily home in on the optimal vertex.

4.3.2 Time-Based Fit

Were the optimal vertex known in advance, the event time T would be best estimated by subtracting from the pulse times t_i the travel times τ_i from the optimal vertex to the tubes registering the pulses, and then taking the average:

$$T \approx \langle \tilde{t} \rangle = \frac{1}{N} \sum_{i=1}^N \tilde{t}_i, \quad (4.4)$$

where \tilde{t}_i denotes the corrected time $t_i - \tau_i$. The τ_i may be calculated knowing the detector geometry and the speeds of light in each material traversed. Given an initial guessed vertex, which may be far from the optimal vertex, a more optimal position may be found as follows. First, the event time is estimated as if the vertex were optimal using Equation 4.4. Then the moment of each \tilde{t}_i about $\langle \tilde{t} \rangle$ is used to define

a push $\delta\mathbf{r}_i$ in the direction $\hat{\mathbf{r}}_i$ between the guessed vertex and PMT i : $\tilde{t}_i > \langle \tilde{t} \rangle$ means that light arrived at the tube later than expected for the guessed vertex, and it should therefore be moved closer to the tube, and vice versa for $\tilde{t}_i < \langle \tilde{t} \rangle$. The equation is

$$\delta\mathbf{r}_i = c_{LS}(\tilde{t}_i - \langle \tilde{t} \rangle)\hat{\mathbf{r}}_i, \quad (4.5)$$

where c_{LS} is a tunable parameter that represents the effective speed of light in the liquid scintillator. The average push over all pulses, $\langle \delta\mathbf{r} \rangle$, moves the guessed vertex to a more optimal position, and the process is repeated until $\langle \delta\mathbf{r} \rangle$ becomes small.

The above assumes a point-like event and gaussian errors on the t_i . The various corrections on the t_i described in Section 4.2.4, namely the T_0 and $T-Q$ corrections, remove the major biases that would invalidate the latter assumption, so it is ignored. However, the former assumption of point-like-ness is not strictly correct, not only because the particle generating the scintillation light travels a finite distance, but more importantly because not all of the light in the event is emitted at the same point in time. See the T_0 -corrected pulse-time distribution for a central ^{60}Co calibration run shown in Figure 4.3. The distribution is marked by a sharp peak followed by a long tail with fast (~ 10 ns) and slow (~ 70 ns) decay components typical of organic scintillators. The finite emission time, scattering, reflections, re-emission, etc., complicate the fit by adding hits at times lagging relative to the interaction time. In order to minimize the biasing effects of these phenomena on the vertex fitting, the algorithm selects only pulses appearing in the peak of the corrected pulse-time distribution for the calculations described by Equations 4.4 and 4.5. The \tilde{t}_i are histogrammed into 1 ns bins, and the peak is fit with a gaussian. Pulses with \tilde{t}_i not falling within 10 ns of the peak are removed from the fitting process. It is this step that requires the initial charge-based pre-fit, so that the times near the peak are already those that best constrain the vertex. Also, by reducing the number of pulses used, this step makes the fit rather untrustworthy at energies much below 1 MeV, where the total number of detected photons is less than 100.




Figure 4.3: Distribution of T_0 -corrected pulse time differences from their average for a ^{60}Co calibration run. The data are fit to a gaussian-resolved double-exponential drawn in solid black. The fast and slow decay times fit to ~ 10 ns and ~ 70 ns, respectively.

4.3.3 Calibration and Performance

There are two tunable parameters in the fitter, namely the effective speeds of light in the liquid scintillator, c_{LS} , and the buffer oil, c_{MO} , which appear in the calculation of the τ_i and in Equation 4.5. The fitter was tuned using radioactive calibration source deployments along the z -axis. Sources were suspended from a stainless steel cable whose deployed depth was read out from an encoder. The encoder was calibrated to better than 1 cm accuracy by comparing the readout with marks made at 1 m intervals along the length of the cable. Events from the calibration sources, with energies ranging from 1 to 8 MeV, were selected using cuts on $nsmx_{ID}$ or energy to isolate the source peak. An additional loose cut requiring the reconstructed position to lie within 2 m of the expected source position helps eliminate large backgrounds from activity at the balloon surface for deployments near the very top or bottom of the LS. The reconstructed z -positions were histogrammed, and each distribution was fit to a gaussian. The values of c_{LS} and c_{MO} were varied until the z -biases,




Figure 4.4: Vertex reconstruction biases for various calibration sources. A constant offset of ~ 2 cm may be seen; the variation about this offset is less than 5 cm for positions within 5.5 m.

the differences between the deployed and reconstructed positions, were as small as possible for all sources. For z -positions within 5.5 m, the final z -biases, shown in Figure 4.4, vary by less than 5 cm about a ~ 2 cm offset. Variations with energy, which are particularly evident near the top of the detector, are at the level of 2 to 3 cm.

Figure 4.5 shows the position distribution of reconstructed events for a typical day-long run. The axes are reconstructed z -position and cylindrical-radius ρ , where the latter is divided by the balloon radius (r_{ball}) and squared in order for each bin to represent an identical amount of volume. The shape of the balloon and its deviation from sphericity in the upper half of the detector is illuminated by high activity on the balloon surface. The most active region of the detector is the chimney region, where shielding to external γ -rays is lowest. Activity from the three thermometers deployed just off the z -axis can also be seen in the figure.




Figure 4.5: Distribution of normal physics events plotted in z vs. cylindrical radius squared. Activity is concentrated at the chimney and at the bottom of the balloon. Radioactivity in the balloon surface and from a thermometer deployed near the center of the detector are also evident.

4.4 Energy Estimation

The estimation of an event's energy is a position- and particle-dependent conversion based on the hit pattern and the amount of light collected in the event. The conversion is divided into two steps: first the position dependence is removed by calculating the amount of light that would have been collected had the event occurred at the center of the detector. The result of this conversion is expressed in a unit called "analysis energy", or E_a , in which 0 corresponds to the minimum bias and 2.506 MeV corresponds to the amount of light collected in a typical ^{60}Co γ event when the source is deployed at the center of the detector. The second step involves a particle- and energy-dependent transformation from E_a to the "real" energy deposited by the particle assumed to be involved in the event.

4.4.1 Calculation of E_a

The calculation of E_a is based on a maximization of the likelihood of the pattern of hit PMT's around the detector. In order to minimize the dependence on waveform analysis, the likelihood function is written as a function of only which tubes did and did not register hits:

$$L(E_a, \mathbf{r}) = \prod_{i=1}^{1879} \begin{cases} 1 - e^{-\mu_i(E_a, \mathbf{r})} & \text{if hit} \\ e^{-\mu_i(E_a, \mathbf{r})} & \text{if not hit} \end{cases} \quad (4.6)$$

Here $\mu_i(E_a, \mathbf{r})$ is the energy- and position-dependent average number of photoelectrons for the i^{th} PMT, where the position of the event \mathbf{r} is taken from the vertex fitting procedure described in Section 4.3. The functional form of the $\mu_i(E_a, \mathbf{r})$ is dictated by the definition that E_a be proportional to the amount of light emitted by the event, and by the requirement that $E_a = 0$ corresponds to noise levels. We can immediately write

$$\mu_i(E_a, \mathbf{r}) = \alpha_i(\mathbf{r})E_a + \delta_i, \quad (4.7)$$

where $\alpha_i(\mathbf{r})$ is the fractional amount of light detected at the i^{th} PMT for an event at position \mathbf{r} , and δ_i is the dark rate in that tube. In order that the $\mu_i(E_a, \mathbf{r})$ have units of photoelectrons and E_a be in MeV, the units of $\alpha_i(\mathbf{r})$ are [p.e. MeV⁻¹]. The scale of $\alpha_i(\mathbf{r})$ is arbitrarily set so that $E_a = 2.506$ MeV corresponds to the average number of photoelectrons collected in ⁶⁰Co calibration source event at the center of the detector.

The $\alpha_i(\mathbf{r})$ depend on the light transport properties of the scintillator as well as the quantum efficiency and orientation of the PMT in question. The light transport model implemented incorporates absorption only; reflections and re-emission are ignored, as is asymmetrical shadowing from the ropes. Switching to a coordinate system with origin at PMT i and positive z -axis pointing toward the center of the detector, we write

$$\alpha_i(\mathbf{r}) = \eta_i \frac{\Omega(r, \theta)e^{-r/\Lambda}}{\Omega(r_{PMT}, 0)e^{-r_{PMT}/\Lambda}}. \quad (4.8)$$

Here η_i accounts for both the quantum efficiency of the tube and the arbitrary rescaling constant described above; Λ is the absorption length, r and θ are the spherical coordinates of the event in the chosen coordinate system, $r_{PMT} = 8.5$ m is the distance between the PMT and the center of the detector, and $\Omega(r, \theta)$ is the solid angle subtended by the PMT relative to the event vertex. Monte Carlo studies using the detailed shape of the PMT's show that the solid angle is well approximated by the relation $\Omega(r, \theta) \approx \frac{A}{r^2}(0.1 + 0.9 \cos \theta)$, where A is the cross-sectional area of the photocathode when viewed head-on. The term in the denominator of Equation 4.8 normalizes the the position-dependent contribution to $\alpha_i(\mathbf{r})$ relative to its value for an event at the center of the detector, weakening the correlation between η_i and Λ .

It should be noted that, when all tubes are hit, the likelihood function of Equation 4.6 is maximized when $E_a \rightarrow \infty$. However, this is not a problem for energies relevant for reactor $\bar{\nu}_e$ studies, for which the μ_i are all $\lesssim 1$, ensuring that a significant number of tubes are not hit so that the energy estimation is accurate.

4.4.2 Calibration and Performance

All together, there are 2 parameters to be determined for each tube, η_i and δ_i , and additionally one global parameter Λ . The calculation of E_a is found to depend weakly on the value of Λ , with values between 15 m and 40 m giving roughly the same results. So for computational simplicity the value of 25 m is used. The δ_i are determined on a run-by-run basis by averaging the hit rate for each tube during the period in the event time window prior to the arrival of light from the event itself (represented by the parameter "baseline" in the fit in Figure 4.3). That hit rate is then rescaled by the amount of time a channel is available to take data during an event. The η_i are determined from fits to ^{60}Co calibration data taken along the z -axis of the detector. For each deployment at position $\mathbf{r}_j = (0, 0, z_j)$, the number of source events N are counted, as well as the number of events $n_i(\mathbf{r}_j)$ with hits in PMT i . The values of the η_i are obtained by maximizing the likelihood function

$$L = \prod_{i,j} (1 - e^{-\mu_i(2.506\text{MeV}, \mathbf{r}_j)})^{n_i(z_j)} (e^{-\mu_i(2.506\text{MeV}, \mathbf{r}_j)})^{N-n_i(z_j)}. \quad (4.9)$$

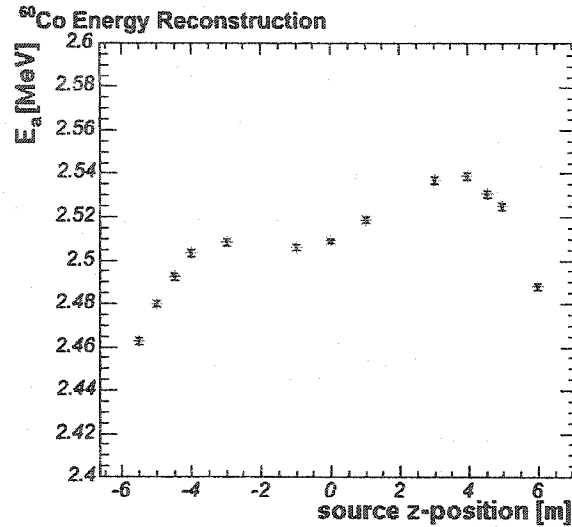


Figure 4.6: Variation of ^{60}Co energy reconstruction versus z -position of the deployed source.

For a typical calibration, the average of the η_i is 0.220 ± 0.004 p.e. MeV^{-1} for the 17-inch tubes, and 0.293 ± 0.005 p.e. MeV^{-1} for the 20-inch tubes. The average δ_i for 17-inch tubes is 0.023 ± 0.002 , while for 20-inch tubes it is 0.036.

Figure 4.6 shows the resulting variation in reconstructed energies along the z -axis for ^{60}Co calibration sources. The variation is $\sim 1\%$ for positions within 5.5 m. Shadowing by the Kevlar ropes constraining the balloon may contribute to the variation and is the subject of current study.

The variation of the energy estimation off-axis is tested by reconstructing the energies of spallation neutron capture γ 's following muons. The muons are tagged according to their large light emission, as discussed in Section 4.5.1. Events with $nsmx_{ID} > 250$ occurring between $150 \mu\text{s}$ and 2 ms following a muon are tagged as spallation neutron candidates; events in the following 100 ms are tagged for background subtraction. Events within the first $150 \mu\text{s}$ following a muon are ignored due to noise associated with ringing in the electronics that drown out physics events. Figure 4.7 plots the background-subtracted distribution of the energies of these events vs. spherical radius cubed (r^3), normalized to the radius of the balloon. The 2.2 MeV




Figure 4.7: Energy of spallation neutron events versus r^3 . Gaussian fits to slices of constant r^3 give a variation of 0.58% for $r < 5.5$ m.

np -capture γ peak can be clearly seen. Gaussian peaks are fit to the γ peak in slices of constant radius; the variation of these peaks for positions within 5.5 m was found to be 0.58%.

The stability of the energy reconstruction over the data taking period is tested with periodic calibration data. Figure 4.8 shows the results for ^{60}Co calibration data. The variation is at the fraction-of-a-percent level, not surprising since the energy estimation is calibrated to the ^{60}Co γ -line.

Figure 4.9 shows the energy distribution for normal, prompt-trigger physics events collected in a typical day-long run. A $150 \mu\text{s}$ veto has been applied after muon events to remove noise. The $nsmx_{ID} > 200$ threshold appears as a wall at ~ 1 MeV. The contributions from ^{210}Pb and ^{208}Tl visible in the $nsmx_{ID}$ distribution (Figure 3.4) appear sharper here. The muon peak is at energies $\gg 10$ MeV and is not shown. The flat spectrum above 5 MeV is due to muon spallation products with lifetimes longer than $150 \mu\text{s}$.

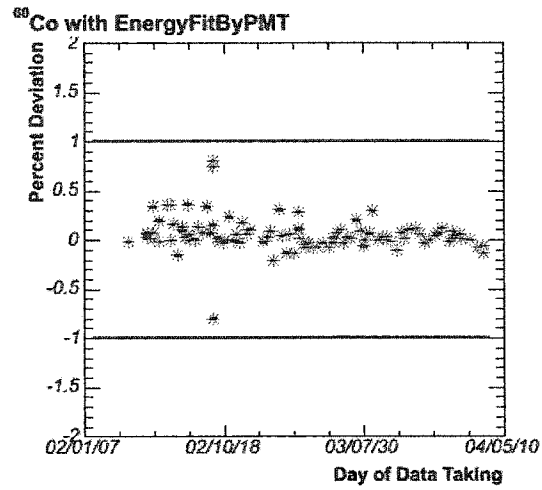


Figure 4.8: Time variation of ^{60}Co energy reconstruction. The variation is at the fraction-of-a-percent level.




Figure 4.9: Energy distribution of physics events. The lack of events below 1 MeV is due to the triggering threshold. The tail to higher energies are primarily spallation products.




Figure 4.10: A fit of the resolution function (Equation 4.10) to calibration points. The parameter a has been fixed to 0.1 MeV to match dark rate levels indicated by the calibration constants.

4.4.3 Energy Resolution

The energy resolution function is approximated as a gaussian with width σ dependent on the square-root of the number of collected photons. It may be written

$$\sigma(E_a) \approx k \sqrt{E_a + a}, \quad (4.10)$$

where k and a are constants to be determined from the calibration data. The constant a accounts for the contribution to σ from the dark rate that is subtracted off by the δ_i . A fit of $\sigma(E_a)$ to the values of σ extracted from gaussian fits to the calibration source event distributions is shown in Figure 4.10. The best fit (with a constrained to be non-negative) yields a value of 0 for a ; however, small positive values of a still fit the data reasonably well. Since the δ_i are found to be roughly 10% of the η_i , a is fixed by hand to equal 0.1 MeV. The fit value for k is $6.4 \pm 0.2\%$.

4.4.4 Conversion from E_a to E

The “real” energy deposited by a particle passing through KamLAND’s scintillator is not quite the same as the analysis energy returned by the energy estimator. While E_a is a quantity that is proportional to the light yield Y of an event, the amount of light emitted is not strictly proportional to the incident particle’s energy due to two effects, quenching and Cherenkov emission.

Quenching is the saturation of scintillation emission for highly ionizing particles. This saturation is well described by the semi-empirical Birks’ Law [64], which relates the infinitesimal light yield dY during a step dx to the energy lost during the step $\frac{dE}{dx}$.

$$\frac{dY}{dx} \propto \frac{\frac{dE}{dx}}{1 + k_B \frac{dE}{dx}}. \quad (4.11)$$

The constant k_B is called the Birks constant and depends on the chemical composition of the scintillator. k_B is positive, so the denominator is greater than unity. The light yield is lower when $\frac{dE}{dx}$ is large, so that for example a highly ionizing α particle will emit much less scintillation light than a proton or electron of the same energy.

A charged particle emits Cherenkov radiation if its velocity exceeds the speed of light in the medium being traversed. Cherenkov emission is the electromagnetic analogy of the formation of a wake behind a boat traveling faster than the speed of waves on a body of water, or the formation of a sonic boom from a jet traveling faster than the speed of sound. The opening angle θ_c of the “light-wake”, or the Cherenkov cone, of a particle traveling with velocity v in a medium with refractive index n is given by

$$\theta_c = \arccos\left(\frac{c}{nv}\right), \quad (4.12)$$

where c is the speed of light. The spectrum of light emitted per unit length is

$$\frac{d^2 N(E)}{dE dx} = \frac{\alpha z^2}{\hbar c} \sin^2 \theta_c, \quad (4.13)$$

where α is the fine-structure constant, z is the particle’s charge in units of the electron charge, and \hbar is Plank’s constant divided by 2π . The energy dependence of the

spectrum is inherited through θ_c from the frequency dependence of the refractive index.

The relationship between analysis and real energy may be written as

$$\frac{E_a}{E} \propto 1 - \delta_q(E, k_B) + k_c \delta_c(E), \quad (4.14)$$

where δ_q and δ_c account for the fractional losses and gains due to quenching and Cherenkov emission, respectively. For a given particle and for a particular value of k_B , $\delta_q(E, k_B)$ may be calculated by integrating Equation 4.11; $\delta_c(E)$ is likewise determined by Equation 4.13. The constant k_c represents the fractional collection of the Cherenkov radiation, accounting for the absorption and re-emission of the radiation as well as photocathode coverage and the quantum efficiency of the PMT's. Both k_B and k_c must be determined from the data.

The functions $\delta_q(E, k_B)$ and $\delta_c(E)$ were calculated for energies up to 30 MeV using the EGSnrc [65] Monte Carlo package for γ 's and e^\pm 's in KamLAND's scintillator. In the simulation, the quenching and Cherenkov emission are summed step-by-step along the path of a simulated particle until it reaches a tracking threshold. This tracking threshold necessitates the addition of another particle- and energy-dependent factor $k_0 \delta_0(E)$ to the right-hand-side of Equation 4.14 to allow part of the energy lost below the tracking threshold to be recovered. The constant k_0 performs the same function for $\delta_0(E)$ that k_c performs for $\delta_c(E)$, and likewise must be determined from the data. The functions $\delta_0(E)$ are a byproduct of the simulations.

