

Double Beta Decay in Xenon-136: Measuring the Neutrino-Emitting Mode and Searching for Majoron-Emitting Modes

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: Steven Herrin

DOUBLE BETA DECAY IN XENON-136:
MEASURING THE NEUTRINO-EMITTING MODE
AND SEARCHING FOR MAJORON-EMITTING MODES

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

Steven Herrin

June 2013

© 2013 by Steven M. Herrin. All Rights Reserved.
Re-distributed by Stanford University under license with the author.



This work is licensed under a Creative Commons Attribution-
Noncommercial 3.0 United States License.
<http://creativecommons.org/licenses/by-nc/3.0/us/>

This dissertation is online at: <http://purl.stanford.edu/ss855rt9766>

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.

Martin Breidenbach, Primary Adviser

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

Patricia Burchat

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

Aaron Roodman

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumpert, Vice Provost Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

Observations of neutrino flavor oscillations have demonstrated that neutrinos have mass. Since the discovery of these oscillations, much progress has been made at measuring the neutrino mass-squared differences and lepton mixing angles that characterize them. However, the origin and absolute scale of neutrino masses remain unknown. Unique among fermions, neutrinos can be Majorana particles, which could provide an explanation for neutrino masses. Discovery of a hypothetical process known as neutrinoless double beta decay would show that neutrinos are Majorana particles and determine the mass scale for neutrinos.

The Enriched Xenon Observatory (EXO) is a series of experiments searching for the neutrinoless double beta decay of ^{136}Xe . The first experiment, EXO-200, began operation in 2011 and makes use of 200 kg of xenon enriched to 80.6 % in ^{136}Xe . The analysis presented here makes use of data from EXO-200 to obtain a more precise measurement of the half-life for the two-neutrino-emitting mode of double beta decay than previously reported. The analysis also sets limits on the half-lives for exotic, Majoron-emitting modes of neutrinoless double beta decay. Data from EXO-200 is also used to produce a measurement of the cosmic muon flux at the WIPP underground site where EXO-200 is located.

Acknowledgements

I'd like to begin by thanking my adviser, Martin Breidenbach. From working with and watching Marty, I've learned a lot about physics, and even more about building and running a successful physics experiment. I'd also like to thank the rest of my committee, including Patricia Burchat, Giorgio Gratta, Aaron Roodman, and Boris Murmann for their time and feedback as I prepared and defended this dissertation.

My family, including my mother Doyn Kellerhals, my father Stanley Herrin, and my sister Janet Herrin were all so supportive of me throughout graduate school, and I thank them deeply.

My SLAC labmates, both past and present, have been great to work with. Al Odian, Charles Prescott, and Peter Rowson shared with me their experience and knowledge, and I have learned a lot from them. Working together with Nicole Ackerman, Derek Mackay, Ryan MacLellan, Jesse Wodin, and Liang Yang to actually build EXO-200 was a great adventure, especially when we finally got everything working.

The EXO collaboration is large, and while every member deserves acknowledgement, I only have enough space to thank a subset. Phil Barbeau and Simon Slutsky were both awesome to work with while building EXO-200 and then analyzing some of the first data. Tim Daniels, Michelle Dolinski, and Delia Tosi deserve praise for co-ordinating EXO-200 runs and being great to work with. So much work has gone into this analysis, and I only begin to scratch the surface by thanking David Auty, Joshua Albert, Sean Daugherty, Clayton Davis, Ralph DeVoe, David Moore, Tessa Johnson, Sereres Johnston, Michael Marino, Igor Ostrovsky, Josiah Walton, and Liangjian Wen. The Gratta Group (including Giorgio) on campus have been almost like a second lab group, and so I must also thank Francisco LePort, Maria Montero-Diez,

Russell Neilson, Kevin O'Sullivan, and Karl Twelker for their work and friendship.

I would also like to thank the SLAC staff. Bob Conley deserves special mention for teaching me a lot about plumbing, machining, and welding. I thank Matthias Wittgen for taking over maintenance for the slow controls and freeing up a lot of my time for physics. Ken Fouts, Tony Johnson, J. J. Russell, Knut Skarpaas VIII, Matt Swift, and Anthony Waite have done a lot to get EXO-200 running and keep it going.

The administrative staff at both SLAC and Stanford have been very helpful. Traci Kawakami and Marcia Keating have made my life so much easier. Maria Frank and Judy Meo have also helped me handle a lot of the paperwork that goes with being a graduate student.

The technicians at WIPP work hard to keep EXO-200 from breaking down. I appreciate Jon Eric Davis and Adam Rivas and their years spent working on EXO-200.

I am grateful to the United States Department of Energy for their support of EXO-200. I would like to especially thank SLAC National Accelerator Laboratory for being a great place to work over these years. I would also like to thank the Waste Isolation Pilot Plant for their hospitality. Much computing power from the National Energy Research Scientific Computing Center went into this work. The National Science Foundation also supports EXO in the United States. The National Sciences and Engineering Research Council of Canada, the Swiss National Science Foundation, the National Research Foundation of Korea, and the Russian Foundation for Basic Research provide support for EXO worldwide.

Contents

Abstract	v
Acknowledgements	vi
1 Introduction	1
2 Neutrinos	4
2.1 History	4
2.2 The Nature of Neutrinos	5
2.2.1 Dirac Particles	5
2.2.2 Majorana Particles	6
2.2.3 Majorons	7
2.3 Neutrino Mixing and Oscillation	8
2.4 Measuring Neutrino Masses	9
2.4.1 Beta Decay Endpoint	9
2.4.2 Cosmology	11
2.4.3 Double Beta Decay	11
2.4.4 Summary	17
3 Liquid Xenon for Double Beta Decay Experiments	19
3.1 Double Beta Decay Isotopes	19
3.2 The Choice of Xenon	20
3.3 Measuring Radiation	21
3.3.1 Ionization	21

3.3.2	Scintillation	22
3.3.3	Combining Ionization and Scintillation	23
3.4	Radiopurity	25
4	The EXO-200 Detector	26
4.1	Creating a Sensitive Detector	26
4.2	Time Projection Chamber	27
4.2.1	Scintillation Readout	28
4.2.2	Ionization Readout	29
4.2.3	Construction	30
4.3	Calibration	31
4.4	Infrastructure	32
4.4.1	Cryostat and Clean Rooms	32
4.4.2	Muon Veto	34
4.4.3	WIPP	34
4.5	Xenon and Recirculation	35
4.6	Summary	36
5	Data Collection and Processing	37
5.1	Signal Readout	37
5.2	Signal Reconstruction	39
5.2.1	Signal Finding	39
5.2.2	Signal Parameter Extraction	40
5.2.3	Signal Clustering	41
5.2.4	Topology	43
5.3	Corrections	43
5.3.1	Wire Gain Correction	44
5.3.2	Shielding Grid Correction	45
5.3.3	Electron Lifetime Correction	47
5.3.4	Light Collection Correction	47
5.4	Energy Calibration	47
5.4.1	Fit Model	47

5.4.2	Rotation Angle	48
5.4.3	Thallium Peak from Thorium Runs	49
5.4.4	Overall Energy Scale	50
5.5	Data Quality Cuts	54
6	Electron Lifetime	55
6.1	Electron Capture on Impurities	55
6.2	Measuring Electron Lifetime	57
6.2.1	Method	57
6.2.2	Comparison to Simulation	60
6.2.3	Practical Considerations	61
6.3	Effects of Electron Lifetime on the Energy Resolution	62
6.3.1	Position Uncertainty	62
6.3.2	Electron Lifetime Uncertainty	63
6.3.3	Rate of Change	65
6.3.4	Overall	67
6.4	Measurements of Electron Lifetime in EXO-200	67
6.4.1	Time Variation and Correction Function	67
6.4.2	Comparison with Recirculation Rate	68
6.4.3	Comparison with Gas Purity Monitor Readings	69
6.4.4	Electron Lifetime for Different Electric Fields	71
7	Muons	73
7.1	Motivation	73
7.2	Identifying Muons	74
7.2.1	Identifying Muons with the Hough Transform	75
7.2.2	Validation with Monte Carlo Simulations	77
7.2.3	Reconstruction Accuracy	78
7.2.4	Efficiency	79
7.3	The Muon Flux at WIPP	82
7.3.1	Comparison with Previous Results	83

8 Measuring Double Beta Decay	84
8.1 Event Selection	85
8.1.1 Timing-based Veto	85
8.1.2 Other Veto	86
8.1.3 Fiducial Volume	87
8.1.4 Quantities of Interest	88
8.2 Monte Carlo Simulations of Signals and Backgrounds	90
8.2.1 Simulations	90
8.2.2 PDF Generation	91
8.2.3 Agreement between Simulation and Data	94
8.3 Systematic Errors	99
8.3.1 Rate Uncertainty	99
8.3.2 Other PDFs	104
8.4 Maximum Likelihood Method	106
8.4.1 Theory	106
8.4.2 Application to EXO-200	108
8.5 Constraints	109
8.5.1 Theory	109
8.5.2 Single Site Fraction	110
8.5.3 Radon in the Xenon	110
8.5.4 Radon in the Air Gap	111
8.5.5 Normalization	111
8.6 Measurement of $2\nu\beta\beta$	111
8.6.1 Effects of Systematic Errors	113
8.6.2 Consistency Checks	114
8.6.3 Comparison to Previous Results	115
8.7 Limits on $0\nu\beta\beta\chi^0(\chi^0)$	116
8.7.1 Effects of Systematic Errors	119
8.7.2 Coverage Tests and Sensitivity	119
8.7.3 Comparison to Previous Results	122