The values of k_B , k_c , k_0 , and the proportionality constant in Equation 4.14 are obtained from fits to calibration data using MINUIT [66]. In order to weaken the correlation between the fit parameters, the actual fit function was normalized by its value at the energy of an np -capture γ , $E_{np} = 2.2$ MeV, giving

$$\frac{E_a(E, \Theta_E)}{E} = a \frac{1 - \delta_q(E, k_B) + k_c \delta_c(E) + k_0 \delta_0(E)}{1 - \delta_q(E_{np}, k_B) + k_c \delta_c(E_{np}) + k_0 \delta_0(E_{np})}. \quad (4.15)$$

The constant a accounts for both the normalization and the arbitrary scaling of E_a ; the symbol Θ_E is a shorthand representation of the four parameters to be fit. The




Figure 4.11: Fit of Equation 4.15 to γ calibration source energies. The grey shaded region depicts the 1σ error bars.

results of the fit are shown in Figure 4.11, where k_B is constrained primarily by the α lines detected in ^{212}Bi - ^{212}Po and ^{214}Bi - ^{214}Po coincidences, and the remaining parameters are determined from calibration source γ lines. The value of k_B was found to be $(1.046 \pm 0.039) \times 10^{-2}$ g/cm²/MeV, but the fit to the γ data depends very weakly on this parameter. The remaining three parameters fit to $a = 1.122 \pm 0.0029$, $k_c = 0.29 \pm 0.11$, and $k_0 = 0.34 \pm 0.13$. The grey shaded region depicts the 1σ error bars on the energy scale transformation.

Equation 4.15 is transcendental in E and must be numerically solved to obtain the desired transformations for each particle type, $E_{\gamma,e^\pm}(E_a, \Theta_E)$. It should be noted that the transformation function for positrons necessarily includes the contribution of the two 511 keV annihilation γ 's; it will be labeled $E_p(E_a, \Theta_E)$ to match the notation used in Equation 2.2 for the energy of the prompt event in a $\bar{\nu}_e$ coincidence pair. The uncertainty in $E_p(E_a, \Theta_E)$, propagated from the uncertainty in fit for $E_\gamma(E_a, \Theta_E)$, at $E_p = 2.6$ MeV is 2.4%.

4.5 Muon Track Fitting

Not all events in KamLAND make point-like vertices whose positions and energies can be reconstructed with the techniques described in Sections 4.3 and 4.4. A second class of events are highly energetic particles that make scintillation tracks. This class consists mostly of cosmic rays, or muons. These muons may activate the material they traverse, producing radioactivity along their tracks that can mimic the coincidence neutrino signal. Devoted algorithms were developed to reconstruct these muon tracks so that these backgrounds could be efficiently removed.

4.5.1 Muon Tagging

Muons are identified primarily by their Cherenkov production in the OD or buffer oil, or by the scintillation light they emit in the ID. Events registering at least 10 hits in the OD are called “OD muons”, whereas events with significant ID activity are labeled “ID muons”. ID muons are separated into two classes, LS and oil muons. LS muons, which traverse the scintillator, are easily identified by their enormous light yield. Oil muons are muons that pass through the buffer oil between the balloon surface and the PMT’s, generating primarily Cherenkov radiation at a fraction of the intensity of the LS muons. ID muons are identified and distinguished by their value of n_{pe} (number of photoelectrons) in the ID. Figure 4.12 plots $nsmx_{OD}$ (the sum of the individual values of $nsmx$ for each OD section) vs. n_{pe} collected in the ID for a typical day-long run. Clusters corresponding to OD, oil, and LS muons are labeled. The cuts used to categorize these muon events are shown as solid lines; note that oil and LS muons generating significant OD activity are classified as both OD and ID muons. The efficiency for detecting an ID muon by the OD is 0.9931 ± 0.0004 . The inefficiency for tagging muons passing through the LS by their ID activity is assumed to be negligible, owing to their large energy deposit.

Another subclassification involves LS muons creating more than 10^6 p.e., identified by the sum of the charges over all pulses in the event, denoted Q_{tot} (differences in the Q_0 are neglected for these high charge events). These events are tagged as “showering muons”, since they are thought to be associated with muon-induced hadronic showers




Figure 4.12: A plot of $nsmx_{OD}$ vs. n_{pe} collected in the ID for a typical day-long run, with clusters corresponding to the three main muon classes labeled. The solid lines show the cuts used to select each class. Note that oil and LS muons with $nsmx_{OD}$ exceeding the OD muon threshold are also classified as OD muons.

in the LS, and therefore enhanced activation of the scintillator. These make up only 10.6% of LS muons but account for much of the muon induced backgrounds. A fifth muon class describes anomalous events resulting from a bug in the early triggering logic in which concurrent high activity in the ID and OD failed to issue a trigger signal to collect ID waveforms. These events are easily identified as having high $nsmx_{ID}$ and no waveforms, and are treated as if they are showering muons by the analysis. The selection criteria for the various muon classes referred to in this analysis are listed in Table 4.1.

4.5.2 Track Reconstruction

Track reconstruction is attempted for all events with $\log_{10}(Q_{tot}) > 4.5$. Like vertex reconstruction, track reconstruction begins with a pre-fit to obtain a reasonable initial guess at the track direction, followed by a more precise determination based on the arrival times registered by each PMT. The second step is not performed for events

Table 4.1: Muon Classification and Rates

| Class | Selection | Rate |
|-----------|---|------------------------|
| OD | $nsmx_{OD} > 10$ | ~ 0.65 Hz |
| LS | $\log_{10}(n_{pe}) > 4.8$ | 0.2044 ± 0.0001 Hz |
| Oil | $3.5 < \log_{10}(n_{pe}) \leq 4.8$ AND $[nsmx_{OD} - 10 > 20(4 - \log_{10}(n_{pe})) \text{ OR } \log_{10}(n_{pe}) > 4]$ | 0.1259 ± 0.0001 Hz |
| Showering | $\log_{10}(Q_{tot}) > 7.2$ | 0.0224 ± 0.0001 Hz |
| Missed | $nsmx_{ID} > 1250$ AND no waveform data | $\ll 1$ Hz |

which do not give a reasonable pre-fit.

The pre-fit is based on the empirical observation that muons create regions of high light collection, called “wounds”, at their entrance and exit points. For an example, see Figure 4.13, in which the PMT positions on the inside of the containment vessel have been mapped to 2D using a Mollweide projection and have been colored according to their charge. The entrance and exit wounds are identified by ignoring all but the largest 10% of the pulses, sorting the remaining pulses according to their times, and taking the average positions of the first and last 25 (or the first and last half, if there are less than 25 pulses). Then charges more than 3 m beyond the wound position are removed from consideration and the wound positions are recalculated. For muons with $Q_{tot} < 2.5 \times 10^6$ (oil muons, roughly), the muon’s track direction is set to point from the entrance wound to the exit wound. For higher charge muons (LS muons), the entrance wound actually appears near the PMT’s pointed to by the balloon-surface normal vector originating at the muon’s LS entrance point, since Cherenkov light generated in the BO propagates primarily in the direction of the muon and hence away from the tubes. The same is true for the exit point of an LS muon since the wound analysis is restricted to tubes with the highest charges. Hence pre-fits for LS muon tracks point between the projections of the entrance and exit wounds onto the balloon surface.

The refinement of the track direction is based on a light transport model for scintillating relativistic particles. A charged particle ionizing material in KamLAND’s LS emits scintillation light isotropically. If the particle is non-relativistic or doesn’t

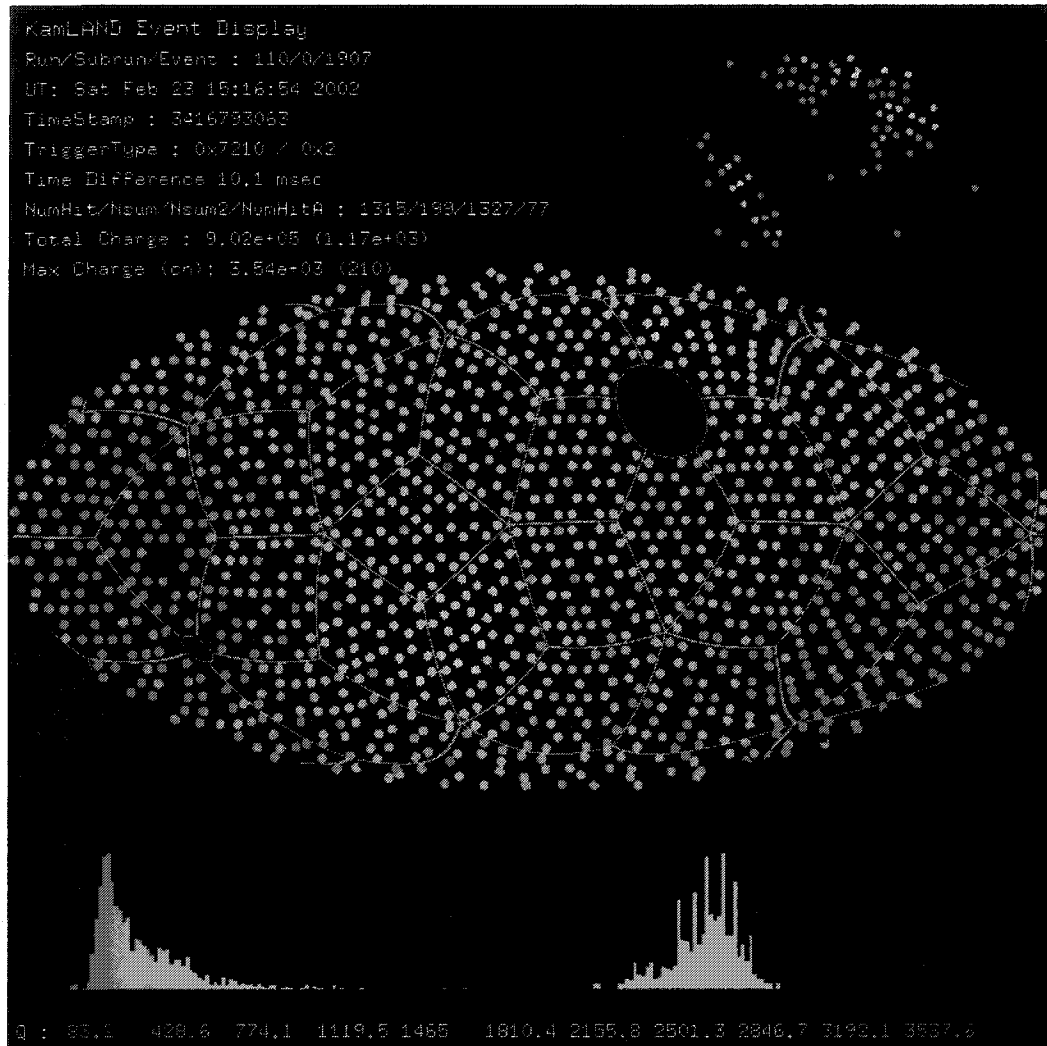


Figure 4.13: PMT hits for a typical muon event showing entrance and exit wounds. The PMT positions on the inside of the containment vessel have been mapped to 2D using a Mollweide projection and have been colored according to their charge. The lower left and right panels show the charge- and time-distributions of the hits, respectively. The smaller display in the upper-right corner shows the associated OD activity using the same projection.

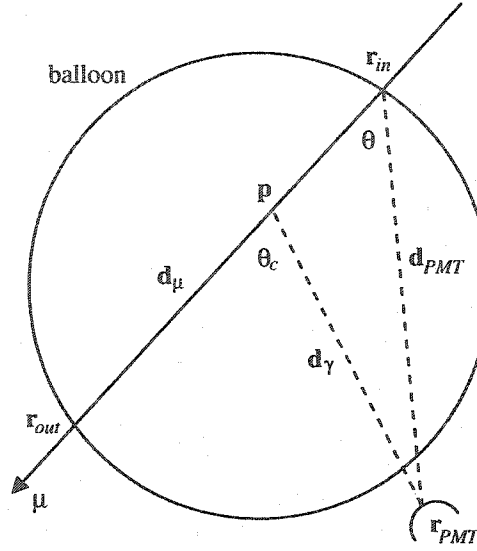


Figure 4.14: Definitions of geometrical parameters for muon tracks.

travel very far, the wavefront of the scintillation light is a sphere centered on the particle's position, as is assumed for vertex fitting. However, if the particle is a highly relativistic muon crossing the detector, the wavefront will form a wake behind the particle, analogous to the Cherenkov cone, with the same opening angle θ_c given by Equation 4.12. The first scintillation light arriving at a particular PMT from the muon will have been generated at the first point \mathbf{p} in the LS along the track whose difference vector from the PMT position lies at angle θ_c relative to the muon track direction. Referring to Figure 4.14 in which, for simplicity, the balloon is approximated as a perfect sphere, let \mathbf{r}_{in} and \mathbf{r}_{out} be the muon's entrance and exit points in the LS, $\mathbf{d}_\mu = \mathbf{r}_{out} - \mathbf{r}_{in}$ be the track vector, \mathbf{r}_{PMT} be the position of the PMT in question, and θ be the angle between \mathbf{d}_μ and $\mathbf{d}_{PMT} = \mathbf{r}_{PMT} - \mathbf{r}_{in}$. Writing $\mathbf{p} = \mathbf{r}_{in} + \alpha \mathbf{d}_\mu$, α may be calculated using the Law of Sines,

$$\alpha(\mathbf{r}_{in}, \mathbf{r}_{out}) = \max \left(0, \min \left(1, \frac{d_{PMT} \sin(\theta_c - \theta)}{d_\mu \sin \theta_c} \right) \right), \quad (4.16)$$

where d_{PMT} and d_μ are the magnitudes of \mathbf{d}_{PMT} and \mathbf{d}_μ , respectively. α is written as a function of \mathbf{r}_{in} and \mathbf{r}_{out} since once these two vectors are set, there is a unique value

of α for each PMT. This suggests a procedure to find the most likely values of \mathbf{r}_{in} and \mathbf{r}_{out} and the entrance time t_{in} for LS muons: given the pre-fit guesses for \mathbf{r}_{in} and \mathbf{r}_{out} , α is calculated, as is $\mathbf{d}_\gamma = \mathbf{r}_{PMT} - \mathbf{p}$, the distance the scintillation light travels, for each PMT. Next, with an estimate from the pre-fit for t_{in} , the expected arrival time of the earliest scintillation light at the i^{th} PMT, \tilde{t}_i , is calculated according to

$$\tilde{t}_i = t_{in} + \alpha_i \frac{d_\mu}{c} + n \frac{d_{\gamma,i}}{c}, \quad (4.17)$$

where n is the measured refractive index of the LS. Differences between the refractive indices of the LS and the buffer oil are ignored. The \tilde{t}_i are compared to the T_0 -subtracted arrival times t_i of the earliest large pulse at each PMT, and from the differences $\tilde{t}_i - t_i$ the log-likelihood is constructed as

$$\log L = \sum_{i=1}^N -\frac{(\tilde{t}_i - t_i)^2}{2\sigma^2}, \quad (4.18)$$

where N is the number of hit 17-inch PMT's, and $\sigma = 2$ ns is the error on the t_i , assumed to be gaussian and identical for all 17-inch PMT's. The log-likelihood is then maximized over the variation of \mathbf{r}_{in} , \mathbf{r}_{out} , and t_{in} .

4.5.3 Performance

Unfortunately there is no unambiguous way to calibrate the muon track reconstruction. However, there are several ways to evaluate its performance. First, the zenith-angle distribution of the reconstructed tracks may be compared to simulations of the muon flux at KamLAND. Such a simulation was performed using the MUSIC software toolkit [67] based on the sea-level muon flux and topographic data of Mt. Ikenoyama. Figure 4.15 shows the comparison between data and Monte Carlo for azimuthal (ϕ) and zenith (θ) angles. The agreement is quite good, with both data and MC showing essentially the same features in each variable. The excess flux in the Western and Southern directions in azimuthal angle, for example, corresponds to directions along which the mountain is not as thick. Differences between data and MC are attributable




Figure 4.15: Angular distribution of reconstructed muon tracks. The crosses are data, the histogram is Monte Carlo. The peaks and valleys in the azimuthal angle (ϕ) are due to variations in the rock depth in the different directions. Differences between the two distributions are attributable to the resolution of the reconstruction algorithm.

to resolution effects.

Another sanity check tests whether the reconstructed muon flux is uniform across the detector. This is done by checking for flatness in the distribution of impact parameter squared (b^2) for the reconstructed tracks relative to the center of the detector. This distribution is plotted in Figure 4.16, in which b has been normalized to the balloon radius r_{ball} . There appears to be a problem near the edge of the balloon, as if some oil muons are being systematically pushed into the LS. This is probably related to the algorithm's different treatment of events it assumes to be oil or LS muons, an identification that is difficult to make for tracks that just graze the LS. The fact that the track directions still agree well with MC implies that the misreconstruction is likely an issue of shifted impact parameters.




Figure 4.16: Distribution of reconstructed muon impact parameter squared. There appears to be a problem near the balloon radius, which is likely due to simple shifts in b .

4.6 Total Reconstruction Efficiency

It is possible that one of the steps in the event reconstruction will fail for some anomalous events. For example, the location of an event with a strange pulse-time distribution might be pushed far outside the detector volume, or a few spurious large pulses might give an otherwise normal low-energy event an enormous energy. A contribution must be included in the $\bar{\nu}_e$ detection efficiency to account for the probability that an event is reconstructed with approximately the right position and energy.

The total reconstruction efficiency is obtained for physics events by counting the fraction of calibration source γ events that survive reconstruction. The raw event count is made by counting events during a source run with $nsmx_{ID}$ within a certain range, and subtracting off the background rate estimated in a normal physics run using the same cuts on $nsmx_{ID}$. The reconstructed events are counted the same way, except with cuts on energy and position to be near those of the source. The reconstructed vertex is additionally required to have converged normally.

To cover the range of reactor $\bar{\nu}_e$ energies, events from ^{68}Ge , ^{60}Co , and $^{241}\text{Am}/^9\text{Be}$

Table 4.2: Cuts for Evaluating the Total Reconstruction Efficiency

| Source | Trigger Cuts | Energy Cut | Position Cut |
|-------------------------------|--------------------------------|---|--|
| ^{68}Ge | $160 \leq nsmx_{ID} \leq 300$ | $0.5 \text{ MeV} < E_a < 3 \text{ MeV}$ | $ \mathbf{r} - \mathbf{r}_{source} < 2 \text{ m}$ |
| ^{60}Co | $300 \leq nsmx_{ID} \leq 600$ | $0.5 \text{ MeV} < E_a < 6 \text{ MeV}$ | $ \mathbf{r} - \mathbf{r}_{source} < 2 \text{ m}$ |
| $^{241}\text{Am}/^9\text{Be}$ | $600 \leq nsmx_{ID} \leq 1000$ | $2 \text{ MeV} < E_a < 13 \text{ MeV}$ | $ \mathbf{r} - \mathbf{r}_{source} < 3 \text{ m}$ |

calibrations are used. The cuts for each are listed in Table 4.2. The efficiencies for ^{60}Co and $^{241}\text{Am}/^9\text{Be}$ are 1.0006 ± 0.0002 and 0.998 ± 0.001 , respectively. These efficiencies showed no significant dependence on the z -position of the source in question. The efficiency for ^{68}Ge is 1.040 ± 0.005 , although this value was found to be highly sensitive to the background subtraction. The total reconstruction efficiency is therefore taken to be the average of the ^{60}Co and $^{241}\text{Am}/^9\text{Be}$ efficiencies, 99.9%, with a 0.1% systematic error allotted to account for the difference between these two measurements.