9 Conclusions	123
A Tagging Electronic Noise	124
A.1 Motivation	124
A.2 Types of Noise Events	125
A.2.1 Unphysically Negative Signals on the Collection Wires	125
A.2.2 “Glitch” Events	125
A.2.3 APD “Bouncing” Events	126
B Correcting for Light Collection Efficiency	130
B.1 Motivation	130
B.2 Building the Light Map	131
B.2.1 Calibration Runs	131
B.2.2 Event Selection	131
B.2.3 Binning the Detector Volume	131
B.3 Correction Function	132
B.3.1 Interpolating in Cylindrical Polar Coordinates	132
B.3.2 Correcting for Light Response	134
Bibliography	140

List of Tables

2.1	Current limits on $0\nu\beta\beta$	14
2.2	Current limits on $0\nu\beta\beta\chi^0(\chi^0)$	15
5.1	Electronic shaping times	37
5.2	Shielding grid correction parameters	45
5.3	Fit function biases	48
5.4	Calibration peak ratios	52
6.1	Effect of drift time uncertainty on resolution for ideal events	63
6.2	Effect of drift time uncertainty on resolution for realistic events	64
6.3	Effect of electron lifetime uncertainty on resolution	65
6.4	Effect of electron lifetime time variance on resolution	66
8.1	Impact of timing-based vetoes	86
8.2	Energy resolution parameters	93
8.3	Calibration source rate agreement	98
8.4	Contributions to systematic uncertainty on rate	100
8.5	Effects of uncertainties on the half-life measurement	113
8.6	Limits on Majoron-emitting mode half-lives	117
8.7	Effects of uncertainties on the limits for Majoron-emitting modes	119
B.1	Resolution improvement due to correction	135

List of Figures

2.1	Neutrino mixing and mass hierarchy	10
2.2	Beta decay spectrum endpoint for massive neutrinos	10
2.3	$2\nu\beta\beta$ decay	12
2.4	Neutrinoless double beta decay modes	13
2.5	Comparison of $2\nu\beta\beta$ and $0\nu\beta\beta$ energy spectra	13
2.6	Energy spectra for Majoron-emitting double beta decay modes	15
2.7	Observation of $0\nu\beta\beta$ implies Majorana neutrinos	16
2.8	Effective Majorana mass vs. smallest neutrino mass	17
3.1	Nuclear decay structure for $A = 136$	20
3.2	Combining signals to improve energy resolution in liquid xenon	24
4.1	A conceptual drawing of EXO-200	28
4.2	Geometry of the ionization readout wire planes	29
4.3	The EXO-200 detector	30
4.4	The calibration system	31
4.5	Cutaway view of EXO-200 infrastructure	33
4.6	The WIPP site	34
4.7	The EXO-200 xenon recirculation system	35
5.1	The EXO-200 electronics	38
5.2	The signal-finding process	39
5.3	The pair-production peak on individual u wire channels	44
5.4	Fits to determine the shielding grid correction parameters	46

5.5	The result of applying the shielding grid correction to data	46
5.6	Finding the optimal rotation angle	49
5.7	Variation of the optimal rotation angle over time	50
5.8	Variation of full-absorption peak energies and their ratios over time .	51
5.9	Residuals after energy calibration	53
6.1	A fit to the full-absorption peak in a drift time bin	58
6.2	A fit to the exponential attenuation of the full-absorption peak . . .	58
6.3	A profile scan around the best-fit electron lifetime to determine confidence intervals	59
6.4	Error on reconstructed electron lifetime from simulation	60
6.5	Confidence band widths for electron lifetime measurements	61
6.6	Fits to confidence band widths for electron lifetime measurements .	64
6.7	Fits to time-varying electron lifetime	68
6.8	Electron lifetime vs. recirculation rate	69
6.9	Time history of electron lifetime and recirculation rate	70
6.10	GPM 3 reading vs. electron lifetime	71
6.11	Electron lifetime vs. electric field	72
7.1	A muon passing through EXO-200	74
7.2	Identifying a muon with the Hough transform	76
7.3	Rate for poorly reconstructing muon angles as a function of angle .	80
7.4	Efficiency of reconstructing muons as a function of angle	81
8.1	Fiducial volume	87
8.2	Increased event rates near cathode and anodes	88
8.3	Energy response due to increasing fiducial volume apothem	89
8.4	EXO-200 geometry simulated in GEANT4	91
8.5	Parameterized energy resolution functions	93
8.6	Examples of 2D PDFs	95
8.7	Single site fraction agreement between simulation and data	96
8.8	Agreement between simulation and data for a ^{228}Th calibration source	97

8.9	Ratio of simulated spectra to data as a function of energy	98
8.10	Evidence that β and γ events might have different energy scales	102
8.11	Lack of statistical power to distinguish the source of uranium backgrounds	105
8.12	Comparison of default PDF to PDF with missing channel	106
8.13	The best fit for $2\nu\beta\beta$ to the data	112
8.14	$2\nu\beta\beta$ half-life vs. time	114
8.15	Number of $2\nu\beta\beta$ decays vs. position	115
8.16	Comparison of measurements of the half-life for $2\nu\beta\beta$	116
8.17	Profile likelihoods for $0\nu\beta\beta\chi^0(\chi^0)$	117
8.18	Fits for $0\nu\beta\beta\chi^0(\chi^0)$ with 90% limits shown	118
8.19	Coverage tests for $0\nu\beta\beta\chi^0(\chi^0)$	120
8.20	Sensitivity studies for $0\nu\beta\beta\chi^0(\chi^0)$	121
8.21	KamLAND-Zen spectrum with limits for Majoron-emitting modes . .	122
A.1	Collection wires going negative noise	127
A.2	“Glitch” noise	128
A.3	APD “bouncing” noise	129
B.1	Event selection for the light map	132
B.2	Scintillation spectrum before and after correction	135
B.3	Light map ϕ slices	136
B.4	Light map z slices	137
B.5	Correction function ϕ slices	138
B.6	Correction function z slices	139

Chapter 1

Introduction

In recent decades, experiments such as Super-Kamiokande [1], SNO [2], KamLAND [3], and Daya Bay [4] have observed neutrino flavor oscillation. Together, these and other experiments have measured the mixing angles between the three neutrino flavor states, as well as two of the three mass-squared differences between the mass states. However, despite this progress, there are still unanswered questions about neutrinos: Are neutrinos Dirac or Majorana particles? Is the hierarchy of neutrino masses similar to other massive fermions, or is it inverted? What is the absolute mass scale for neutrinos, and why is their mass so much smaller than other fermions?

A process known as neutrinoless double beta decay ($0\nu\beta\beta$) may provide the answers for all of these questions. This process, which has yet to be observed, requires neutrinos to be Majorana fermions. A family of “see-saw” mechanisms could explain why Majorana neutrinos have such small masses compared to charged leptons and quarks. Furthermore, the rate of neutrinoless double beta decay depends on the neutrino mass scale and mass hierarchy.

The Enriched Xenon Observatory (EXO) is a series of experiments that seek to observe neutrinoless double beta decay in ^{136}Xe . The first experiment, EXO-200, began taking data in 2011. Data from EXO-200’s “Run 1” was used to make the first observation in ^{136}Xe of the standard-model-allowed mode of two-neutrino-emitting double beta decay ($2\nu\beta\beta$) [5]. Data from EXO-200’s “Run 2a” was used to place a lower limit on the half-life of the neutrinoless mode [6]. The upper limit on neutrino

masses derived from this half-life limit almost completely rules out a controversial claim of observation of $0\nu\beta\beta$ in ^{76}Ge [7].

Searching for an extremely rare process has unique difficulties. Natural radioactivity can create interactions in the xenon that could be mistaken for neutrinoless double beta decay. EXO-200, which was built to have a low rate of these background events, allows rejection the remaining background events through two chief means. The first is through energy resolution. For the $0\nu\beta\beta$ decay mode, only two electrons are emitted, and they carry the full energy of the decay, known as the “Q value”. The signals collected by EXO-200 can be used to precisely determine the energy of an event, rejecting background events with energies away from the Q value of ^{136}Xe . The remaining rate of background events with energies in the “region of interest” near the Q value can be estimated by looking for correlated gamma lines elsewhere in the energy spectrum. Secondly, EXO-200’s time projection chamber design admits the use of event topology to reject some backgrounds with energies in the region of interest.

This dissertation describes a new analysis of the Run 2a data from EXO-200. This analysis makes use of improved analysis techniques and a better understanding of the detector in order to reduce systematic errors. Chapter 2 describes the theoretical motivation for a double beta decay experiment. Chapter 3 details the motivation behind the choice of liquid xenon for the experiment and the physics of a liquid xenon detector. A detector sensitive to double beta decay must satisfy many criteria. Chapter 4 describes these criteria and their realization in EXO-200.

The system that EXO-200 uses to collect data is described in chapter 5. This chapter also describes how data is processed from its raw form into physically meaningful quantities. One of the largest corrections applied in this processing is a correction for the electron lifetime: how long drifting ionization electrons survive in EXO-200. The physics behind that, the measurement of that quantity, and the correction are described in extra detail in chapter 6.

Finally, two analyses are presented. The first, in chapter 7, details a new measurement of the muon flux at the underground WIPP site where EXO-200 is located. The primary benefit of identifying and measuring muons is to reduce and constrain

backgrounds that might interfere with double beta decay measurements. The flux measurement is also of interest to any other experiment at the WIPP site. Lastly, Chapter 8 describes a more precise measurement of the half-life of the two-neutrino-emitting mode than previously reported by EXO-200. It also reports limits on the half-lives of modes of neutrinoless double beta decay that emit hypothetical Majoron particles ($0\nu\beta\beta\chi^0(\chi^0)$).