Chapter 5

Data Reduction

Data reduction refers to the process of sifting through the data in search of electron antineutrinos. The first step in data reduction is to eliminate bad runs. Then anti-neutrinos are identified as event-pairs that are correlated in time and space, and have a delayed event energy consistent with that of a 2.2 MeV np -capture γ . A 5.5 m radius spherical fiducial volume cut is applied to remove backgrounds near the surface of the balloon. A series of spallation cuts are applied to remove cosmogenic backgrounds. These cuts and their efficiencies are described in detail below.

5.1 Data Cleaning and Livetime Calculation

An entire run may be rejected for any of the following reasons:

- It is a testing/engineering or calibration run.
- There is an electronics failure or high-voltage problems, or the magnetic field-canceling coils were off.
- The number of inactive or over-active channels is too high.
- The trigger rate or $nsmx$ distribution is abnormal.
- Any of the muon rates are abnormal.

- The accidental background rate (see Section 6.2) is abnormal.
- The OD hit rate or inefficiency is abnormal.
- The run is shorter than one hour, and the reason it was stopped is unknown.
- The operator taking the run requests in the log book that the run not be analyzed.

When any of these problems can be limited to a short period during a run, only that period is vetoed for analysis.

A period of a run may also be vetoed if it is found to be noisy. The two major types of noise which must be addressed are post-muon noise, and periods of high data rates not associated with muons. The former occurs in early runs during the 150 μs following events with very large n_{pe} . An electronics upgrade in January 2003 shortened the noisy period to 40 μs for subsequent runs. To remove these events, a 150 μs veto is applied after all LS and missed muons. Even after this veto is applied, some noisy periods may still be found in the data. These periods are tagged as having more than 5 analyzable events in a period of 6 ms (equivalent to a trigger rate of almost 1 kHz), and are removed via a cut on “multiplicity”, described in Section 5.2.

The livetime of un-vetoed periods is calculated simply by summing the time between the first and last trigger record in each clean period. The livetime calculation was checked by counting 1PPS triggers issued by the GPS module throughout the run. The two methods agreed to within one second per day. Livetime is subtracted for “gaps” in the data, periods during which no trigger records are issued for more than 0.1 s, a symptom that the trigger module may have disabled itself due to data overflow or network problems. During standard operation, trigger records are recorded at a rate of well over 100 Hz, so this subtraction wrongfully removes $< 0.05\%$ of the livetime. Accounting for this uncertainty, the total livetime for this analysis is 484.75 ± 0.24 days.

5.2 Event Pairing

In order to facilitate the application of cuts on pairs of events, a list is constructed of all event pairs occurring within 1.5 ms of each other. The list is organized into “multiplets”, groups of sequential events in which the time difference between consecutive events never exceeds 1.5 ms. In order for an event to appear in a multiplet, it must have $r < 6.5$ m and $E_a > 0.8$ MeV. The noise cut mentioned in Section 5.1 is applied by requiring that no multiplet have a multiplicity, or number of events, larger than 5. The correction to the livetime from this cut is of order 10^{-7} and is ignored.

5.3 Time Correlation

The neutron emitted in an inverse beta interaction quickly thermalizes and begins a random walk through the scintillator. The probability of the neutron capturing on a proton and producing the 2.2 MeV delayed γ event in any given instant is constant. Therefore the probability of the capture occurring within dt of time t after the emission of the neutron falls exponentially with time-constant τ :

$$P(t)dt = \frac{1}{\tau}e^{-t/\tau}dt. \quad (5.1)$$

The efficiency of a cut on t is given by the integral of Equation 5.1 and therefore depends on the value of τ .

The value of τ itself depends on the density-weighted sum of the cross-section for the neutron to be captured on the various atoms that comprise the scintillator. First-principle calculations of τ give values in the range $211 \pm 5 \mu\text{s}$. Monte Carlo simulations of the neutron walk using the Geant4 software toolkit [68] give a mean τ of $212 \pm 2 \mu\text{s}$.

Using data, τ was measured from both spallation neutrons following muons and $^{241}\text{Am}/^9\text{Be}$ calibration data. Spallation neutrons were selected using the same criteria described in Section 4.4.2. Only runs prior to the electronics upgrade in January 2003 were used. The distribution of time-since-last muon for each event is plotted in Figure 5.1. A fit to an exponential plus a constant background gave a capture




Figure 5.1: Neutron capture time distribution measured from spallation neutrons. The first 150 μs following a muon is vetoed due to electronics noise. The fit function, shown as the solid line, is $N_0 e^{-t/\tau} + bg$. The fit extends out to 100 ms (only the first 2 ms are shown).

time of $\tau = 207.8 \pm 0.5 \mu\text{s}$. $^{241}\text{Am}/^9\text{Be}$ calibration events were paired using the same procedure to pair $\bar{\nu}_e$ candidates, except that the coincidence time was allowed to extend to 5 ms, and no cut was applied on multiplicity. The prompt event was required to have $E_a > 4 \text{ MeV}$, which selects neutron emissions accompanied by 4.44 MeV γ 's from $^{12}\text{C}^*$, and the delayed event was required to have an energy consistent with a 2.2 MeV np -capture γ . Both prompt and delayed events were required to be reconstructed within 3 m of the deployed source position. A fit of an exponential plus constant background to the distribution of the times between the prompt and delayed events gives a capture time $\tau = 203 \pm 3 \mu\text{s}$. The slightly faster capture time could be due to the higher density of the neutron moderating material Delrin encapsulating the source. The combined value of τ considering all measurements and calculations is taken to be $208 \pm 3 \mu\text{s}$, where the error $\delta\tau$ accounts for the spread in the various values obtained for τ .

Antineutrino candidates are required to have a time correlation $0.5 \mu\text{s} < \Delta t < 660 \mu\text{s}$. This variable is plotted in Figure 5.2 for the $\bar{\nu}_e$ candidates after applying




Figure 5.2: Δt distribution of $\bar{\nu}_e$ candidates. The distribution is fit to the function $N_0 e^{-\Delta t/\tau} + bg$. The fit value $\tau = 225 \pm 25 \mu\text{s}$ is consistent with the measured value of $\tau = 208 \pm 3 \mu\text{s}$.

all cuts except the time correlation cut. The efficiency of this cut is determined by simply integrating the exponential capture time distribution,

$$\epsilon_{\Delta t} = \int_{t_1}^{t_2} \frac{1}{\tau} e^{-t/\tau} dt. \quad (5.2)$$

Since $\frac{\delta\tau}{\tau} \ll 1$ and $t_1 \ll \tau < t_2$, the uncertainty of $\epsilon_{\Delta t}$ due to $\delta\tau$ may be approximated as

$$\delta\epsilon_{\Delta t} \approx \left| \frac{\partial\epsilon}{\partial\tau} \right| \delta\tau = \left| \frac{t_1}{\tau} e^{-t_1/\tau} - \frac{t_2}{\tau} e^{-t_2/\tau} \right| \frac{\delta\tau}{\tau}. \quad (5.3)$$

With $t_1 = 0.5 \mu\text{s}$, $t_2 = 660 \mu\text{s}$, and $\tau = 208 \pm 3 \mu\text{s}$, the efficiency is then $\epsilon_{\Delta t} = 0.9557 \pm 0.0019$.

5.4 Space Correlation

Perhaps the most powerful cut for reducing backgrounds after the time correlation cut is the requirement that the prompt and delayed events occur within 1.6 m of




Figure 5.3: Δr distribution of $\bar{\nu}_e$ candidates. Candidates are required to have $\Delta r < 1.6$ m. The increase in the distribution beyond the cut is due to accidental backgrounds.

each other. The distance Δr between the prompt and delayed events is plotted for the $\bar{\nu}_e$ candidates in Figure 5.3, with the cut depicted by the vertical solid line. The increasing distribution outside of the cut is due to accidental backgrounds.

To obtain the efficiency of the cut, it is necessary to investigate the effects that determine the spacial correlation of the prompt and delayed events. One component is the neutron's random walk at thermal energies. Geant4 Monte Carlo simulations indicate that neutrons are captured on average within 9 cm of the interaction vertex. The absorption length of the 2.2 MeV γ and the vertex reconstruction resolution further smear out the distribution of Δr , the distance between the prompt and delayed events.

The radial distribution of γ calibration source events were found to be nicely fit by an exponential absorption convolved with a gaussian resolution. The apparent absorption lengths and resolutions for the various sources are summarized in Table 5.1. As expected, the absorption length is longer for higher energy γ 's, and the resolution improves with the greater light yield. The absorption lengths for ^{68}Ge and ^{60}Co are

Table 5.1: Absorption Lengths and Vertex Resolutions of Calibration Sources

| Source | Energy [MeV] | Absorption Length [mm] | Resolution [mm] |
|-------------------------------|--------------|------------------------|-----------------|
| ^{68}Ge | 1.022 | 94.5 | 173.1 |
| ^{65}Zn | 1.116 | 157.9 | 158.7 |
| $^{241}\text{Am}/^9\text{Be}$ | 2.223 | 206.3 | 146.9 |
| ^{60}Co | 2.506 | 89.3 | 136.6 |
| $^{241}\text{Am}/^9\text{Be}$ | 4.5 | 251.9 | 119.0 |

much smaller than the others since these sources emit two γ 's in different directions, while the other source events consist of only a single γ .

A ^{68}Ge source event may be taken to be representative of the prompt event from a $\bar{\nu}_e$ interaction with energy just at or the threshold for inverse β -decay. For higher energy $\bar{\nu}_e$ events, the apparent absorption length is shorter because the positron deposits its energy much closer to the prompt vertex than the annihilation γ 's. Thus the radial spread of the ^{68}Ge source event distribution may be taken to be an upper limit on that of the $\bar{\nu}_e$ candidates. This distribution is shown as the dashed histogram in Figure 5.4. The delayed event of a candidate, on the other hand, is mimicked by a 2.2 MeV $^{241}\text{Am}/^9\text{Be}$ neutron capture event. The distribution of these events is shown as the dotted histogram in Figure 5.4.

The distribution of the distances Δr between candidate events was approximated by simulating the prompt and delayed radii based on the ^{68}Ge and $^{241}\text{Am}/^9\text{Be}$ event distributions. The angle θ between the events was generated with a flat distribution in $\cos\theta$. The simulated Δr distribution is shown as the solid curve in Figure 5.4. Since the ^{68}Ge distribution gives a maximal absorption length for the prompt event, the fraction of simulated events falling within 1.6 m of each other represents an upper limit on the efficiency of the Δr cut for $\bar{\nu}_e$ candidates. A lower limit can be obtained by taking the prompt event to have no smearing, or in other words by integrating the $^{241}\text{Am}/^9\text{Be}$ Δr distribution below 1.6 m. In this way, an efficiency of $\epsilon_{\Delta r} = 0.990 \pm 0.002$ is obtained.




Figure 5.4: Δr distributions of ^{68}Ge (dashed histogram) and $^{241}\text{Am}/^9\text{Be}$ 2.2 MeV (dotted histogram) source events (relative to the source position). These are used to simulate the Δr distribution of $\bar{\nu}_e$'s, shown as the solid curve. Upper and lower limits on $\epsilon_{\Delta r}$ are obtained from the fraction of events within 1.6 m for $^{241}\text{Am}/^9\text{Be}$ (99.2%) and MC (98.8%), respectively.




Figure 5.5: E_d distribution of the $\bar{\nu}_e$ candidates. Candidates are required to have $1.54 \text{ MeV} < E_d < 2.64 \text{ MeV}$. The non-gaussian tail at lower energies is due to ATWD dead time.

5.5 Delayed Energy Cut

The delayed event in a $\bar{\nu}_e$ -candidate pair is required to have an energy $E_d = E_\gamma(E_a)$ consistent with a 2.2 MeV np -capture γ via the cut $1.54 \text{ MeV} < E_d < 2.64 \text{ MeV}$. The distribution of E_d for the candidate pairs is plotted in Figure 5.5, with the cuts drawn as the solid lines. The np -capture γ peak of the candidates is non-gaussian in its lower tail. This is due to ATWD dead time when the time Δt between the events is less than $30 \mu\text{s}$. During this time, a fraction of ATWD's involved in digitizing the prompt event are unavailable to record new pulses from the delayed event, resulting in a downward bias in the energy reconstruction. This bias can be clearly seen when the the delayed energy of the candidates is plotted versus Δt , as shown in Figure 5.6.

The efficiency of the delayed energy cut is evaluated from the energy distribution of 2.2 MeV np -capture γ 's in $^{241}\text{Am}/^9\text{Be}$ calibration source data. As in Section 5.3, the $^{241}\text{Am}/^9\text{Be}$ events are grouped into multiplets, except with the time correlation condition relaxed to $\Delta t < 5 \text{ ms}$ for precise background assessment. Both events are required to fit to within 2 m of the source position, and a coincidence pair is required




Figure 5.6: E_d versus Δt for the $\bar{\nu}_e$ candidates. ATWD dead time for $\Delta t < 30 \mu\text{s}$ results in a $\sim 10\%$ downward shift in E_d .

to have $0.5 \mu\text{s} < \Delta t < 660 \mu\text{s}$. A background window with $2 \text{ ms} < \Delta t < 5 \text{ ms}$ measures the rate of pairs in which both prompt and delayed events are np -capture γ 's. The background-subtracted prompt energy distribution is shown in Figure 5.7, and shows a low energy hump due to thermalization of the emitted neutrons, and a second peak at higher energy due to the concurrent emission of a 4.44 MeV γ from $^{12}\text{C}^*$. Figure 5.8 shows the delayed energy distribution in the region of the 2.2 MeV np -capture γ peak for prompt energies with $E_a > 4 \text{ MeV}$ (solid histogram), and for $E_a < 4 \text{ MeV}$ (dashed histogram). The peak shows the same tail due to the energy bias at small Δt as the $\bar{\nu}_e$ candidates, which is more prominent when the prompt event energy is large. The solid lines show the position of the delayed energy cuts in the analysis energy range. Its efficiency is measured to be 0.9925 ± 0.0006 by counting the fraction of $^{241}\text{Am}/^9\text{Be}$ events selected. The uncertainty includes a contribution from the difference in efficiencies for low and high energy prompt events. Since the small- Δt energy bias is worse for events at high radius, an additional 0.6% uncertainty is applied due to the difference in $^{241}\text{Am}/^9\text{Be}$ source data taken at $z = +5.5 \text{ m}$.

An additional correction must be made to account for neutrons captured on carbon




Figure 5.7: $^{241}\text{Am}/^9\text{Be}$ prompt energy distribution. Events below 1 MeV are removed by an analysis cut.




Figure 5.8: E_a distribution of $^{241}\text{Am}/^9\text{Be}$ np -capture γ 's. The solid line is for event pairs with prompt energy > 4 MeV; the dashed line is for prompt energy < 4 MeV. The non-gaussian smearing in the lower tail is more prominent for higher energy prompt events. The $\bar{\nu}_e$ delayed energy limits have been converted to the analysis energy scale and are drawn as solid lines.

nuclei in the scintillator, rather than on protons. Capture on ^{12}C , with cross-section 3.53 ± 0.07 mb, emits a 4.95 MeV γ ; capture on ^{13}C , with cross-section 0.00137 ± 0.00004 mb, emits an 8.17 MeV γ [69]. Both of these are rejected by the delayed energy cut. The cross-section for capture on a proton is much larger, 0.3326 ± 0.0007 b [69]. Using a carbon-to-hydrogen ratio of 1:1.969 for the LS and accounting for the natural abundance of ^{13}C , the efficiency for capture on a proton is 0.9947 ± 0.0001 . Combining this with the efficiency for detection of the np -capture γ , the total efficiency of the delayed energy cut is $\epsilon_{E_d} = 0.987 \pm 0.004$.

5.6 Fiducial Volume Cut

While $\bar{\nu}_e$ events are distributed uniformly throughout the detector, accidental backgrounds are concentrated near the balloon surface, where the rate of events from balloon contaminants and external radioactivity is high. This excess is illustrated by the bright regions and the rim at the balloon radius in the singles event distribution, Figure 4.5. In order to remove the regions of high backgrounds at large radius from the analysis, a fiducial volume cut is applied in which $\bar{\nu}_e$ candidate-pairs are required to have an average radius \bar{r} less than $r_{fid} = 5.5$ m. The inefficiency of this cut for $\bar{\nu}_e$ interactions with protons inside of r_{fid} is negligible: if the event spacing and vertex resolution are much smaller than r_{fid} , and if the vertex resolution does not vary appreciably across r_{fid} , then the efficiency of an event pair generated at distance $d \ll r_{fid}$ inside the cut is approximately equal to 1 minus the efficiency for an event pair occurring a distance d beyond the cut. Moreover, by cutting on the average radius, orthogonality with Δr guarantees no correlation with the spacial correlation cut. Figure 5.9 plots \bar{z} versus $\bar{\rho}^2$ for $\bar{\nu}_e$ candidate pairs. The plot is limited to $\bar{\nu}_e$ candidates with average spherical radius within 6.5 m. The solid curve denotes the position of the fiducial volume cut.

Geometrically, the size of the fiducial volume, which is multiplied by the scintillator's proton density to give the number of target protons (n_p in Equation 2.3), is given by $V_{fid}^{geo} = \frac{4}{3}\pi r_{fid}^3$. Unfortunately, z -axis calibrations of the vertex reconstruction give insufficient information on the uncertainty of off-axis positions to assess the




Figure 5.9: Distribution of $\bar{\rho}^2$ vs. \bar{z} for $\bar{\nu}_e$ candidates. A cut has been applied requiring $\bar{r} < 6.5$ m. The solid curve is the fiducial volume cut, $\bar{r} < 5.5$ m. The cluster of events at $\bar{\rho} = r_{ball}$, $z = 0$ mm is mostly accidental backgrounds; this is the region of the ID that is closest to the cylindrical cavern wall.

error on r_{fid} and hence V_{fid}^{geo} . So instead of using the geometrical formula, the fiducial volume size is measured from the uniformity of the reconstructed distributions of muon-spallation events.

If one may assume that the variation of the differential muon flux, $\frac{d\phi_\mu(E, \mathbf{r})}{dE}$, over all ID positions \mathbf{r} is negligible, then the spallation rate at all \mathbf{r} inside the LS may be taken to be a constant. Additionally, the validity of this assumption requires that the difference in spallation production rates between the LS and BO be negligible, that production in the balloon material and the Kevlar ropes may be neglected, and that the BO be large enough to shield the LS from enhanced production rates in the denser materials in the surrounding detector components. Uniform spallation events should yield a step-function distribution in r^3 that is flat inside the LS and drops to zero outside. The sharpness of the step function will be affected by the non-zero inefficiency/efficiency for detecting events just inside/outside the balloon (respectively), and by the resolution of the vertex reconstruction. Events in the chimney region will also contribute to finite levels beyond the 6.5 m balloon radius.

If r_{fid} is well inside the balloon away from the roll-off, and if inefficiencies near the balloon radius and in the chimney region can be neglected, then one may obtain the fractional size of the fiducial volume f relative to the entire LS volume V_{LS} by taking the ratio of spallation events with $r < r_{fid}$ to the total number of spallation events detected. Then the fiducial volume V_{fid} is given simply by $V_{fid} = fV_{LS}$. The error on the fiducial volume is obtained by propagating the errors on V_{LS} and f .