Chapter 2

Neutrinos

2.1 History

Neutrino physics first arose in the study of beta decay. In the most basic form of beta decay, a neutron in a nucleus decays to a proton and emits an electron. Early nuclear physicists observed that unlike alpha particles and gamma rays, which are emitted in monoenergetic lines from nuclear processes, beta decay electrons are emitted with a continuous range of energies up to some maximum energy (known as the Q value). Rather than abandoning the principle of energy conservation, Wolfgang Pauli in 1930 proposed a massless particle that could carry away the “missing energy” in beta decays. In 1933, Enrico Fermi [8] combined the idea with Werner Heisenberg’s nucleon model to form a theory of beta decay. The proposed massless particle was eventually dubbed the “neutrino”.

For several decades after Pauli proposed the neutrino, the only evidence for its existence was indirect. It was not until 1956 that Cowan and Reines [9] observed the electron antineutrino by looking for the inverse beta decay reaction of antineutrinos from a nuclear reactor. In 1962, Lederman, Schwartz, and Steinberger [10] used a beam of pions decaying to muons to create a beam of neutrinos. These neutrinos interacted to form only muons and not electrons. This established the existence of the muon neutrino, distinct from the electron neutrino. These distinct types of neutrino with respect to the weak interaction are referred to as “flavors”.

Beginning in the early 1970s, the Homestake experiment [11] observed a large deficit in the flux of electron neutrinos coming from the sun compared to the expected flux. Eventually, it was suggested that this “solar neutrino problem” was due to neutrinos oscillating between flavors in flight. The Super-Kamiokande experiment first observed evidence for neutrino oscillation in 1998 [1] by looking at neutrinos produced in the upper atmosphere. The solar neutrino problem was not definitively resolved until 2001 when the Sudbury Neutrino Observatory [2] simultaneously measured the flux of electron neutrinos and the total flux of all neutrinos from the sun. The total flux matched expectations, meaning that the missing electron neutrinos oscillated to other flavors in flight. For neutrinos to oscillate, they must have mass, and their mass eigenstates must be mixtures of the flavor eigenstates that participate in the weak interaction. Neutrino oscillations and neutrino masses have now become the subject of much research.

2.2 The Nature of Neutrinos

2.2.1 Dirac Particles

In the standard model of particle physics, neutrinos are massless and have left-handed chirality, forming a doublet with the left-handed charged leptons under the $SU(2)$ symmetry of the weak force. The right-handed charged leptons form an $SU(2)$ singlet, and there are no right-handed neutrinos. The mass of a charged lepton is generated by a Yukawa coupling between the left-handed lepton, the right-handed lepton, and the Higgs field. Since neutrinos have been observed to have mass, then, it seems natural to generate their mass in an analogous way. Masses due to this mechanism are known as Dirac masses. This would introduce right-handed neutrinos (and left-handed antineutrinos) that would not interact weakly. However, given the observed smallness of neutrino masses, this requires Yukawa couplings much, much smaller than those of charged leptons and quarks. Such small Yukawa couplings are difficult to explain.

2.2.2 Majorana Particles

Neutrinos do not carry charge, and so mass terms of the form

$$\mathcal{L}_m^\nu = -\frac{1}{2}\bar{\nu}_L m_{LL}^\nu \nu_L^c + h.c. \quad (2.1)$$

can be added to the standard model Lagrangian. Such terms are not allowed for charged leptons, since they would violate electric charge conservation. No fundamental rule requires total lepton number conservation in the standard model, and so it is not troublesome that these terms violate it. Particles that have such mass terms satisfy the Majorana equation:

$$i\gamma^\mu \partial_\mu \psi - m\psi^c = 0 \quad (2.2)$$

The Majorana equation is similar to the Dirac equation, except that it includes the charge conjugate ψ^c of the spinor ψ . Adding the supplemental condition that $\psi^c = \psi$ results in a single neutral particle solution. This means Majorana particles can be described as their own antiparticles.

Mass terms like those in eq. (2.1) for left-handed neutrinos violate weak isospin symmetry, but can arise from other phenomena that conserve it. In 1979, S. Weinberg [12] proposed a dimension 5 operator coupling the neutrino fields to the standard model Higgs field. However, such a term is not renormalizable and would have to reflect new physics at a mass scale much larger than the weak scale. Introducing massive right-handed Majorana neutrinos provides one possible mechanism. With right-handed neutrinos, the most general way to write the mass terms for one generation is

$$\mathcal{L}_m^\nu = -\frac{1}{2} \begin{pmatrix} \bar{\nu}_L & \bar{\nu}_R^c \end{pmatrix} \begin{pmatrix} m_{LL} & m_{LR} \\ m_{LR} & M_{RR} \end{pmatrix} \begin{pmatrix} \nu_L^c \\ \nu_R \end{pmatrix} + h.c. \quad (2.3)$$

where m_{LR} is a Dirac mass and m_{LL} and M_{RR} are Majorana masses. Diagonalizing this yields two mass eigenstates:

$$m_\pm = \frac{1}{2} \left| (M_{RR} + m_{LL}) \pm \sqrt{(M_{RR} - m_{LL})^2 + 4m_{LR}^2} \right| \quad (2.4)$$

The case in which $m_{LL} = 0$ (which it must be to conserve weak isospin symmetry) and $M_{RR} \gg m_{LR}$ is known as a type I seesaw mechanism. In this case, $m_- \approx m_{LR}^2/M_{RR}$ and $m_+ \approx M_{RR}$. Suppose m_{LR} is on the order of quark or charged lepton masses. If M_{RR} is much greater than the electroweak scale, then the light neutrino masses (m_-) take on small values. The light neutrinos are a mixture of left and right-handed neutrinos, but the right-handed amplitude is suppressed by m_{LR}/M_{RR} and so is negligible.

2.2.3 Majorons

Instead of right-handed neutrinos, suppose there is instead a new set of Higgs-like particles ϕ^0 , ϕ^- and ϕ^{--} that forms a complex triplet under the weak interaction. If, analogously to the standard Higgs, ϕ^0 acquires a vacuum expectation value $\langle\phi^0\rangle$ through the spontaneous breaking of lepton number symmetry, then interactions with the neutrino fields can give Majorana neutrino mass terms with $m_{LL} = \lambda\langle\phi^0\rangle$, where λ is a Yukawa coupling.

If the vacuum expectation value $\langle\phi^0\rangle$ were small, this would lead to small neutrino masses. Furthermore, analogously to the standard Higgs mechanism, the spontaneous symmetry breaking would lead to a Goldstone boson known as the Majoron. This theory, due to Gelmini and Roncadelli [13] and Georgi et al. [14] is disfavored due to precise measurements of the width of the Z^0 decay to invisible channels.

Another theory, due to Chikashige et al. [15] introduces a Higgs-like singlet that acquires a vacuum expectation value and gives a Majorana mass to right-handed neutrinos. The light neutrinos acquire mass through a seesaw mechanism as above, but the symmetry breaking again creates a Majoron. This theory contributes to the Z^0 width at a level less than the current uncertainty on the measurement, and so is not ruled out. Other models [16, 17] have been proposed that are consistent with LEP data on the Z^0 width.

2.3 Neutrino Mixing and Oscillation

Neutrinos are observed to oscillate, which implies that the mass eigenstates that propagate are different from the flavor eigenstates that participate in the weak interaction. Neutrino mixing is described by the unitary Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix U . In the case of three mass eigenstates and 3 flavor eigenstates, U can be parameterized by three Euler rotation angles, one Dirac CP violation phase, and two Majorana CP violation phases (if neutrinos are Majorana particles):

$$U = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \quad (2.5)$$

$$= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{\frac{i\alpha_{21}}{2}} & 0 \\ 0 & 0 & e^{\frac{i\alpha_{31}}{2}} \end{pmatrix}$$

where the s_{ij} denote $\sin \theta_{ij}$ and c_{ij} denote $\cos \theta_{ij}$. δ is the Dirac CP violation phase, and α_{21} and α_{31} are the Majorana CP violation phases.

A flavor eigenstate, such as a neutrino produced by the decay of a lepton, is a superposition of mass eigenstates:

$$|\nu_\ell\rangle = \sum_i U_{\ell i}^* |\nu_i\rangle \quad (2.6)$$

This mixing is illustrated in fig. 2.1.

The time evolution of a mass eigenstate is simply

$$|\nu_i(t)\rangle = e^{-i(E_i t - \vec{p}_i \cdot \vec{x})} |\nu_i(0)\rangle \quad (2.7)$$

Neutrinos have such small masses that they are always observed in a ultrarelativistic state, in which $t \approx L$ and $p \approx E$, where E is the total energy of the neutrino. In this limit, the phase in the exponent for eq. (2.7) becomes $m_i^2 L / (2E)$. Combining this and eq. (2.6), the probability for a neutrino with flavor α and energy E to

oscillate to a neutrino with flavor β after traveling distance L is

$$P_{\alpha \rightarrow \beta} = |\langle \nu_\beta | \nu_\alpha(L) \rangle|^2 = \delta_{\alpha\beta} - 4 \sum_{i>j} \operatorname{Re} (U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin^2 \left(\frac{\Delta m_{ij}^2 L}{4E} \right) + 2 \sum_{i>j} \operatorname{Im} (U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^*) \sin^2 \left(\frac{\Delta m_{ij}^2 L}{2E} \right) \quad (2.8)$$

Thus, the probability depends only on the parameters in the mixing matrix U , the difference between the squares of the masses Δm_{ij}^2 , and the ratio of distance travelled to neutrino energy L/E . If there is no CP violation, U is real and the second sum is zero. Since the Majorana phases only enter U as a diagonal factor, they do not affect the oscillation probability, and thus oscillation experiments cannot determine the Majorana nature of the neutrino.