One class of uniform spallation events used to evaluate f and its error includes the β -decays of ^{12}B , with $Q = 13.4$ MeV and a half life of 20.2 ms, and of ^{12}N , with $Q = 17.3$ MeV and a half life of 11.0 ms. An event is tagged as a $^{12}\text{B}/^{12}\text{N}$ candidate if it has analysis energy in the range $6 \text{ MeV} < E_a < 20 \text{ MeV}$, $nsmx_{ID} > 500$, $nsmx_{OD} < 5$, and $\log_{10}(Q_{tot}) < 5.5$. The latter two requirements ensure that the $^{12}\text{B}/^{12}\text{N}$ event is not itself a muon. It is also required that the time-since-last-muon Δt_μ of the $^{12}\text{B}/^{12}\text{N}$ candidates obeys $2 \text{ ms} < \Delta t_\mu < 52 \text{ ms}$, where a muon in this case is defined as an event having $(\log_{10}(Q_{tot}) > 6)$ OR $(nsmx_{ID} > 600 \text{ AND } nsmx_{OD} > 10)$ (this historical selection criteria essentially selects ID muons). The background for the $^{12}\text{B}/^{12}\text{N}$ candidates is measured with a time window of $300 \text{ ms} < \Delta t_\mu < 500 \text{ ms}$; the time between muons is required to be at least 500 ms to ensure that the background subtraction is valid.

The distribution of r^3 for the $^{12}\text{B}/^{12}\text{N}$ events is histogrammed in Figure 5.10. To within statistical errors, the distribution is flat inside the LS, and it exhibits the expected roll-off at the balloon radius. The solid line in the figure marks the location of the fiducial volume cut, well inside the balloon before the onset of balloon surface complications. The fraction of events falling inside the cut is $f = 0.558 \pm 0.022$.

To test the assumption of uniformity, f is recalculated after removing events following oil muons. This preferentially removes events occurring near the balloon surface, purposefully making the distribution non-uniform. The value of f changes by only 1.7%; this deviation is added in quadrature to the error on f .

The energies of the $^{12}\text{B}/^{12}\text{N}$ events are significantly higher than typical reactor $\bar{\nu}_e$ events. It is possible that the vertex reconstruction algorithm exhibits different behavior at such higher energies, and the error on f should reflect this uncertainty. To obtain a limit on the high energy reconstruction bias, we turn to the position




Figure 5.10: r^3 distribution of $^{12}\text{B}/^{12}\text{N}$ candidates. The solid line shows the position of the fiducial volume cut. The fraction of events reconstructed inside the fiducial volume, including the systematics described in the text, is $f = 0.577 \pm 0.034$.

reconstruction of the prompt and delayed events of the spallation-produced delayed neutron β -decays of ^9Li ($Q = 13.6$ MeV, $\tau_{1/2} = 178$ ms) and ^8He ($Q = 10.7$ MeV, $\tau_{1/2} = 119$ ms). These events, which mimic the $\bar{\nu}_e$ signal (see Section 6.3.2), are selected as $\bar{\nu}_e$ candidates occurring within 2 s following a showering muon, or within 2 s and 3 m of a well-tracked LS muon. This selection naturally contains a small contamination of $\bar{\nu}_e$ candidates, which are mostly removed by requiring that the $^9\text{Li}/^8\text{He}$ events, like the $^{12}\text{B}/^{12}\text{N}$ candidates, have prompt energy $E_d > 6$ MeV. The fiducial volume cut is not applied. If there is a relative radial bias for high-energy events near r_{fid} , then it may be found by comparing the radii r_p and r_d of the prompt and delayed event in $^9\text{Li}/^8\text{He}$ candidate pairs. Figure 5.11 plots $r_p - r_d$ versus the average radius $\bar{r} = \frac{|r_p + r_d|}{2}$. Events within 50 cm (approximately twice the vertex resolution) to either side of the fiducial volume boundary are selected, and the distribution of $r_p - r_d$ for these events is histogrammed and fit to a gaussian in Figure 5.12. The gaussian gives an outward bias of 6.2 ± 3.7 cm for high energy events, amounting to a 3.4% bias in f as measured by the $^{12}\text{B}/^{12}\text{N}$ events. Correcting f for this bias, and including both the size of the




Figure 5.11: Prompt and delayed event radial difference vs. \bar{r}^3 for ${}^9\text{Li}/{}^8\text{He}$ candidates above 6 MeV.

correction and its statistical error in the uncertainty, gives $f = 0.577 \pm 0.034$.

Another uniform source that may potentially be used to measure f is spallation neutrons. Unfortunately, post-muon ringing in the electronics followed by bursts of neutron captures compounds the ATWD dead time problems referred to in Section 5.5, and it is unclear exactly what sort of bias the dead ATWD's have on the vertex reconstruction. While it is easy to find which events are missing waveforms by comparing the number of ATWD's registering hits with $nsmx_{ID}$, these events cannot be simply removed from the data: muons passing through the center of the detector have longer path lengths through the LS, creating more light and more dead time problems. Hence applying cuts to remove events with missing waveforms preferentially removes events near the center of the detector, ruining the uniformity of the distribution essential to the fiducial volume ratio measurement. However, one can simply choose neutrons captured at very long times following the muon, after the electronics have had a chance to recover. Unfortunately, events with missing waveforms are found more than a millisecond following a muon, albeit in decreasing numbers. Removing the first millisecond of neutron captures results in a dramatic




Figure 5.12: Prompt and delayed event radial difference for ${}^9\text{Li}/{}^8\text{He}$ events with \bar{r} within 50 cm of r_{fid} . A slightly positive average indicates that high energy events are reconstructed with slightly larger radius relative to 2.2 MeV np -capture γ 's.

loss of signal. In particular, the signal-to-background ratio near the balloon surface becomes much less than one. Moreover, the loss of signal makes it difficult to determine how much the residual ATWD dead time affects the reconstruction. For these reasons, the measurement with the spallation neutrons is considered as a verification of the measurement with the ${}^{12}\text{B}/{}^{12}\text{N}$ candidates. The spallation neutron candidates are required to have $nsmx_{ID} > 250$ and to have $1 \text{ ms} < \Delta t_\mu < 2 \text{ ms}$ (the muon definition is the same as for the ${}^{12}\text{B}/{}^{12}\text{N}$ candidates). The background window is chosen as $2 \text{ ms} < \Delta t_\mu < 102 \text{ ms}$. The background-subtracted r^3 distribution is shown in Figure 5.13. The small signal-to-background ratio appears as large uncertainties near $r = r_{ball}$. The fiducial volume ratio for these events is $f = 0.569 \pm 0.024$, in good agreement with the ${}^{12}\text{B}/{}^{12}\text{N}$ measurement. Removing the contribution from oil muons changes f by 1.2%.

V_{LS} was measured during detector filling using flow meters on the input pipes for the LS and BO. A secondary measurement obtained V_{LS} by integrating fluid levels in 1 m^3 tanks also employed during oil filling. A third method determined V_{LS} from a




Figure 5.13: r^3 distribution of spallation neutrons. Errors near r_{ball} are large due to an extremely small signal-to-background ratio. ATWD dead time also inhibits proper interpretation of the distribution. The solid line shows the position of the fiducial volume cut. The fiducial volume fraction, including the uniformity systematic, is $f = 0.569 \pm 0.025$.

calculation of the full volume of fluid inside the containment vessel multiplied by the ratio of flow meter measurements for the LS and BO, which pushes to second order many of the systematics related to the flow meter measurements. The three methods give consistent values for V_{LS} , and combined give a volume $1171 \pm 25 \text{ m}^3$.

Combining V_{LS} and f gives a fiducial volume size of $676 \pm 42 \text{ m}^3$. As discussed in Chapter 3, the scintillator has proton number density $(6.611 \pm 0.007) \times 10^{28} \text{ m}^{-3}$. This gives a total of $n_p = (4.47 \pm 0.28) \times 10^{31}$ target protons for the reactor $\bar{\nu}_e$ measurement.

5.7 Spallation Cuts

Muons crossing the detector leave in their wake spallation fragments, including neutrons and longer-lived isotopes such as ^8He and ^9Li that may mimic the $\bar{\nu}_e$ signal [70]. The spallation backgrounds within the reactor $\bar{\nu}_e$ energy range will be described in detail in Chapter 6. In anticipation of these backgrounds, the selection criteria

imposed to remove them are discussed here.

It is beneficial to veto the full detector for a short period after a muon passes to remove the short lived backgrounds, in particular the copiously produced spallation neutrons. Extending these full-detector vetoes to longer times to remove long-life isotopes would dramatically reduce the livetime for the $\bar{\nu}_e$ measurement. Detector livetime may be salvaged while still retaining a high efficiency to remove long-life spallation products by vetoing for an extended period only along the muon's track. When a muon's track can not be reconstructed, or when muons create showering events that generate high levels of activity, the whole detector must be vetoed. It is also prudent to apply such a veto at the start of a run and after gaps. The specific spallation cuts used in this analysis are as follows:

- Veto the full detector for 2 ms following any muon.
- For LS muons whose track could be fit (see Section 4.5), veto a 3 m radius cylinder around the track for 2 s.
- For showering muons, LS muons with poorly reconstructed tracks, and following run starts and gaps, veto the full detector for 2 s.

The small backgrounds remaining after these selection criteria are applied will be discussed in Chapter 6. Since these cuts are in essence detector vetoes, their affect on the $\bar{\nu}_e$ measurement is treated as a correction to the livetime and fiducial volume rather than an efficiency. The dead time introduced by the first and third cuts is simple to calculate by summing vetoed periods following muons. However, the cylindrical cuts veto only a section of the fiducial volume for a short time, reducing the total exposure (volume times time) by an amount that depends on the impact parameter of the muon track. The dead time and exposure loss of all three cuts is calculated using a Monte Carlo method in which events are randomly distributed throughout each run, with a uniform position distribution throughout the 5.5 m radius fiducial volume. The actual muons in the run are used to calculate the fraction of Monte Carlo events not removed by the spallation cuts. This live-fraction is then treated as a livetime correction for the run from which it was calculated. The uncertainty in




Figure 5.14: Live-fraction distribution, weighted by livetime, for all runs in the data set reported here. The group of runs with lower live-fractions are due to lower muon tracking efficiencies resulting from an early ATWD launch setting described in Chapter 3.

the calculation depends essentially on the statistics of the Monte Carlo and is negligible. The live-fraction distribution for the runs used in this analysis is plotted in Figure 5.14. The cluster of runs with lower live-fraction are early runs with lower muon tracking efficiency due to the ATWD launch setting described in Chapter 3. The weighted average live-fraction for this data set is 85.53%. Multiplying the live-fractions of each run by its livetime, calculated as described in Section 5.1, gives a total effective livetime of 414.6 ± 0.2 days. Multiplying by the 527 ton fiducial volume mass gives a total exposure of 599 ton-years.

5.8 Prompt Energy Analysis Threshold

One final cut requires $\bar{\nu}_e$'s to have a prompt energy above 2.6 MeV and less than 8.5 MeV. As discussed in the Chapter 6, this cut eliminates most accidental backgrounds and virtually all potential non-reactor $\bar{\nu}_e$ backgrounds. The distribution of delayed versus prompt event energies for the $\bar{\nu}_e$ candidates is shown in Figure 5.15. The




Figure 5.15: Prompt and delayed energies of $\bar{\nu}_e$ candidates. Events at low energies are mostly accidental backgrounds, which are discussed in Chapter 6. The two events with $E_d \approx 5$ MeV are most likely due to neutron captures on ^{12}C .

energies have been corrected with the appropriate conversion functions, $E_p(E_a)$ for the prompt event and $E_\gamma(E_a)$ for the delayed event. The cuts are drawn as solid lines. The event pairing step described in Section 5.2 imposes an implicit cut of $E_a > 0.8$ MeV in the Figure. The contamination from accidental backgrounds at low energies can be clearly seen. Above 2.6 MeV, the $\bar{\nu}_e$ events show a clean separation from the accidentals. The two events with $E_d \approx 4.5$ -5 MeV are most likely due to neutron captures on ^{12}C . The prompt event energy distribution is presented later in Figure 8.2.

5.9 Total $\bar{\nu}_e$ Detection Efficiency

The total $\bar{\nu}_e$ detection efficiency is the product of the triggering efficiency, the reconstruction efficiency, and the efficiencies of the $\bar{\nu}_e$ selection criteria described above (except for the spallation cut correction, which is applied to the livetime). The triggering efficiency is the only of these with significant energy dependence; it is obtained




Figure 5.16: Efficiency of the trigger $nsmx > 200$ versus analysis energy. The efficiency above 1 MeV is ~ 1 .

Table 5.2: $\bar{\nu}_e$ Detection Efficiency

| | |
|---|-------------------|
| Triggering efficiency ($nsmx_{ID} > 200$) | ~ 1 |
| Reconstruction efficiency | 0.999 ± 0.001 |
| Time correlation ($0.5\mu s < \Delta t < 660\mu s$) | 0.956 ± 0.002 |
| Space correlation ($\Delta r < 1.6$ m) | 0.990 ± 0.002 |
| Delayed energy cut (1.54 MeV $< E_d < 2.64$ MeV) | 0.987 ± 0.004 |
| Fiducial volume cut ($r < 5.5$ m) | ~ 1 |
| Total | 0.933 ± 0.004 |

from special low-threshold runs by dividing the total number of events within the fiducial volume at a particular energy by the number with $nsmx > 200$. Figure 5.16 plots the triggering efficiency versus analysis energy. As can be seen, the energy dependence is confined to energies below 1 MeV, and the efficiency above E_p 2.6 MeV is ~ 1 . The parameter $\epsilon(E_{\bar{\nu}_e})$ in Equation 2.3 is therefore taken to be a constant over the energy range of interest and may be moved outside the integral. The various contributions to ϵ and their uncertainties are listed in Table 5.2. The full $\bar{\nu}_e$ detection efficiency is 0.933 ± 0.004 .

Chapter 6

Backgrounds

The selection criteria described in Chapter 5 are tuned to select reactor $\bar{\nu}_e$ events with as high efficiency as possible, and as little as possible of anything else. The backgrounds which nevertheless survive the selection criteria fall into three categories. The first is $\bar{\nu}_e$'s not emitted by known reactors. The next category is accidental backgrounds, pairs of single events which randomly happen to pass the selection criteria. The final category are non- $\bar{\nu}_e$ correlated backgrounds which produce a prompt-delayed coincidence pair with some efficiency for passing the $\bar{\nu}_e$ selection criteria. The candidate backgrounds for each category and their expected rates are discussed below.

6.1 $\bar{\nu}_e$ Backgrounds

6.1.1 Man-made $\bar{\nu}_e$ Sources

Potential man-made $\bar{\nu}_e$ background sources include nuclear material or reactors that are not accounted for in the signal estimation presented in Chapter 7. In particular, the estimation ignores contributions from spent nuclear fuel stored in Japan, and naval reactors approaching KamLAND.

Nuclear fuel continues to emit a significant amount of $\bar{\nu}_e$'s for years after being processed in a reactor [71]. The accumulation of spent nuclear fuel over many years of reactor operations pose a potential background to reactor $\bar{\nu}_e$ studies. However, the

highest energy decay from spent fuel has a Q of 3.5 MeV. The contribution of spent fuel $\bar{\nu}_e$'s above the 2.6 MeV prompt energy threshold is very small, and hence this background may be ignored in this analysis.

The operation of naval reactors on board nuclear-powered vessels such as submarines or aircraft carriers is classified, as are typically their positions. There is no consistent way to include these reactors in the signal estimation. However, given the size of these reactors, it has been shown that they are only a problem for KamLAND if a vessel enters Toyama Bay (the closest body of water to the detector) and remains there for an extended period [72]. A multi-reactor aircraft carrier in Toyama Bay would draw much attention, and its period of stay could be easily vetoed in the data. A submarine may go unnoticed, and would comprise up to 10% of KamLAND's signal if it were to remain stationary, running its reactors at full power. It is assumed that this is an unreasonable mode of operation for military vessels, and this background is hence ignored.

6.1.2 Geological $\bar{\nu}_e$ Sources

There are two potential significant geological antineutrino backgrounds: "geo-neutrinos" [54], and $\bar{\nu}_e$'s emitted by a hypothetical "geo-reactor" at the Earth's core [73]. Geo-neutrinos are $\bar{\nu}_e$'s emitted in the decay chain of uranium and thorium deposits throughout the Earth. The geo-neutrino spectrum extends up to energies of $E_{\bar{\nu}_e} = 3.3$ MeV (see Figure 6.1), with some geological models predicting tens of $\bar{\nu}_e$'s per 10^{32} protons yr for 100% efficiency (see, for example, [74]). The 2.6 MeV threshold imposed on the prompt events eliminates this background for the reactor $\bar{\nu}_e$ analysis presented here (this is, in fact, the purpose of the threshold). As KamLAND stands to be the first experiment to detect geo-neutrinos, data below the threshold are not presented in anticipation of a future analysis of $\bar{\nu}_e$ energies in the geo-neutrino range.

A geo-reactor would have a $\bar{\nu}_e$ spectrum indistinguishable from that produced by man-made reactors. However, while the signal from Japanese reactors varies in time as reactors are turned off and on (see Figure 7.4), a geo-reactor should give




Figure 6.1: Prompt energy spectrum of geo-neutrino backgrounds. The normalization of each spectrum is arbitrary. The number of events above the 2.6 MeV analysis threshold is negligible.

a constant rate. The expected background at KamLAND for a typical 3 TW_{th} geo-reactor is only ~ 5 events over the data taking period reported here, much smaller than the Japanese reactor $\bar{\nu}_e$ flux and on the same order as other backgrounds discussed in this chapter. This background is addressed with a time-variation analysis (see Section 8.2), but since the existence of the geo-reactor has not been experimentally verified, its potential contribution is otherwise ignored.

6.1.3 Cosmic $\bar{\nu}_e$ Sources

The first extraterrestrial sources of $\bar{\nu}_e$'s to be discussed are the sun and cosmic rays which are the sources of solar and atmospheric neutrino oscillation experiments. Other sources include supernovae, with close supernovae producing event bursts and far supernovae producing supernova relic neutrinos [75], and the cosmic neutrino background [14]. However, both supernova relic neutrinos and the cosmic neutrino background give rates too low to be detected or $\bar{\nu}_e$ energies far above the reactor $\bar{\nu}_e$ range, and an event burst from a near-by supernova, should KamLAND be lucky

enough to detect one, could easily be vetoed from the reactor $\bar{\nu}_e$ analysis.

The conventional reactions that fuel the sun produce only neutrinos, but unconventional reactions such as spin-flavor precession and neutrino decay [14] may convert the neutrinos into antineutrinos that may pose a background for reactor $\bar{\nu}_e$ studies. A search for $\bar{\nu}_e$'s from the sun has been performed with KamLAND in the energy range $8.3 \text{ MeV} < E_{\bar{\nu}_e} < 14.8 \text{ MeV}$ [53]. No candidates were observed for an expected background of 1.1 ± 0.4 events. Assuming that the $\bar{\nu}_e$ spectral shape follows the standard solar ${}^8\text{B}$ spectrum [76], this limits the $\bar{\nu}_e$ flux to be less than $3.7 \times 10^2 \text{ cm}^{-2} \text{ s}^{-1}$ at the 90% confidence level. Since solar $\bar{\nu}_e$'s have never been observed, this background is ignored in this analysis.