2.4 Measuring Neutrino Masses

From oscillation experiments, it is known that the mass-squared splittings in eq. (2.8) are nonzero. The mass state ν_1 has been defined as the one with the largest ν_e component. Under this convention, the sign of Δm_{21}^2 is known from the effect of solar matter on neutrinos emitted by the sun (known as the MSW effect), while only the magnitude of Δm_{31}^2 is known. The mass splittings provide lower bounds on the masses of two mass eigenstates. This does not determine the absolute mass scale, however, and the lightest eigenstate may have zero mass. Since the sign of Δm_{31}^2 remains unknown, it is unknown which mass eigenstate is the lightest. Figure 2.1 summarizes the situation. Other experiments, described below, may be able to measure the absolute mass scale and possibly determine the hierarchy of the mass eigenstates. However, to date, such experiments have only provided upper bounds on the mass.

2.4.1 Beta Decay Endpoint

When a nucleus beta decays, it emits an electron antineutrino, which carries away some energy from the decay. However, this amount may be very small, and so in rare

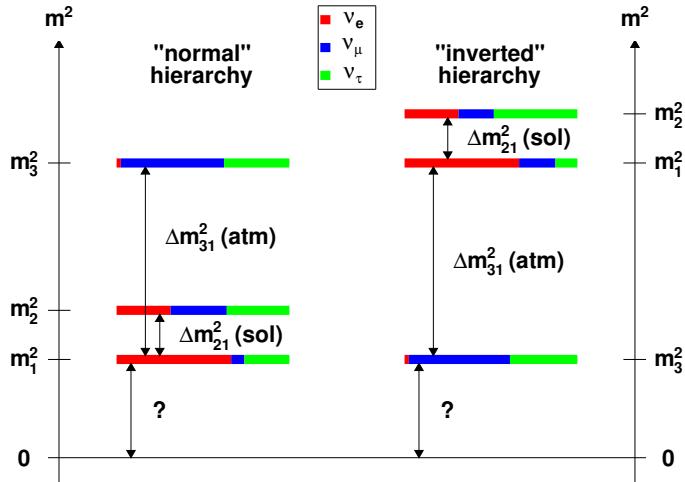


Figure 2.1: The neutrino mass eigenstates (represented by the horizontal bands) are each a mixture of the flavor eigenstates (represented by the different colors). The mixing angles in the PMNS matrix provide the flavor compositions of each mass state. The splittings between the squares of the masses of the mass eigenstates are known (but not shown to scale), though the sign of Δm_{31}^2 remains unknown, allowing for the possibility of an “inverted” mass hierarchy. The absolute scale of the neutrino masses is also unknown.

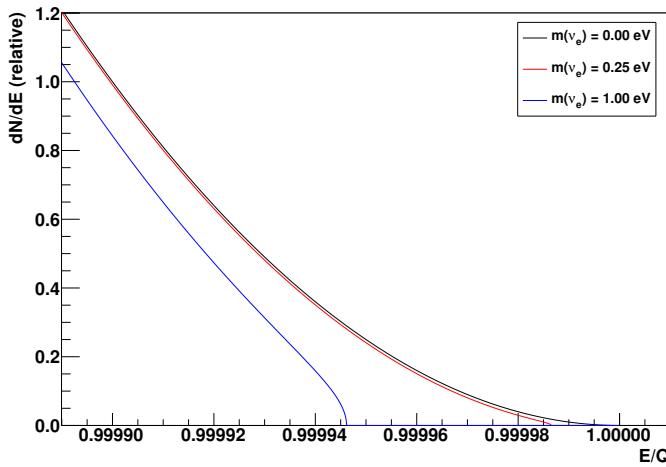


Figure 2.2: The endpoint of the electron energy spectrum for beta decay for several neutrino masses. Experiments like Mainz [18], Troitsk [19], and KATRIN [20] aim to measure the neutrino mass by looking for the slight distortion of the spectrum close to the Q value of the decay due to the neutrino mass.

cases, the electron may carry away nearly all of the energy. A neutrino with nonzero rest mass will always take some energy to create, however, limiting the maximum energy of the electron and distorting the spectrum near the upper end of its range. This is illustrated in fig. 2.2.

An experiment looking for this distortion in the beta decay spectrum will measure the effective mass squared of the electron neutrino:

$$m^2(\nu_e) = \sum_j |U_{ej}|^2 m^2(\nu_j) \quad (2.9)$$

The best limits come from the Mainz and Troitsk experiments, which examined the spectrum of tritium. Mainz measured $m(\nu_e) \leq 2.3 \text{ eV}/c^2$ [18], while Troitsk measured $m(\nu_e) \leq 2.05 \text{ eV}/c^2$ [19], both at the 95% confidence level. The KATRIN experiment, which will begin operation soon and also examine the spectrum for tritium, hopes to be sensitive to masses greater than $0.2 \text{ eV}/c^2$ [20].