Unlike the sun, the cosmic rays that generate the atmospheric neutrino signal also produce $\bar{\nu}_e$'s via reactions like $\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu$. The flux of atmospheric $\bar{\nu}_e$'s at the Kamioka site has been measured above 50 MeV in [77]; the flux decreases for lower energies. Taking the flux above 50 MeV as a conservative upper limit, the expected atmospheric $\bar{\nu}_e$ rate at KamLAND is some 6 orders of magnitude below the reactor signal. This background may therefore also be safely ignored.

6.2 Accidental Backgrounds

The accidental background is comprised of pairs of single events that randomly pass the $\bar{\nu}_e$ selection criteria. For example, two 2.6 MeV γ 's from decays of ${}^{208}\text{Tl}$ occurring within 660 μs and 1.6 m of each other will look very much like a $\bar{\nu}_e$. As shown in Figures 4.5 and 4.9, the accidental background rate is highest at low prompt energies and near the balloon edge, where event rates due to radioactivity are highest. Most of these are removed by the fiducial volume and prompt energy cuts.

The Δt distribution for the candidates after all cuts have been applied, Figure 5.2, consists of an exponential neutron capture time distribution on top of an approximately flat background. A flat background suggests either correlated event pairs with a very long time constant, or events with no time correlation at all. The latter comprise the residual accidental background whose rate we want to measure. Assuming that the former make up a negligible fraction of the flat background, the accidental




Figure 6.2: Analysis energy distribution of accidental background events. The integrated rate above $E_p = 2.6$ MeV is 0.0081 ± 0.0001 events per day.

background rate is obtained by applying the usual selection criteria except changing the time-correlation cut to a later, or “off-time”, window. The analysis energy distribution of events selected with the off-time window $0.01 \text{ s} < \Delta t < 20 \text{ s}$ is shown in Figure 6.2. This method gives a total accidental background rate for $E_p > 2.6$ MeV of 0.0081 ± 0.0001 events per day. The total expected accidental background count for the data period reported here is obtained by multiplying by the livetime uncorrected for spallation cuts (since those cuts are already applied to the accidental candidates), giving a total of 3.91 ± 0.05 events.

6.3 Correlated Non- $\bar{\nu}_e$ Backgrounds

In addition to the $\bar{\nu}_e$ backgrounds described in Section 6.1, there are several classes of correlated events with non-zero efficiency to pass the $\bar{\nu}_e$ selection criteria. The most problematic interactions are those with a neutron in the final state, since they will automatically create event pairs mimicking a $\bar{\nu}_e$. Such interactions may occur from natural radioactivity in the detector, or may be induced by muon spallation. In

regard to natural radioactivity, background levels were assessed for (α, n) reactions, (γ, n) reactions, spontaneous fission, neutron emission, and short-life nuclei. Of these, only the reaction $^{13}\text{C}(\alpha, n)^{16}\text{O}$ is significant. The muon-induced backgrounds are pushed to very low levels thanks to KamLAND's 2700 m.w.e. overburden and the spallation cuts described in Section 5.7. Short-lived spallation products, such as the copiously-produced spallation neutrons, are removed with high efficiency by the 2 ms spallation cut following all muons. Still, untagged muons or long-life spallation products may pose non-ignorable backgrounds [70]. Of all possible nuclei that may be spalled from the elements composing KamLAND's LS, the only three that pose a threat for KamLAND are ^9Li , ^8He , and fast neutrons. The backgrounds from these and the $^{13}\text{C}(\alpha, n)^{16}\text{O}$ reaction are discussed in the following sections.

6.3.1 Backgrounds from $^{13}\text{C}(\alpha, n)^{16}\text{O}$

The primary source of α 's in KamLAND's scintillator is decays of the radon daughter ^{210}Po . When it does not interact with ^{13}C , the 5.304 MeV α emitted in the decay of ^{210}Po is quenched by a factor of ~ 13 , well below the triggering threshold. The rate of these decays in the fiducial volume is obtained from special runs with very low $n\text{sum}_{ID}$ thresholds in which no waveforms were collected. As shown in Figure 6.3, the α -peak appears at $n\text{sum}_{ID} \approx 80$, and is fit nicely by a gaussian plus a linear background. For early runs, 5.490 MeV α 's from decays of ^{222}Rn also appear in this peak. To disentangle the ^{210}Po events from the ^{222}Rn and to obtain the time variation of their rates, the ^{210}Po α -peak height was measured for several low-threshold runs throughout the data taking period. The peak heights are plotted versus time since detector filling as the solid points in Figure 6.4. Under the assumptions that all of the ^{210}Po comes from ^{210}Pb , and that levels of ^{222}Rn have been decaying in time t since the end of detector filling, the height of the peak, A_{Po+Rn} , is fit to the function

$$A_{Po+Rn}(t) = A_{Rn}e^{-\lambda_{Rn}t} + A_{Po}\lambda_{Bi}\lambda_{Po} \left[\frac{e^{-\lambda_{Pb}t}}{(\lambda_{Bi} - \lambda_{Pb})(\lambda_{Po} - \lambda_{Pb})} + \frac{e^{-\lambda_{Bi}t}}{(\lambda_{Pb} - \lambda_{Bi})(\lambda_{Po} - \lambda_{Bi})} + \frac{e^{-\lambda_{Po}t}}{(\lambda_{Pb} - \lambda_{Po})(\lambda_{Bi} - \lambda_{Po})} \right], \quad (6.1)$$




Figure 6.3: $^{210}\text{Po}/^{222}\text{Rn}$ quenched α decay peak in the $nsmx_{ID}$ distribution of low-threshold data. The peak fits nicely to a gaussian plus a linear background.

where λ_{Rn} , λ_{Bi} , λ_{Po} , and λ_{Pb} are the decay constants for ^{222}Rn , ^{210}Bi , ^{210}Po , and ^{210}Pb , respectively, and A_{Rn} and A_{Po} are the fit parameters. The production and decay of the various isotopes involved in the fit are plotted as solid lines in the figure. The beginning of data taking occurs after the ^{222}Rn line falls to zero so the α rate is assumed to be due entirely to decays of ^{210}Po .

The activity obtained from Figure 6.4 gives the ^{210}Po decay rate throughout the entire LS. A correction must be made to obtain the activity within the fiducial volume only. To do this, pre-scaled waveform data taken during a special low threshold run were reconstructed. The events were divided according to their position into spherical shells, and the ^{210}Po rate in each shell was determined again from its $nsmx_{ID}$ peak. The results are plotted in Figure 6.5, in which an excess at the balloon surface can be seen. The biases in the vertex fitting at such low energies were limited to less than 1 m from comparisons of prompt and delayed event vertices in ^{212}Bi - ^{212}Po coincidences. Combining the position information with the total rate determined by the fit to Equation 6.1, the total ^{210}Po rate in the fiducial volume was determined to be 33 ± 3 Bq.




Figure 6.4: Variation in time of the height of the $^{210}\text{Po}/^{222}\text{Rn}$ peak in the $nsm.x_{ID}$ distribution (Figure 6.3). The dashed green line is the result of a fit to Equation 6.1, while the other lines denote the production/decay of the various isotopes involved in the fit. The ^{222}Rn decays away before the start of data taking.




Figure 6.5: Reconstructed radial distribution of ^{210}Po events. The rate is relatively constant throughout the fiducial volume, with a slight excess at the balloon surface.

The $^{13}\text{C}(\alpha, n)^{16}\text{O}$ reaction has a Q of 2.216 MeV. The prompt event in a $^{13}\text{C}(\alpha, n)^{16}\text{O}$ background event consists of any scintillation produced by the incident alpha prior to the interaction, along with scintillation of recoil protons involved in thermalization of the neutron in the final state. If the neutron energy is high enough, it may excite an atom of ^{12}C during thermalization, emitting a 4.4 MeV γ . If the incident α does not lose too much energy before interacting with the ^{13}C , the ^{16}O may be left in its first or second excited states $^{16}\text{O}^*$, releasing 6.049 MeV or 6.130 MeV in γ 's, respectively. The full $^{13}\text{C}(\alpha, n)^{16}\text{O}$ spectrum was determined from Geant4 [68] Monte Carlo simulations using cross-section data from [78]. The alpha particle was tracked to its interaction energy, and its energy deposition was quenched using the Birks constant determined in Section 4.4.4. The angular distribution of the outgoing neutron was simulated using the Legendre polynomial coefficients given in [79]. Then the neutron was passed back to Geant4 for tracking and quenching was again applied to the recoil protons. Since data exist only for quenching of α 's (i.e. from ^{212}Po and ^{214}Po), the uncertainty in the quenching factor for protons is estimated to be 10%. The prompt energy spectrum is then normalized according to the ^{210}Po rate, and the result is shown in Figure 6.6 with the detector resolution applied. The uncertainty in the normalization of the 4.4 MeV peak and the lower energy peak is estimated to be 32% and is dominated by the uncertainties in the ^{210}Po rate, the $^{13}\text{C}(\alpha, n)^{16}\text{O}$ cross-section, and the neutron angular distribution. The normalization of the 6 MeV peak is much less certain due to large uncertainties in the $^{13}\text{C}(\alpha, n)^{16}\text{O}^*$ branching ratio. The total event count above the 2.6 MeV analysis energy threshold due to $^{13}\text{C}(\alpha, n)^{16}\text{O}$ events is estimated to be 8.6 ± 5.9 events.

That some small fraction of the α 's emitted in ^{210}Po decays may go on to interact with the 1.1% natural abundance of ^{13}C in the scintillator and produce a significant background for reactor $\bar{\nu}_e$ studies was somewhat unexpected. Due to its second-order nature, this background went unnoticed in earlier KamLAND analyses [55].




Figure 6.6: Prompt energy spectrum of $^{13}\text{C}(\alpha, n)^{16}\text{O}$ background events with detector resolution applied. The low energy peak is from thermalization of the neutron and is subject to 10% uncertainties in the proton quenching factor. Its normalization, along with that of the 4.4 MeV peak due to $^{12}\text{C}(n, n + \gamma)^{12}\text{C}$, has an uncertainty of 32%. The peak at 6 MeV is due to $^{13}\text{C}(\alpha, n)^{16}\text{O}^*$ and has a large uncertainty in height.




Figure 6.7: Δt_μ distribution of ${}^9\text{Li}/{}^8\text{He}$ and $\bar{\nu}_e$ events. The distribution is fit to a constant background plus an exponential decay with $\tau = \tau_{{}^9\text{Li}}$.

6.3.2 ${}^9\text{Li}/{}^8\text{He}$ Events

The long-life spallation products ${}^9\text{Li}$ and ${}^8\text{He}$ have 49.5% and 16% branching ratios, respectively, for delayed-neutron β -decay, which mimics the $\bar{\nu}_e$ signal. Since the ${}^9\text{Li}$ decay has $Q = 13.6$ MeV and $\tau_{1/2} = 178$ ms and ${}^8\text{He}$ has $Q = 10.7$ MeV and $\tau_{1/2} = 119$ ms, the two decays are largely indistinguishable and are handled simultaneously. The purpose of the 2 second cylindrical and full volume vetoes following muons described in Section 5.7 is to remove these backgrounds. If the cut is removed and the time Δt_μ to all muons in the past 2 seconds is histogrammed for all $\bar{\nu}_e$ candidates, as has been done in Figure 6.7, the ${}^9\text{Li}/{}^8\text{He}$ contamination can be clearly seen. The decay time is consistent with that of ${}^9\text{Li}$. The energy spectrum of ${}^9\text{Li}/{}^8\text{He}$ events selected by inverting the 2 s spallation cuts is shown in Figure 6.8. The reactor $\bar{\nu}_e$ background was subtracted by normalizing the $\bar{\nu}_e$ candidate spectrum by 1 minus the weighted average live-fraction after the spallation cuts. The analysis energies were converted with the electron energy function $E_{e^-}(E_a)$. The β^- energy spectrum of the ${}^9\text{Li}$ decay, shown as the solid line, fits the data reasonably well; there is no indication that any significant amount of ${}^8\text{He}$ is produced.




Figure 6.8: Energy distribution of ${}^9\text{Li}/{}^8\text{He}$ events. The solid line shows the energy spectrum of the β^- emitted in the ${}^9\text{Li}$ decay; its normalization was fit to the data. The largest contributions to the χ^2 are from bins at high energy. There is no indication of ${}^8\text{He}$.

The fit of the Δt_μ distribution shown in Figure 6.7 to a constant background plus an exponential with time constant corresponding to the ${}^9\text{Li}$ decay time gives a total of 292 ± 23 events prior to spallation cuts. To evaluate the background remaining when the spallation cuts are applied, the events are separated into statistically independent classes: events occurring within 2 s of a showering muon or a poorly tracked muon, and events among the remaining candidates falling within 3 m of a muon track. Fits to the Δt_μ distributions give 266 ± 19 events in the former class, and 30 ± 8 in the latter, indicating that most of the ${}^9\text{Li}/{}^8\text{He}$ candidates are generated by showering muons.

Again assuming the ${}^9\text{Li}$ decay time, the efficiency of a 2 s veto for removing these events is 0.99958. The efficiency is even higher for the faster ${}^8\text{He}$ decay. The efficiency of the 3 m cylinder cut is assessed by measuring the fraction of spallation neutrons reconstructed more than 3 m from their generating muon track. Spallation neutrons are selected as events having $nsmx_{ID} > 250$ occurring between 0.4 and 2 ms following a non-showering ID muon with a good track reconstruction status. The lower limit




Figure 6.9: L_μ distribution of spallation neutrons. The fraction of events falling within 3 m is used to estimate the efficiency of the cylindrical veto along muon tracks for selecting ${}^9\text{Li}/{}^8\text{He}$ events.

of 0.4 ms on the time correlation eliminates possibly biased events missing waveforms due to ATWD dead time. A background sample is selected with the same criteria, except with $2 \text{ ms} < \Delta t_\mu < 102 \text{ ms}$. The background subtracted distribution of the neutron distances L_μ from the muon track is shown in Figure 6.9. The fraction with $L_\mu < 3 \text{ m}$ is 0.984 ± 0.014 . Since neutrons are lighter than ${}^9\text{Li}$ or ${}^8\text{He}$ nuclei and are hence expected to have a broader L_μ distribution, this fraction is taken as a conservative estimate of the fraction of ${}^9\text{Li}/{}^8\text{He}$ events with $L_\mu < 3 \text{ m}$. Including the timing cut, the total efficiency of the cylinder cut for selecting ${}^9\text{Li}/{}^8\text{He}$ events is 0.983 ± 0.018 .

Combining the ${}^9\text{Li}/{}^8\text{He}$ event counts (N) before the spallation cuts are applied with the removal efficiencies (ϵ) of the cuts leaves $N(1 - \epsilon)/\epsilon$ events in the data set. For the highly efficient 2 s full-volume veto following showering muons and poorly reconstructed muons, the residual ${}^9\text{Li}/{}^8\text{He}$ background is 0.111 ± 0.008 events. For the cylindrical cut around well-reconstructed tracks, the remaining background is 0.5 ± 0.1 events. The total ${}^9\text{Li}/{}^8\text{He}$ background among the $\bar{\nu}_e$ candidates is 0.6 ± 0.1

events.

6.3.3 Fast Neutrons

Muons that go untagged by the detector can produce spallation neutrons which enter the ID, thermalize, then undergo their usual random walk and capture on a proton. Since the efficiency for detecting a muon in the ID is negligibly close to 100% due to their high light yield, such a neutron would have to be generated either by an OD-crossing muon missed by the OD, or by a muon passing through the rock just outside of the OD. In order to make it all the way into the LS, through meters of water and oil buffers, and still produce enough scintillation before thermalizing to create an event above the 2.6 MeV analysis threshold, the neutron would have to have a rather high energy, hence the name “fast” neutrons. Fast neutrons are a background for reactor $\bar{\nu}_e$'s when the scintillation from recoil protons during thermalization produces a prompt event in the analysis energy range.

The number of fast neutrons from muons passing through the OD water but not the ID can be obtained by dividing the rate of $\bar{\nu}_e$ candidates immediately following OD-only muons by the OD muon efficiency ϵ_{OD} . By reversing the OD muon veto in the candidate selection, i.e. requiring that candidates occur within 2 ms of an OD muon, 4 events were found. The OD efficiency was studied with Monte Carlo simulations and was found to be 0.9951 with all tubes functioning, but had dropped to 0.9883 by summer 2003 due to loss of tubes in the OD [81]. Allowing for continued efficiency loss for data taking through the end of 2003 reported here, an OD efficiency of 0.990 ± 0.005 is used. This gives a count of $N_{water} = 4 \pm 2$ fast neutrons generated in the water that enter the ID, of which $N_{water}(1 - \epsilon_{OD}) = 0.04 \pm 0.03$ go untagged.

An upper limit can be obtained on the number of fast neutrons spalled by muons passing through the rock just outside of the OD by assuming that production rates scale according to density and using the maximum attenuation length of neutrons in rock, $\lambda_{rock} = 40$ cm [80]. The active volume of rock producing fast neutrons is less than $A_{cavern}\lambda_{rock}$, where $A_{cavern} = 1835$ m² is the area of the cavern wall. Considering that less than 1/4 of neutrons spalled by muons passing through rock are

Table 6.1: Total Backgrounds

| Background | Number of Events |
|---|------------------|
| Accidental | 3.91 ± 0.05 |
| $^{13}\text{C}(\alpha, n)^{16}\text{O}$ | 8.5 ± 5.9 |
| $^9\text{Li}/^8\text{He}$ | 0.6 ± 0.1 |
| Fast n | < 0.1 |
| Total | 13.2 ± 5.9 |

aimed towards the ID, the limit on fast neutrons from the rock is obtained by scaling the rate from the water according to

$$N_{\text{rock}} < \alpha \frac{A_{\text{cavern}} \lambda_{\text{rock}} \rho_{\text{rock}}}{4V_{\text{water}} \rho_{\text{water}}} N_{\text{water}}, \quad (6.2)$$

where ρ denotes density, $V_{\text{water}} = 2600 \text{ m}^3$ is the volume of water in the OD, and α is a factor that accounts for the extra average attenuation of a neutron emitted from the rock with respect to neutrons emitted throughout the OD water; its value must lie between 1 and $e^{-d_{\text{OD}}/\lambda_{\text{rock}}} \sim 0.3$, where $d_{\text{OD}} 50 \text{ cm}$ is typical path length across the OD for a neutron emanating from the rock. Since N_{rock} here is the total number of spallation neutrons produced in the rock, it includes neutrons spalled by muons that go on to cross the OD. Consideration of solid angle shows that less than 5% of muons passing within $2\lambda_{\text{rock}}$ can do so without also crossing the detector. The number of neutrons spalled in the rock that go untagged is then $N_{\text{rock}} \times 5\% = 0.03 \pm 0.02$ events. Adding the contributions from the rock and the water, the total fast neutron rate in the data taking period reported here is less than 0.1 events.

6.4 Total Backgrounds

The backgrounds from non-negligible sources discussed in this chapter are listed in Table 6.1. Small backgrounds from geo-neutrinos and naval reactors or other hypothetical sources such as a geo-reactor are ignored. The total background for the data period reported here is 13.2 ± 5.9 events. In this estimate, the upper limit on the fast neutron background has been included in the uncertainty only.

Chapter 7

Reactor Signal Estimation

7.1 Fission Yields

Nuclear reactors create power by converting the heat released in the fissions and subsequent decay chains of heavy isotopes such as ^{235}U into electricity. In addition to heat, the reactions also produce the antineutrinos that comprise KamLAND's signal. For example, in the most common fission of ^{235}U , two large fragments and two neutrons are emitted in the final state:



Typically, one of the fragments X_1 or X_2 is heavier than the other; the mass number distribution for both fragments is shown in Figure 7.1 [35]. The stable elements nearest the peaks of the distribution are ^{94}Zr and ^{140}Ce , which together have 98 protons and 136 neutrons. Considering that the original ^{235}U nucleus contains 92 protons and 143 neutrons, on average about six neutrons will be converted to protons in order to reach stable matter. If the energy spectrum of the roughly six emitted $\bar{\nu}_e$'s is known, then calculating the signal at KamLAND amounts to counting the number




Figure 7.1: Mass number distribution of fission fragments of ^{235}U (Equation 7.1), in percent-yield. Since the distribution accounts for the masses of both fission fragments, it is normalized to 200%. [35]

of fissions of each isotope at each reactor. Referring back to Equation 2.3, we write

$$I_i(E_{\bar{\nu}_e}, t) = \sum_k^{\text{isotopes}} f_{i,k}(t) \frac{dN_{\bar{\nu}_e,k}(E_{\bar{\nu}_e})}{dE_{\bar{\nu}_e}}, \quad (7.2)$$

where $f_{i,k}(t)$ is the instantaneous fission rate of isotope k at reactor i , and $\frac{dN_{\bar{\nu}_e,k}(E_{\bar{\nu}_e})}{dE_{\bar{\nu}_e}}$ is the differential neutrino emission of isotope k , measured in $[\bar{\nu}_e]$'s per unit energy per fission]. The sum over k runs over the six isotopes ^{235}U , ^{238}U , ^{239}Pu , ^{241}Pu , ^{106}Ru , and ^{144}Ce . The uranium and plutonium isotopes together account for more than 99.9% of $\bar{\nu}_e$ -emitting fissions in the reactor. The last two isotopes are longer lived and are important for periods immediately following reactor shut-downs. Their contribution can be calculated from the uranium and plutonium fissions using published fission yields (see, for example, [82]) and applying the appropriate decay times.

Since each fission and the chain of decays following it releases energy that heats up the reactor core, the instantaneous thermal power of the reactor $P_{th,i}(t)$ acts as a constraint on the sum of the fission rates $f_{i,k}(t)$ (for the power constraint, fissions from

^{242}Pu must be included, although they do not emit $\bar{\nu}_e$'s). If the initial fuel composition is known then the $f_{i,k}(t)$ can be determined by tracking the depletion of the uranium isotopes and the breeding of the plutonium isotopes. Previous work has shown that detailed simulations of the reactor core can track the burn-up to within the error on $P_{th,i}(t)$ [44]. However, it is practically impossible in a finite time to simulate all 53 commercial Japanese reactors contributing to KamLAND's signal¹. Instead, a burn-up model constructed by Tokyo Electric Power Company, Inc. (TEPCO) and Tohoku University is used to estimate the $f_{i,k}(t)$. The inputs to the model are the initial fraction of new fuel and ^{235}U concentration, and the instantaneous and integrated thermal power output of the reactor. These data are provided for all 53 reactors and are utilized according to special agreements between the reactor operators and Tohoku University as a member of KamLAND. During stable reactor operation, data are provided on a weekly basis; when the reactors are starting or stopping, hourly data are provided. The model reproduces detailed calculations to within an accuracy of 1%. The simulation results for a few power cycles of a typical Japanese 3 GW_{th} boiling water reactor (BWR) are shown in Figure 7.2. As expected, plutonium isotopes are generated as the uranium isotopes are depleted. The periods during which the rates fall to zero correspond to reactor shut-downs for maintenance and refueling, during which a fraction of the fuel is replaced while the remaining fuel is redistributed to maximize burn-up efficiency.

The power is determined from a calculation of the energy balance across the steam generator or reactor vessel, the uncertainty of which is dominated by the feed-water flow measurement. The time dependence of $P_{th,i}(t)$ is cross checked with publicly available independent records of electricity generation [83]. The uncertainty on the power measurements is quoted by the reactor operators as 2%, in accordance with the margin of error for safe operation recognized by regulatory agents. This conservative estimate assumes that measurements made at different reactors may be correlated; efforts are underway to reassess the validity of such assumptions, and in the future a smaller uncertainty may be achievable.

¹Much of the difficulty lies in the fact that existing detailed simulations are proprietary or require proprietary information. Reactor operators are reluctant to grant third parties access to these codes and information.




Figure 7.2: Calculated fission rates of heavy isotopes over a few power cycles of a typical 3 GW_{th} BWR. The periods during which the rates fall to zero correspond to reactor shut-downs for maintenance and refueling.

The contribution from reactors in Korea is estimated with a 10% error from electricity generation records [84] under the assumption that they have the same average fuel composition as Japanese reactors. The contribution from Japanese research reactors and from reactors around the world is estimated with a conservative 50% uncertainty from publicly available data [51] with the same fuel composition assumption used for Korean reactors. The fractional contribution of all non-Japanese reactors and Japanese research reactors combined is estimated to be $(4.5 \pm 0.6)\%$.

The simulated average relative fission yields for the run period used in this analysis are $^{235}\text{U} : ^{238}\text{U} : ^{239}\text{Pu} : ^{241}\text{Pu} = 0.563 : 0.079 : 0.301 : 0.057$. The total uncertainty on the fission yield calculations is estimated to be 2.3%, obtained by adding in quadrature the 2.0% error on the reactor power, the 1.0% error on the burn-up model relative to detailed reactor core simulations, and the 0.6% uncertainty in the contribution from non-Japanese commercial reactors and research reactors.




Figure 7.3: $\bar{\nu}_e$ energy spectra of ^{235}U , ^{238}U , ^{239}Pu , and ^{241}Pu above the 1.8 MeV inverse β -decay threshold.

7.2 $\bar{\nu}_e$ Spectra

With the $f_{i,k}(t)$ in hand, the next step is to obtain the $\frac{dN_{\bar{\nu}_e,k}(E_{\bar{\nu}_e})}{dE_{\bar{\nu}_e}}$. For ^{235}U , ^{239}Pu and ^{241}Pu , data are available on β -decay spectra measured at the ILL High Flux Reactor [85, 86]. Converting these β spectra into $\bar{\nu}_e$ spectra is non-trivial since it involves an inversion of the sum of many spectra in long decay chains with many branches. No such measurements have been made for ^{238}U , so we rely on theoretical calculations [87]. Comparisons between calculations and measurements of the other $\bar{\nu}_e$ spectra indicate that the uncertainty of the calculations is on the order of 10%. However, since ^{238}U makes up only 10% of the $\bar{\nu}_e$ signal, this amounts to only a 1% error on the total flux. The $\bar{\nu}_e$ spectra above the 1.8 MeV inverse β -decay threshold is plotted for each of the four heavy isotopes in Figure 7.3. The uncertainty on the $\bar{\nu}_e$ rate above the analysis energy threshold due to these spectra is estimated to be 2.5%.

7.3 $\bar{\nu}_e$ Flux and Baseline

The $\bar{\nu}_e$ flux at KamLAND weights the contribution from each reactor according to the inverse-square of the distance L_i from KamLAND:

$$\phi_{\bar{\nu}_e}(t) = \sum_i^{\text{reactors}} \phi_{\bar{\nu}_e,i}(t) = \sum_i^{\text{reactors}} \frac{\int_0^\infty I_i(E_{\bar{\nu}_e}, t) dE_{\bar{\nu}_e}}{4\pi L_i^2}. \quad (7.3)$$

The longitudes and latitudes of Japanese commercial reactor cores were provided by TEPCO; those of KamLAND were obtained from Kamioka Mining and Smelting Co., Ltd. The errors on the L_i are estimated to be less than 70 m. The time variation of $\phi_{\bar{\nu}_e}(t)$ at KamLAND is shown in Figure 7.4. The flux is also broken down into reactor groups. As can be seen in the figure, many reactors around Japan were powered down during 2003 for safety maintenance, resulting in a modulation of the total flux at KamLAND by about 50%. As reactors at different distances varied their power, the typical baseline distance $\bar{\nu}_e$'s traveled on their way to KamLAND also changed. The baseline variation for Japanese reactors, given by

$$\bar{L}_{Jp}(t) = \frac{\sum_i^{Jp \text{ reac}} L_i \phi_{\bar{\nu}_e,i}(t)}{\sum_i^{Jp \text{ reac}} \phi_{\bar{\nu}_e,i}(t)}, \quad (7.4)$$

is plotted in Figure 7.5. The sum is limited to reactors in Japan so that \bar{L}_{Jp} is closer to the peak of the full baseline distribution rather than its true average (see Figure 2.3). The position of the peak is a more physical quantity than the average baseline because the baseline distribution is essentially the “power spectrum” for neutrino oscillations, and it is the position of the lowest-frequency peak in a power spectrum that determines the fundamental mode of a periodic signal.

7.4 Total Reactor $\bar{\nu}_e$ Signal

The remaining ingredients in the signal calculation (see Equation 2.3) are the resolution function $R(E_p, E'_{\bar{\nu}_e})$ described by Equations 2.2 and 4.10, the number of target protons n_p calculated in Section 5.6, the cross-section $\sigma(E_{\bar{\nu}_e})$ taken from [52], the $\bar{\nu}_e$




Figure 7.4: $\bar{\nu}_e$ flux at KamLAND, broken down into different reactor groups.




Figure 7.5: Variation of the effective baseline $\bar{L}_{Jp}(t)$, defined by Equation 7.4

Table 7.1: Systematic Uncertainties

| | |
|------------------------------------|-------|
| Number of target protons | 6.2% |
| Detection efficiency | 0.5% |
| Energy threshold | 2.3% |
| Livetime | 0.05% |
| Fission yields | 2.3% |
| $\bar{\nu}_e$ spectra [85, 86, 87] | 2.5% |
| Cross-section [52] | 0.2% |
| Total | 7.4% |

detection efficiency $\epsilon = 0.933 \pm 0.004$ calculated in Section 5.9, and the oscillation probability, given by Equation 1.11. Assuming no oscillations ($\Delta m^2, \sin^2 2\theta = 0$), the total expected number of detected $\bar{\nu}_e$'s is given by the energy- and time-integral of Equation 2.3, yielding $N_{no\ osc} = 310. \pm 23$ events. The 7.4% systematic error is broken down in Table 7.1, and includes a 2.3% error in the energy integration (assuming no oscillation) due to the 2.4% uncertainty in the energy scale at 2.6 MeV. The no-oscillation expected energy spectrum (the time-integral of Equation 2.3) is plotted in Figure 7.6.




Figure 7.6: Expected reactor $\bar{\nu}_e$ prompt energy spectrum in the absence of neutrino oscillation, given by the time-integral of Equation 2.3. The total expected event count above 2.6 MeV is $310. \pm 23$ events.

Chapter 8

Analysis

The data reported here were collected between March 2002 and December 2003, with a total exposure of 599 ton-years. Applying the event selection criteria described in Chapter 5 leaves 226 events. Analyses of the total event rate, the time variation, and the energy distribution are presented below.

8.1 $\bar{\nu}_e$ Disappearance and Average Survival Probability

One analysis that can be performed with the KamLAND data is a test of $\bar{\nu}_e$ disappearance comparing the full observed event count, $N_{obs} = 226$ events, and its expected background contribution from Chapter 6, $N_{bg} = 13.2 \pm 5.9$ events, to the expected event count in the absence of neutrino oscillations, calculated in Chapter 7 to be $N_{no\ osc} = 310. \pm 23$ events. Assuming gaussian systematic and poisson statistical uncertainties, the deficit of N_{obs} relative to $N_{no\ osc} + N_{bg}$ is evidence for $\bar{\nu}_e$ disappearance with 99.96% statistical significance. KamLAND is the first reactor antineutrino experiment to detect such disappearance. The average $\bar{\nu}_e$ survival probability is given by the ratio

$$\frac{N_{obs} - N_{bg}}{N_{no\ osc}} = 0.687 \pm 0.049(stat) \pm 0.055(sys) \quad (8.1)$$

The statistical error is on the same order as systematics. When interpreting the survival probability in terms of neutrino oscillation, it should be kept in mind that baseline variation depicted in Figure 7.5 implies a time-dependent survival probability. The number reported above is the flux-weighted average survival probability over the data taking period.

8.2 Time Variation

The time variation of the $\bar{\nu}_e$ rate shown in Figure 7.4 implies that the detected event rate should vary in time with the same pattern if the events detected are indeed $\bar{\nu}_e$'s produced by nuclear reactors. To extract this correlation, the data were binned according to data taking periods of similar expected event rate, and the detected event rate was measured for each bin. The correlation is plotted in Figure 8.1, in which the horizontal position of each point is determined by the time-weighted average of the expected rate in the bin. The vertical error bars express the statistical uncertainty on the measured rate.

For the case of no oscillation, the measured event rate should vary linearly with the expected event rate with slope 1. For finite Δm^2 and $\sin^2 2\theta$, the expectation value in each bin is the expected event rate times the average survival probability for the data taking periods included in the bin. While the average survival probability is different from bin to bin, for the values of Δm^2 and $\sin^2 2\theta$ near those determined in Section 8.3 below, that variation is slight. As a result, the trend is well fit by a first-order polynomial, with slope approximately equal to the average survival probability and y-intercept equal to the rate of background events.

A linear fit to the data is shown as the dashed line in Figure 8.1. The fit has a χ^2 of 0.56 for 6 degrees of freedom, and the blue curves depict the 90% CL. The solid black line is a linear fit with the y-intercept bound to the measured background rates, with $\chi^2 = 3.8$ for 5 degrees of freedom. The slope of 0.71 is in good agreement with the average $\bar{\nu}_e$ survival probability reported in Section 8.1. Due to KamLAND's low event rate, even a flat line with no slope (i.e. all background) fits the data reasonably well, with $\chi^2 = 2.0$ for 5 degree of freedom. However, the positive trend does indicate




Figure 8.1: Correlation of observed and expected $\bar{\nu}_e$ rates. The slope of a linear fit gives roughly the survival probability (see text), while the y -intercept gives the background level. The dashed line shows the best linear fit with $\chi^2/NDF = 0.56/4$; the 90% CL is drawn as the dotted curves. Constraining the fit to known background levels (box) gives a slope of 0.71 with $\chi^2/NDF = 3.8/5$.

that the event rate varies as expected.

As discussed in Section 6.1.2, a hypothetical geo-reactor at the Earth's core would produce $\bar{\nu}_e$'s with an energy spectrum indistinguishable from that of power reactor $\bar{\nu}_e$'s. Since the signal from this geo-reactor would remain constant in time, this time variation analysis could in principle be used to test the existence of the geo-reactor. However, a few-TW_{th} geo-reactor would produce a background rate of only a few events per year, roughly the same level as other backgrounds in KamLAND (see the box in Figure 8.1). Again, due to KamLAND's low event rate, this time variation analysis is not sensitive enough to place a meaningful upper limit on the geo-reactor. A more sophisticated analysis is underway. For the analysis described in Section 8.3, the geo-reactor is assumed to not exist.

8.3 Spectral Analysis

KamLAND's true power lies in its ability to measure the energy spectrum of the incident $\bar{\nu}_e$'s. It is able to detect any spectral distortions that might appear due to the energy dependence of the neutrino oscillation probability. In order to use the spectrum to test the hypothesis of neutrino oscillation, we start with the unbinned likelihood function for the data,

$$L(E_{a1}, E_{a2} \dots E_{aN}, \hat{\mathbf{p}}|\Theta) = P(N|\Theta) \prod_{i=1}^N \frac{dP(E_{ai}|\Theta)}{dE_a} P(\hat{\mathbf{p}}|\Theta). \quad (8.2)$$

Here the E_{ai} are the measured event energies in the analysis energy scale, and i runs from 0 to $N = 226$ observed events. The vector $\hat{\mathbf{p}}$ represents various parameters measured from external studies, such as background rates and energy scale parameters, with "true" values \mathbf{p} and error matrix $\tilde{\sigma}_{\mathbf{p}}^2$. Θ represents the underlying parameters determining the probability distribution function for the E_{ai} and $\hat{\mathbf{p}}$, namely Δm^2 , $\sin^2 2\theta$, \mathbf{p} , and any other parameters \mathbf{q} for which no external information is available. $P(N|\Theta)$ is the poisson probability of getting N events considering that N_0 were

expected:

$$P(N|\Theta) = \frac{N_0(\Theta)^N e^{-N_0(\Theta)}}{N!}. \quad (8.3)$$

The value of N_0 is calculated for a given value of Θ from the time- and energy-integral of Equation 2.4 over the data taking period and the analysis energy range, respectively. Note that the energy integral depends on the transformation function from analysis energy to real energy. The $\frac{dP(E_{ai}|\Theta)}{dE_a}$ are given by the normalized time-integral of Equation 2.4 converted to the analysis energy scale,

$$\frac{dP(E_{ai}|\Theta)}{dE_a} = \frac{1}{N_0(\Theta)} \frac{dN_0(E_{ai}|\Theta)}{dE_a}, \quad (8.4)$$

The third term in the likelihood function, $P(\hat{\mathbf{p}}|\Theta)$, is determined by the external measurements:

$$P(\hat{\mathbf{p}}|\Theta) = \frac{1}{\sqrt{2\pi|\tilde{\sigma}_{\mathbf{p}}^2|}} e^{-\frac{1}{2}(\hat{\mathbf{p}}-\mathbf{p})^T(\tilde{\sigma}_{\mathbf{p}}^2)^{-1}(\hat{\mathbf{p}}-\mathbf{p})}. \quad (8.5)$$

With the likelihood function in hand, we simply maximize it under the variation of Θ .

For computational purposes it is much more practical to maximize the logarithm of the likelihood function, $\log L$. The log-likelihood function avoids round-off errors that could otherwise be encountered with the extremely small likelihoods yielded by the product of probabilities expressed by Equation 8.2. Additionally, confidence intervals are easily calculated since excursions from the maximum value $\log L_{best}$, defined as $\Delta\chi^2 = -2(\log L - \log L_{best})$, obey a χ^2 distribution with number of degrees of freedom corresponding to the number of unconstrained parameters in the fit. The logarithm of the first two terms in Equation 8.2 simplifies greatly:

$$\begin{aligned} \log \left(P(N|\Theta) \prod_{i=1}^N \frac{dP(E_{ai}|\Theta)}{dE_a} \right) &= N \log N_0(\Theta) - N_0(\Theta) + C + \sum_{i=1}^N \log \frac{dP(E_{ai}|\Theta)}{dE_a} \\ &= -N_0(\Theta) + C + \sum_{i=1}^N \log \frac{dN_0(E_{ai}|\Theta)}{dE_a}, \end{aligned} \quad (8.6)$$

where C is a constant (i.e. it doesn't depend on Θ). Defining

$$\chi_{\mathbf{p}}^2 \equiv (\hat{\mathbf{p}} - \mathbf{p})^T (\hat{\sigma}_{\mathbf{p}}^2)^{-1} (\hat{\mathbf{p}} - \mathbf{p}), \quad (8.7)$$

the function to maximize becomes

$$\log L = \sum_{i=1}^N \log \frac{dN_0(E_{ai}|\Theta)}{dE_a} - N_0(\Theta) - \frac{1}{2} \chi_{\mathbf{p}}^2 + C', \quad (8.8)$$

where all constant terms have been absorbed into C' .

The conversion, when necessary, of the components of $\frac{dN_0(E_{\bar{\nu}_e})}{dE_{\bar{\nu}_e}}$ into analysis energy scale is performed using the chain rule. For example, for the reactor $\bar{\nu}_e$'s the conversion is

$$\frac{dN_{\bar{\nu}_e}(E_a)}{dE_a} = \frac{dN_{\bar{\nu}_e}(E_{\bar{\nu}_e}(E_a))}{dE_{\bar{\nu}_e}} \frac{dE_{\bar{\nu}_e}}{dE_a} = \frac{dN_{\bar{\nu}_e}(E_p(E_a) + E_{\Delta})}{dE_p} \frac{dE_p}{dE_a}, \quad (8.9)$$

where Equation 2.2 has been used, and $E_p(E_a)$ refers to the transformation function discussed at the end of Section 4.4.4. The integration limits that determine the value of N_0 use the same transformation function. The energy scale transformation uncertainty would be properly handled by allowing the parameters Θ_E in Equation 4.15 to vary, weighting them in $\chi_{\mathbf{p}}^2$ by their covariance matrix $\hat{\sigma}_E^2$. Since this is technically difficult to implement, the energy scale error is instead approximated as a linear departure from the central energy scale transformation function evaluated at the best-fit value of Θ_E , denoted $\hat{\Theta}_E$. Defining $\hat{E}_p(E_a) \equiv E_p(E_a|\Theta_E = \hat{\Theta}_E)$, we approximate

$$E_p(E_a, \Theta_E) \approx E_0 + \alpha_E \hat{E}_p(E_a), \quad \frac{dE_p}{dE_a} \approx \alpha_E \frac{d\hat{E}_p(E_a)}{dE_a}. \quad (8.10)$$

The values of E_0 and α_E are constrained in $\chi_{\mathbf{p}}^2$ by calculating the χ^2 of the deviation from the central energy scale transformation relative to the overall energy scale uncertainty at three sampled energies, $E_s = 1.022$ MeV, 2 MeV, and 3 MeV. The choice of three points gives the contribution to $\chi_{\mathbf{p}}^2$ from the energy rescaling roughly the same weight as if the three dominant parameters a , k_0 , and k_c in Θ_E were being

varied (recall that wide variations in k_B have little effect on the energy scale transformation function). The particular choice of energies between 1 and 3 MeV is intended to restrict the energy rescaling in the region over which the detector is most precisely calibrated.

The specific spectra appearing in $\frac{dN_0(E_{\alpha}|\Theta)}{dE_{\alpha}}$ and N_0 are the reactor $\bar{\nu}_e$ spectrum (Equation 2.3), the accidental background spectrum (Figure 6.2), the ${}^9\text{Li}/{}^8\text{He}$ spectrum (the solid line in Figure 6.8), and the ${}^{13}\text{C}(\alpha, n){}^{16}\text{O}$ spectrum (Figure 6.6). The handling of each is discussed below. The rates of other backgrounds in the analysis energy range are assumed to be negligible and are ignored.

The reactor $\bar{\nu}_e$ spectrum is given a variable normalization α_{reac} , which is constrained by a term in χ_p^2 to remain within its systematic error σ_{reac} of 1. The contributions to σ_{reac} are the same as those broken down in Table 7.1, except that the energy scale uncertainty at the analysis threshold is not included. It is naturally accounted for by the energy rescaling described above. Uncertainties in the energy dependence of the resolution function, the inverse β -decay cross-section, and the raw heavy isotope spectra $\frac{dN_{\bar{\nu}_e, k}(E_{\bar{\nu}_e})}{dE_{\bar{\nu}_e}}$ appearing in Equation 7.2 may additionally distort the spectral shape. However, these distortions are assumed to be insignificant compared to the statistical uncertainties in the measured spectral shape and the energy scale uncertainty, and are ignored in this analysis.

The accidental spectrum is the best-determined spectrum appearing in Equation 2.4: not only is it the background measured with the highest precision ($\sigma_{\text{acc}} = 1.3\%$), but its natural units are analysis energy, and therefore the conversion between analysis and real energy represented by Equation 8.9 is not necessary. A variable normalization parameter for the accidental spectrum, α_{acc} , is included in $\log L$, and it is constrained in χ_p^2 to lie within σ_{acc} of 1.

The ${}^9\text{Li}/{}^8\text{He}$ contribution, like the accidental spectrum, is given a variable normalization α_{LH} constrained within the $\sigma_{LH} = 74\%$ uncertainty on its measured rate. Unlike the accidental spectrum, energy scale uncertainties are applied because the shape of the ${}^9\text{Li}/{}^8\text{He}$ spectrum is obtained from a calculation of the ${}^9\text{Li}$ β^- -decay spectrum. Uncertainties in the shape of the ${}^9\text{Li}/{}^8\text{He}$ spectrum are ignored relative to the statistical and energy scale uncertainties.

The $^{13}\text{C}(\alpha, n)^{16}\text{O}$ spectrum is divided into three sub-spectra, each of which is handled differently in the analysis. The three sub-spectra are the low-energy hump from the thermalization of the neutron alone, the peak near 4.5 MeV consisting of an additional $^{12}\text{C}^*$ excitation γ , and the peak just above 6 MeV from γ 's emitted by the $^{13}\text{C}(\alpha, n)^{16}\text{O}^*$ states. The energy scale uncertainty is applied to all three spectra, while for the low-energy region an additional energy rescaling factor, α_{pq} , is included to account for uncertainty in the proton quenching factor. This rescaling factor is constrained to lie within $\sigma_{pq} = 10\%$ of 1.0 in χ^2_{p} . The normalization of the low-energy hump and the 4.5 MeV peak are tied to each other but are allowed to vary via the parameter $\alpha_{\alpha n_{gs}}$ by the uncertainty $\sigma_{\alpha n_{gs}} = 32\%$ in the $^{13}\text{C}(\alpha, n)^{16}\text{O}_{gs}$ rate. The normalization of the peak above 6 MeV is allowed to vary freely due to the large uncertainty in the branching ratio for $^{13}\text{C}(\alpha, n)^{16}\text{O}^*$. Its normalization parameter $\alpha_{\alpha n^*}$ is the only parameter comprising \mathbf{q} in the likelihood function.

Putting all of these details together, the factors in $\log L$ become

$$\begin{aligned}
\frac{dN_0(E_{ai}|\Theta)}{dE_a} = & \alpha_{\text{reac}}\alpha_E \frac{dN_{\bar{\nu}_e}(E_0 + \alpha_E \hat{E}_p(E_{ai})|\Delta m^2, \sin^2 2\theta) d\hat{E}_p(E_a)}{dE_p dE_a} + \\
& \alpha_{\text{acc}} \frac{dN_{\text{acc}}(E_{ai})}{dE_a} + \alpha_{LH}\alpha_E \frac{dN_{LH}(E_0 + \alpha_E \hat{E}_p(E_{ai})) d\hat{E}_p(E_a)}{dE_p dE_a} + \\
& \alpha_{\alpha n_{gs}}\alpha_E \frac{dN_{\alpha n \text{ low}}(\alpha_{pq}(E_0 + \alpha_E \hat{E}_p(E_{ai}))) d\hat{E}_p(E_a)}{dE_p dE_a} + \\
& \alpha_{\alpha n_{gs}}\alpha_E \frac{dN_{\alpha n \text{ mid}}(E_0 + \alpha_E \hat{E}_p(E_{ai})) d\hat{E}_p(E_a)}{dE_p dE_a} + \\
& \alpha_{\alpha n^*}\alpha_E \frac{dN_{\alpha n^*}(E_0 + \alpha_E \hat{E}_p(E_{ai})) d\hat{E}_p(E_a)}{dE_p dE_a}
\end{aligned} \tag{8.11}$$

$$\begin{aligned}
N_0(\Theta) = & \alpha_{\text{reac}}N_{\bar{\nu}_e}(\Delta m^2, \sin^2 2\theta, E_0, \alpha_E) + \alpha_{\text{acc}}N_{\text{acc}} + \alpha_{LH}N_{LH}(E_0, \alpha_E) + \\
& \alpha_{\alpha n_{gs}}N_{\alpha n \text{ low}}(E_0, \alpha_E, \alpha_{pq}) + \alpha_{\alpha n_{gs}}N_{\alpha n \text{ mid}}(E_0, \alpha_E) + \\
& \alpha_{\alpha n^*}N_{\alpha n^*}(E_0, \alpha_E)
\end{aligned} \tag{8.12}$$




Figure 8.2: Measured prompt energy spectrum with best-fit to neutrino oscillation. The shaded regions are the best-fit accidental and $^{13}\text{C}(\alpha, n)^{16}\text{O}$ background contributions. The dashed line marks the 2.6 MeV analysis energy threshold.

$$\chi_{\text{p}}^2 = \frac{(\alpha_{\text{reac}} - 1)^2}{\sigma_{\text{reac}}^2} + \frac{(\alpha_{\text{acc}} - 1)^2}{\sigma_{\text{acc}}^2} + \frac{(\alpha_{\text{LH}} - 1)^2}{\sigma_{\text{LH}}^2} + \frac{(\alpha_{\alpha n_{gs}} - 1)^2}{\sigma_{\alpha n_{gs}}^2} + \frac{(\alpha_{pq} - 1)^2}{\sigma_{\text{reac}}^2} + \sum_1^3 \frac{(E_0 + E_{s_i}(\alpha_E - 1))^2}{\sigma_{s_i}^2} \quad (8.13)$$

The value of $\log L$ was maximized under the variation of the 10 parameters Δm^2 , $\sin^2 2\theta$, E_0 , α_E , α_{reac} , α_{acc} , α_{LH} , $\alpha_{\alpha n_{gs}}$, α_{pq} , and $\alpha_{\alpha n^*}$. The maximum was found at $\Delta m^2 = 8.05 \times 10^{-5} \text{ eV}^2$, $\sin^2 2\theta = 0.68$. Each constrained parameter fit within its uncertainty of its central value. The normalization of the $^{13}\text{C}(\alpha, n)^{16}\text{O}^*$ spectrum (determined by the free parameter $\alpha_{\alpha n^*}$) fit to a value of 2.3 events. The best-fit prompt energy spectrum is plotted together with the histogrammed data in Figure 8.2. The best-fit accidental and $^{13}\text{C}(\alpha, n)^{16}\text{O}$ background contributions are also shown.

To determine the confidence intervals for Δm^2 and $\sin^2 2\theta$, the value of $\Delta\chi^2 = -2(\log L(\Delta m^2, \sin^2 2\theta) - \log L_{\text{best}})$ was obtained for every point $(\Delta m_i^2, \sin^2 2\theta_j)$ on a grid in Δm^2 - $\sin^2 2\theta$ parameter space, where $\log L_{\text{best}}$ is the value of $\log L(\Delta m^2, \sin^2 2\theta)$ at the best-fit point. $\log L(\Delta m_i^2, \sin^2 2\theta_j)$ was computed by fixing $\Delta m^2 = \Delta m_i^2$ and




Figure 8.3: Confidence levels for a fit of neutrino oscillation to the KamLAND data (shaded regions) and solar neutrino experiments (lines) [26]. The joint- 1σ intervals within LMA I are $\Delta m^2 = 8.05_{-0.72}^{+0.95} \times 10^{-5} \text{ eV}^2$, $\sin^2 2\theta > 0.46$.

$\sin^2 2\theta = \sin^2 2\theta_j$ while the remaining 8 parameters were varied to maximize $\log L$. Contours in $\Delta\chi^2$ are drawn in Figure 8.3 corresponding to the 95% ($\Delta\chi^2 = 5.99$), 99% ($\Delta\chi^2 = 9.21$), and 99.73% ($\Delta\chi^2 = 11.83$) confidence levels (CL's). There are three inclusion regions, all in the proximity of the LMA MSW solution to the solar neutrino problem (drawn in solid lines in the figure). The two at larger values of Δm^2 have been named “LMA I”, and “LMA II”, in order of increasing Δm^2 [88]. The one at lower Δm^2 will be referred to as “LMA 0”. The best-fit point lies in LMA I, and the joint- 1σ intervals for the two oscillation parameters in this region are $\Delta m^2 = 8.05_{-0.72}^{+0.95} \times 10^{-5} \text{ eV}^2$, $\sin^2 2\theta = 0.68_{-0.22}^{+0.20}$. The minima of LMA 0 and LMA II lie at the 44.0% and 91.6% CL's, respectively. The minimum along the line at maximal mixing ($\sin^2 2\theta = 1.0$) lies at the 51.8% CL.

As is well known, while the unbinned likelihood function provides a powerful way to determine the best-fit parameters and their confidence intervals for some theory, the value of $\log L$ at its maximum cannot be used to test the goodness of the fit to the data [89]. To test the fit, we follow the procedures outlined in [89] and bin the data into $n = 20$ energy bins whose widths are chosen such that the probability, according to the best-fit to neutrino oscillation, of an event landing in any particular bin is a constant. The goodness-of-fit statistic is chosen to be the Pearson- χ^2 (χ_P^2) of the data with respect to the theory in this binning:

$$\chi_P^2 = \sum_{i=1}^n \frac{(N_i - N_{0i})^2}{N_{0i}} = 25.3, \quad (8.14)$$

where N_i and N_{0i} are the observed and expected event counts in bin i , respectively. Naively one would expect χ_P^2 to vary as χ^2 with $n - [2 + \dim(\mathbf{q})] = 17$ degrees of freedom, since Δm^2 , $\sin^2 2\theta$, and $\mathbf{q} = \{\alpha_{\alpha n^*}\}$ are varied freely in the maximization of $\log L$. However, this is not strictly the case, partly because $\log L$ is not binned like χ_P^2 , and also because the χ_P^2 terms in $\log L$ do not appear in χ_P^2 . In case χ_P^2 for the fitting procedure used in this analysis is slightly different from the standard χ^2 distribution, 10000 Monte Carlo (MC) energy spectra are simulated from the best-fit spectrum. As is done for the data, $\log L$ is maximized to obtain a best-fit to neutrino oscillation for each MC spectrum. The best-fit oscillation parameters for the MC spectra are plotted in Figure 8.4. Then χ_P^2 is calculated for each MC spectrum relative to its own best-fit. The distribution of χ_P^2 for the MC spectra is histogrammed in Figure 8.5. Although the distribution nicely follows the χ^2 distribution with 17 degrees of freedom, drawn as the solid curve, its mean is slightly different from the expected 17. Of the MC spectra, 907 have $\chi_P^2 > 25.3$, giving a goodness-of-fit of 9.1% for neutrino oscillation. While this suggests that the fit might not be particularly good, it does not indicate that anything is seriously wrong with the assumption that the energy spectrum can be described by neutrino oscillation. It should also be kept in mind that statistics derived from binned data are inherently sensitive to the choice of binning. The choice of 20 bins was made by approximating Equation 30.72 in [89]. As illustrated in Figure 8.6, the goodness of fit varies widely for $5 < N_{bins} < 35$; the a priori choice




Figure 8.4: Best-fit oscillation parameters to MC spectra.

$N_{bins} = 20$ appears to be somewhat “unlucky” in this respect.

8.4 Spectral Distortion

Exploring further the significance of the distortion of the energy spectrum measured at KamLAND, we plot in Figure 8.7 the histogrammed data along with the no-oscillation spectrum (light grey line). The disagreement between overall amplitudes of the data and the no-oscillation hypothesis is readily apparent. Next we allow the no-oscillation spectrum to rescale itself to fit the data better by setting $\Delta m^2 = 0$, $\sin^2 2\theta = 0$ and removing the constraint on α_{reac} from χ_p^2 in $\log L$. This rescaled no-oscillation spectrum is plotted as the dark line in Figure 8.7. The best-fit rescaling parameter was 0.825 ± 0.046 . Even by eye it appears to still be a poor fit to the data, particularly in the first two bins above threshold. The procedure described in Section 8.3 to obtain the goodness-of-fit for neutrino oscillation was used to test the rescaled no-oscillation fit. The data had $\chi_p^2 = 32.7$, less than only 1.8% of the MC spectra. Taking the scaled no-oscillation theory as a null-hypothesis, its goodness-of-fit may be interpreted as establishing the significance of spectral distortion at 98.2%




Figure 8.5: χ^2_P distribution of MC spectra. A χ^2 distribution with 17 degrees of freedom is normalized to the number of MC spectra and drawn for comparison.




Figure 8.6: Variation of goodness of fit with choice of number of bins (N_{bins}).




Figure 8.7: Best-fit of a scaled no-oscillation hypothesis to the KamLAND data (dark line). The fit is poor, particularly in the first two bins. The unscaled no-oscillation reactor $\bar{\nu}_e$ spectrum is drawn as the grey line.

significance.

Keeping the constraint on α_{react} out of χ^2 , the values of Δm^2 and $\sin^2 2\theta$ were once again allowed to vary, yielding a “shape-only” fit to neutrino oscillation. The $\Delta\chi^2$ contours are drawn for this fit in Figure 8.8. These contours show that KamLAND’s sensitivity in Δm^2 is contributable entirely to the shape of the spectrum. The joint- 1σ error bars on Δm^2 within the LMA I island are $\Delta m^2 = (8.14^{+0.83}_{-0.73}) \times 10^{-5} \text{ eV}^2$. The shape-only fit weakly prefers maximal mixing in LMA 0, with the minimum in LMA I lying at the 59.1% CL. It is the additional constraint of the rate that pulls Δm^2 up into LMA I and $\sin^2 2\theta$ to smaller values.

While much emphasis has been placed on the apparent spectral distortion, it is difficult to simply look at the KamLAND spectrum and claim that its difference from the no-oscillation spectrum looks particularly “oscillation-like”. This difficulty is rooted in the fact that the oscillation probability, Equation 1.11, is really a function of $L/E_{\bar{\nu}_e}$ rather than simply $E_{\bar{\nu}_e}$. So, as a visual tool, the event distribution in $L_{eff}/E_{\bar{\nu}_e}$ is histogrammed in Figure 8.9 as a fraction of the no-oscillation expectation.




Figure 8.8: Confidence levels for a shape-only fit of neutrino oscillation to the KamLAND data. The spectral shape weakly prefers maximal mixing in LMA 0; it is the added constraint of the total rate that pushes the best-fit point into LMA I. LMA 0 is inconsistent with solar neutrino experiments.

$E_{\bar{\nu}_e}$ is the neutrino energy, calculated from E_p using Equation 2.2. The “effective” baseline L_{eff} is a constant, acting as a rescaling, and is taken to be 174 km, given by the time-average of $\bar{L}_{Jp}(t)$ defined by Equation 7.4. The shaded regions correspond to energies outside of the prompt energy analysis range. The blue histogram corresponds to the best-fit neutrino oscillation expectation in this variable; the data nicely follow its sinusoidal pattern. To highlight this oscillatory pattern, the sinusoidal survival probability versus $L_{eff}/E_{\bar{\nu}_e}$ for a single $\bar{\nu}_e$ source at baseline L_{eff} is sketched as the dotted curve. The difference between the best-fit oscillation pattern and this “ideal” oscillation is due primarily to the spread in the reactor baseline distribution, Figure 2.3, which causes the amplitude of the oscillation pattern to wash out with increasing distance or lower energy. If Δm^2 were much larger or smaller than the value measured by KamLAND, no spectral distortion would be present, and the limits on the parameter would be much weaker.




Figure 8.9: Ratio of detected event rate to the no-oscillation expectation versus $L_{\text{eff}}/E_{\bar{\nu}_e}$. The best-fit oscillation curve is drawn as the solid line. The dotted line shows the ideal oscillation pattern for a single reactor at distance L_{eff} .

8.5 Solar + KamLAND Analysis

Assuming CPT invariance, which implies that $P_{\bar{\nu}_e \rightarrow \bar{\nu}_e} = P_{\nu_e \rightarrow \nu_e}$, the KamLAND result may be used to further constrain the LMA MSW solution to the solar neutrino problem. A two-flavor analysis was performed combining the KamLAND reactor $\bar{\nu}_e$ spectral measurement with the observed solar neutrino fluxes assuming the Standard Solar Model [22, 90]. The systematic uncertainties in these two sets of experiments were assumed to be completely uncorrelated. A fit to Δm^2 and $\tan^2 \theta$ was performed by adding a penalty term to KamLAND's log-likelihood function, $\log L_{KL}$ defined by Equation 8.8, for the solar data:

$$\log L_{\text{tot}} = \log L_{KL} + \log L_{\text{sol}}. \quad (8.15)$$

The contributions to $\log L_{\text{sol}}$ are seven coupled equations corresponding to the gaussian probabilities of each solar neutrino experiment to observe its reported flux given a




Figure 8.10: Confidence levels for a combined fit of neutrino oscillation to solar [90] + KamLAND data. The best-fit point is at $\Delta m^2 = 7.86_{-0.59}^{+0.75} \times 10^{-5} \text{ eV}^2$, $\tan^2 \theta = 0.391_{-0.034}^{+0.036}$.

particular set of oscillation parameters, with a constraint for the solar irradiance. The full likelihood was minimized as before on a grid of points in parameter space, giving the contours shown in Figure 8.10. The joint- 1σ errors on the oscillation parameters for the combined analysis are $\Delta m^2 = 7.86_{-0.59}^{+0.75} \times 10^{-5} \text{ eV}^2$, $\tan^2 \theta = 0.391_{-0.034}^{+0.036}$.

Chapter 9

Conclusions

9.1 Results and Implications

As presented in Chapter 8, KamLAND has observed the disappearance of reactor electron antineutrinos with 99.96% statistical significance. A time variation analysis shows weak but positive correlation between the detected signal and reactor power variations. Interpreted in terms of two-flavor neutrino oscillation, the KamLAND spectrum restricts the mixing parameters Δm^2 and $\sin^2 2\theta$ to 3 islands, named LMA 0, I, and II in order of increasing Δm^2 . The best fit is in LMA I and is consistent with the LMA MSW solution to the solar neutrino problem. The joint- 1σ confidence intervals for the mixing parameters within LMA I are $\Delta m^2 = 8.05^{+0.95}_{-0.72} \times 10^{-5} \text{ eV}^2$ and $\sin^2 2\theta = 0.68^{+0.20}_{-0.22}$.

While the goodness of fit for the hypothesis of neutrino oscillation is only 9.1%, there is no indication of fundamental problems with the fit. On the other hand, an un-oscillated spectral shape rescaled to fit the data is a very poor fit, giving evidence for spectral distortion at 98.2% significance. A shape-only fit to the KamLAND data reveals that the sensitivity in Δm^2 is determined essentially by this shape distortion. Hence the KamLAND result does not rely on the reactor power data, a precise determination the fiducial volume size, or any other quantity that affects the overall rate only. When the KamLAND data are plotted as a ratio of the expected signal, the

spectral distortion resembles very much an oscillation pattern, although the significance of the oscillation is not strong enough at this point to claim that the KamLAND data is evidence of such.

Assuming CPT invariance, KamLAND eliminates all neutrino oscillation solutions to the solar neutrino problem except LMA at $> 99.73\%$ CL. A combined two-flavor oscillation analysis of solar neutrino measurements and the KamLAND spectrum leaves only LMA I. The joint- 1σ confidence intervals for the mixing parameters deduced from the combined analysis are $\Delta m^2 = 7.86_{-0.59}^{+0.75} \times 10^{-5} \text{ eV}^2$ and $\tan^2 \theta = 0.391_{-0.034}^{+0.036}$. Essentially, the sensitivity on Δm^2 comes from the KamLAND spectrum, while the $\tan^2 \theta$ limits are determined by the solar neutrino fluxes.

KamLAND is the first and only experiment so far to have observed electron antineutrino disappearance using man-made sources, marking a step forward for the neutrino oscillation industry. Perhaps more significant, however, is the fact that the mixing parameters deduced by KamLAND agree so beautifully with solar neutrino results. KamLAND uses scintillation detection to investigate essentially vacuum oscillations of antineutrinos emitted by nuclear reactors; solar neutrino experiments use radio-chemical and water Cherenkov detection to investigate matter oscillations of neutrinos from the sun. These experiments could not be more different, yet the physics of neutrino oscillation unites them in one consistent theoretical framework.

9.2 KamLAND Prospects

In the present analysis, statistical uncertainties are on the same level as systematics. KamLAND continues to take data, and will likely double its present livetime. Efforts are underway to address several of the larger contributions to the systematic error.

A “ 4π ” calibration system capable of deploying γ sources to off-axis positions is in its final stages of development. It will be used to calibrate vertex reconstruction algorithms throughout the fiducial volume and hence obtain an error on r_{fid} . The target precision for the deployed source position of $< 5 \text{ cm}$ at $r = 5 \text{ m}$ could reduce the current 6.2% error on the fiducial volume to the 3% level or better.

Another project well underway is the installation of a muon tracking chamber in

the dome above the detector. With a secondary measurement of a sample of muon tracks crossing the detector, the muon track reconstruction software can be improved and its performance consistently evaluated. The muon tracking chamber can also be used to verify the OD efficiency and fast neutron background estimates.

As explained in Chapter 7, the 2% uncertainty due to the reactor power is obtained from regulatory safety margins rather than a detailed analysis of the power measurement itself. Preliminary work to reassess this error indicates that a smaller uncertainty may be reasonable. In particular, reactors which have installed the latest ultrasonic flow-measurement technology can measure the power to better than a percent. In addition, a significant fraction of this error may be uncorrelated between reactor sites, allowing for partial cancelations in the summation of the contributions from many reactors.

There are many studies presented in this analysis which would benefit greatly from a full Monte Carlo (MC) simulation of KamLAND. Difficulties in simulating the optical properties of the world's largest scintillator detector have confounded KamLAND MC efforts to date. However, a push is being made to develop a more useful and detailed detector simulation. A good MC could drastically improve detector understanding and reconstruction codes. In particular, it could be used to explore the effect of electronics dead time on reconstruction, or to better estimate the systematics of the fiducial volume ratio measurements with $^{12}\text{B}/^{12}\text{N}$ candidates and spallation neutrons. The MC work will hopefully lead to reductions of systematics in many areas of the experiment.

Most of the systematic uncertainties listed in Table 7.1 affect the rate only. However, due to the nature of the oscillation probability function (Equation 1.11), the rate systematics have little impact on the final Δm^2 sensitivity. As demonstrated by the shape-only analysis reported in Section 8.4, the Δm^2 sensitivity comes essentially from KamLAND's spectral shape. As statistics increase, the Δm^2 error bars may continue to shrink, unhindered by the rate systematics. This suggests a growing need to revisit the uncertainties in the spectral shape. In the present analysis, uncertainties in the $\bar{\nu}_e$ spectra (Section 7.2), errors in the resolution function parameters, and

non-linearities in the energy scale error were ignored as insignificant relative to statistical errors. The largest of these is the $\bar{\nu}_e$ spectral uncertainties, which are currently treated as an error on their normalization only. But recall that they are derived from inversions of β -spectra measurements, so errors in different energy regions are not independent. It may be necessary to parameterize these uncertainties and incorporate their variation into the likelihood function.

In addition to the reassessment of systematic errors, KamLAND is preparing for a second phase aimed at measuring the ${}^7\text{Be}$ peak in the solar neutrino spectrum, see Figure 1.1. These monoenergetic neutrinos scatter elastically off of electrons in the scintillator, giving a spectrum of single events characterized by a sharp edge at the maximum e^- recoil energy of 665 keV. In order to observe these neutrinos, KamLAND would need to decrease backgrounds in the sub-MeV range by some 6 orders of magnitude. Most of these backgrounds are due to ${}^{85}\text{Kr}$ and ${}^{210}\text{Pb}$, with some contributions also from ${}^{40}\text{K}$. Scintillator re-purification research and development teams are devising ways to efficiently remove these contaminants from the scintillator. If they are successful, KamLAND may be capable of detecting an expected few hundred ${}^7\text{Be}$ neutrinos per day and constrain further the Standard Solar Model. Reactor neutrino studies would continue during the solar neutrino detection phase.

Meanwhile, a wealth of other physics topics is being explored with KamLAND. Currently under investigation are geo-neutrinos in the energy range 0.9 to 2.6 MeV. A more sophisticated time variation analysis may lead to a meaningful upper limit on the power of a geo-reactor at the center of the Earth. A measurement of spallation neutron production rates at the KamLAND depth is also being prepared. And with at least a factor of 3 increase in accumulated livetime since the initial publication, the sensitivity to the solar $\bar{\nu}_e$ flux has improved and is being explored. Including future updates to the reactor $\bar{\nu}_e$ result and the possibility of solar neutrino detection, many exciting physics results are still to come from KamLAND.

Bibliography

- [1] J.N. Bahcall *et al.*, *Astrophys. J.* **555**, 990 (2001).
- [2] R. Davis *et al.*, *Phys. Rev. Lett.* **20** 1205 (1968).
- [3] J.N. Bahcall *et al.*, *Phys. Rev. Lett.* **20** 1209 (1968).
- [4] S. Eidelman *et al.*, *Phys. Lett. B* **592**, 1 (2004).
- [5] J.N. Bahcall and R. Davis, *Science* **191**, 264 (1976); J.N. Bahcall, *Neutrino Astrophysics* (Cambridge University Press, Cambridge, UK, 1989);
- [6] J.N. Bahcall, *Astrophys. J.* **467**, 475 (1996).
- [7] B. Pontecorvo, *Sov. Phys. JETP* **26**, 984 (1968).
- [8] B.T. Cleveland *et al.*, *Astrophys. J.* **496**, 505 (1998).
- [9] M. Gell-Mann *et al.*, in "Supergravity", Ed. by D. Freedman *et al.*, North Holland (1979); T. Yanagita, *Prog. Theor. Phys.* **64**, 1103 (1980).
- [10] B. Pontecorvo, *Sov. Phys. JETP* **6**, 429 (1958);
- [11] B. Pontecorvo, *Sov. Phys. JETP* **7**, 172 (1959); Z. Maki *et al.*, *Prog. Theor. Phys.* **28**, 870 (1962).
- [12] R.V. Wagoner *et al.*, *Astrophys. J.* **148**, 3 (1967);
- [13] R.H. Dicke *et al.*, *Astrophys. J.* **142**, 414 (1965); A.A. Penzias and R. Wilson, *Astrophys. J.* **142**, 419 (1965).

- [14] M. Fukugita and T. Yanagida, "Physics of Neutrinos and Applications to Astrophysics", Springer-Verlag, Berlin (2003).
- [15] C. Giunti, *Phys. Scripta* **67**, 29 (2003); M. Beuthe, *Phys. Rept.* **375**, 105 (2003).
- [16] L. Wolfenstein, *Phys. Rev. D* **17**, 2369 (1978).
- [17] S.P. Mikheev and A.Yu. Smirnov, *Sov. J. Nucl. Phys.* **42**, 913 (1985).
- [18] R.D. McKeown and P. Vogel, hep-ph/0402025.
- [19] W. Hampel *et al.*, *Phys. Lett. B* **447**, 127 (1999).
- [20] M. Altmann *et al.*, *Phys. Lett. B* **490**, 16 (2000).
- [21] J.N. Abdurashitov *et al.*, *JETP* **95**, 181 (2002).
- [22] J.N. Bahcall and C. Peña-Garay, *New J. Phys.* **6**, 63 (2004).
- [23] Updated from Figure 2 of [6] with numbers taken from [1] and [26], available from <http://www.sns.ias.edu/jnb> (2004).
- [24] Y. Fukuda *et al.*, *Phys. Rev. Lett.* **77**, 1683 (1996).
- [25] S. Fukuda *et al.*, *Phys. Lett. B* **539**, 179 (2002).
- [26] S.N. Ahmed *et al.*, *Phys. Rev. Lett.* **92**, 181301 (2004).
- [27] K.S. Hirata *et al.*, *Phys. Rev. Lett.* **63**, 16 (1989).
- [28] H. Murayama, <http://hitoshi.berkeley.edu/neutrino>; based on M.B. Smy, *Nucl. Phys. B Proc. Suppl.* **118**, 25 (2003).
- [29] K.S. Hirata *et al.*, *Phys. Lett. B* **280**, 146 (1992).
- [30] R.A. Becker-Szendy *et al.*, *Phys. Rev. D* **46**, 3720 (1992).
- [31] T.J. Haines *et al.*, *Phys. Rev. Lett.* **57**, 1986 (1986).
- [32] W.W.M. Allison *et al.*, *Phys. Lett. B* **391**, 491 (1997).

- [33] Y. Fukuda *et al.*, Phys. Rev. Lett. **81**, 1562 (1998).
- [34] Y. Ashie *et al.*, Phys. Rev. Lett. **93**, 101801 (2004).
- [35] C. Bemporad *et al.*, Rev. Mod. Phys. **74**, 297 (2002).
- [36] H. Kwon *et al.*, Phys. Rev. D **24**, 1097 (1981).
- [37] G. Zacek *et al.*, Phys. Rev. D **34**, 2621 (1986).
- [38] A.I. Afonin *et al.*, Sov. Phys. JETP **67**, 213 (1988).
- [39] G.S. Vidyakin *et al.*, JETP Lett. **59**, 390 (1994).
- [40] B. Achkar *et al.*, Nucl. Phys. B **434**, 503 (1995).
- [41] Z.D. Greenwood *et al.*, Phys. Rev. D **53**, 6054 (1996).
- [42] B. Achkar *et al.*, Phys. Lett. B **374**, 243 (1996).
- [43] F. Boehm *et al.*, Phys. Rev. D **64**, 112001 (2001).
- [44] L. Miller, Ph.D. thesis, Stanford University (2000).
- [45] M. Apollonio *et al.*, Eur. Phys. J. C **27**, 331 (2003).
- [46] A. Aguilar *et al.*, Phys. Rev. D **64**, 112007 (2001).
- [47] B. Armbruster *et al.*, Phys. Rev. D. **65**, 112001 (2002).
- [48] A.O. Bazarko, Nucl. Phys. B Proc. Suppl. **117**, 33 (2003).
- [49] D.N. Spergel *et al.*, Astrophys. J. Suppl. **148**, 175 (2003).
- [50] J. Busenitz *et al.*, "Proposal for US Participation in KamLAND" (1999).
- [51] International Nuclear Safety Center Web Page, <http://www.insc.anl.gov> (2004).
- [52] P. Vogel and J.F. Beacom, Phys. Rev. D **60**, 053003 (1999); radiative correction from A. Kurylov, *et al.*, Phys. Rev. C **67**, 035502 (2003).

- [53] K. Eguchi *et al.*, Phys. Rev. Lett. **92**, 071301 (2004); K. McKinney, Ph.D. thesis, University of Alabama (2003); H. Ogawa, Ph.D. thesis, Tohoku University (2003).
- [54] R.S. Raghavan *et al.*, Phys. Rev. Lett. **80**, 635 (1998).
- [55] K. Eguchi *et al.*, Phys. Rev. Lett. **90**, 021802 (2003); T. Iwamoto, Ph.D. thesis, Tohoku University (2003); O. Tajima, Ph.D. thesis, Tohoku University (2003).
- [56] T. Araki *et al.*, hep-ex/0406035.
- [57] F. Suekane *et al.*, physics/0404071.
- [58] J.B. Benziger *et al.*, Nucl. Instrum. Methods Phys. Res. A **417**, 278 (1998).
- [59] H. Kume *et al.*, Nucl. Instrum. Methods **205**, 443 (1986).
- [60] J.F. Beacom *et al.*, Phys. Rev. D **66**, 033001 (2002).
- [61] A. Savitzky and M.J.E. Golay, Anal. Chem. **36**, 1627 (1964).
- [62] W.H. Press *et al.*, "Numerical Recipes in C: The Art of Scientific Computing", Cambridge U. Press, Cambridge (1992).
- [63] J. Friedman, Proc. of the 1974 CERN School of Computing, Norway (1974).
- [64] J.B. Birks, "The Theory and Practice of Scintillation Counting", Pergamon, London (1964).
- [65] I. Kawrakow and D.W.O. Rogers, "The EGSnrc Code System: Monte Carlo Simulation of Electron and Photon Transport", unpublished. The EGSnrc package can be found at <http://www.irs.inms.nrc.ca/inms/irs/EGSnrc/EGSnrc.html>.
- [66] F. James and M. Roos, Comp. Phys. Comm. **10**, 343 (1975).
- [67] P. Antonioli *et al.*, Astropart. Phys. **7**, 357 (1997).
- [68] S. Agostinelli *et al.*, Nucl. Instrum. Methods Phys. Res. A **506**, 250 (2003).

- [69] NuDat2.0, National Nuclear Data Center, Brookhaven National Laboratory, available at <http://www.nndc.bnl.gov/nudat2> (2004).
- [70] T. Hagner *et al.*, *Astropart. Phys.* **14**, 33 (2000).
- [71] V.I. Kopeikin *et al.*, *Physics of Atomic Nuclei* **64**, 849 (2001).
- [72] J. Detwiler *et al.*, *Phys. Rev. Lett.* **89**, 191802 (2002).
- [73] J.M. Herndon, *Proc. Nat. Acad. Sci.* **100**, 3047 (2003).
- [74] F. Mantovani *et al.*, *Phys. Rev. D* **69**, 013001 (2004).
- [75] G.V. Domogatskii, *Sov. Astron.* **28** (1984).
- [76] C.E. Ortiz *et al.*, *Phys. Rev. Lett.* **85**, 2909 (2000).
- [77] T.K. Gaisser *et al.*, *Phys. Rev. D* **38**, 85 (1988).
- [78] K.K. Sekharan *et al.*, *Phys. Rev.* **156**, 1187 (1967); we actually used an R-matrix evaluation of this measured cross section from <http://wwwndc.tokai.jaeri.go.jp/ftpnd/jendl/jendl-an-2003.html> (2003).
- [79] G.W. Kerr *et al.*, *Nucl. Phys. A* **110**, 637 (1968); R.B. Walton *et al.*, *Phys. Rev.* **170**, 1065 (1957).
- [80] K. Hagiwara *et al.*, *Phys. Rev. D* **66**, 010001-220, Chapter 20 (2002).
- [81] J.A. Messimore, Ph.D. Thesis, North Carolina State University (2003).
- [82] T.R. England and B.F. Rider, ENDF-349, LA-UR-94-3106 (1994).
- [83] Japan Atomic Industrial Forum, Inc. Web Page, <http://www.jaif.or.jp> (2004); The Federation of Electric Power Companies of Japan Web Page, <http://www.fepc-atomic.jp> (2004).
- [84] Korea Hydro & Nuclear Power Company Web Page, <http://www.khnp.co.kr> (2004).

- [85] A.A. Hahn *et al.*, Phys. Lett. B **218**, 365 (1989).
- [86] K. Schreckenbach *et al.*, Phys. Lett. B **160**, 325 (1985).
- [87] P. Vogel *et al.*, Phys. Rev. C **24**, 1543 (1981).
- [88] G.L. Fogli *et al.*, Phys. Rev. D **67**, 073002 (2003).
- [89] A. Stuart *et al.*, "Kendall's Advanced Theory of Statistics", Vol 2A, Oxford University Press, New York (1999).
- [90] M. Maltoni *et al.*, New J. Phys. **6**, 122 (2004).