2.4.2 Cosmology

Due to their small masses and large velocities, neutrinos in the early universe did not clump together as much as other matter. The distribution of matter in the universe, then, is sensitive to the mass ratio of neutrinos to other matter. Cosmological surveys of structure and anisotropy can potentially measure $\sum m_\nu$. Cosmic microwave background data from the Planck satellite provide the best limit to date. The Planck data yield $\sum m_\nu < (0.23 - 1.08) \text{ eV}/c^2$. The spread is due to the choice of model and which datasets are combined in the analysis [21]. Further data from other surveys may push the sensitivity down to the eV scale [22].

2.4.3 Double Beta Decay

Two-Neutrino-Emitting Double Beta Decay

Standard double beta decay ($2\nu\beta\beta$) is a standard model process in which two neutrons simultaneously decay to two protons, emitting two electrons and two antineutrinos.

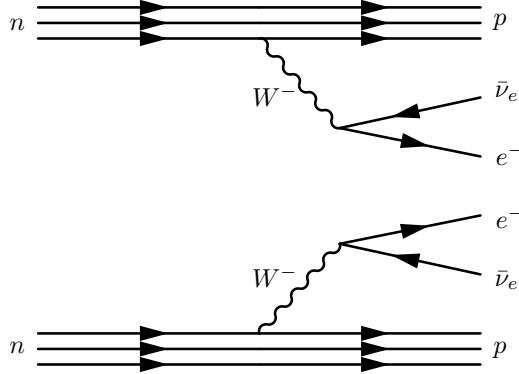


Figure 2.3: The standard-model-allowed mode of two-neutrino-emitting double beta decay ($2\nu\beta\beta$).

Figure 2.3 illustrates this process. It has been observed in a number of isotopes with half lives between 7.1×10^{18} yr for ^{100}Mo [23] and 2.2×10^{21} yr for ^{136}Xe [6]. The rate goes like:

$$[T_{1/2}^{2\nu}]^{-1} = G^{2\nu}(E, Z) |\mathcal{M}^{2\nu}|^2 \quad (2.10)$$

where $G(E, Z)$ is a known phase space factor and \mathcal{M} is the nuclear matrix element, which can be calculated using various nuclear structure models.

Any experiment looking for more exotic forms of double beta decay should also observe this standard mode. Studying the two-neutrino-emitting mode can constrain it as a background to these searches. Moreover, measuring the rate in turn measures the nuclear matrix element. For more exotic modes, models must be used to predict the nuclear matrix elements in order to translate the rates (or limits on the rates) into statements about neutrinos. These models can be evaluated by comparing their predictions to the measurement of the $2\nu\beta\beta$ matrix element.

Neutrinoless Double Beta Decay

If neutrinos are Majorana particles, then there is the potential to observe neutrinoless double beta decay ($0\nu\beta\beta$). The simplest mechanism, shown in fig. 2.4a, has a light Majorana neutrino being exchanged. In this process two electrons are emitted, but unlike $2\nu\beta\beta$, the electrons carry away the entire energy of the decay. Thus, this decay

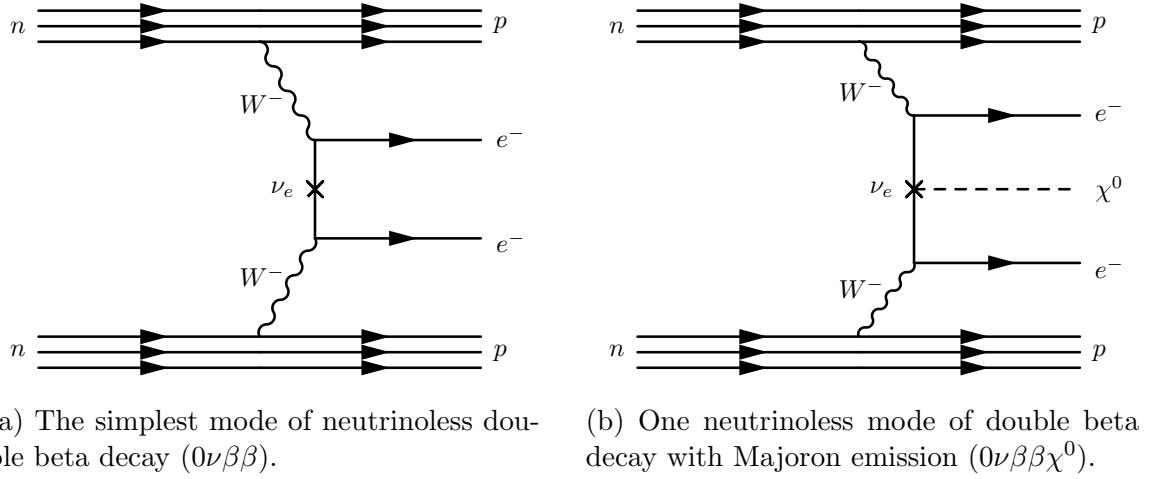


Figure 2.4: Neutrinoless double beta decay modes

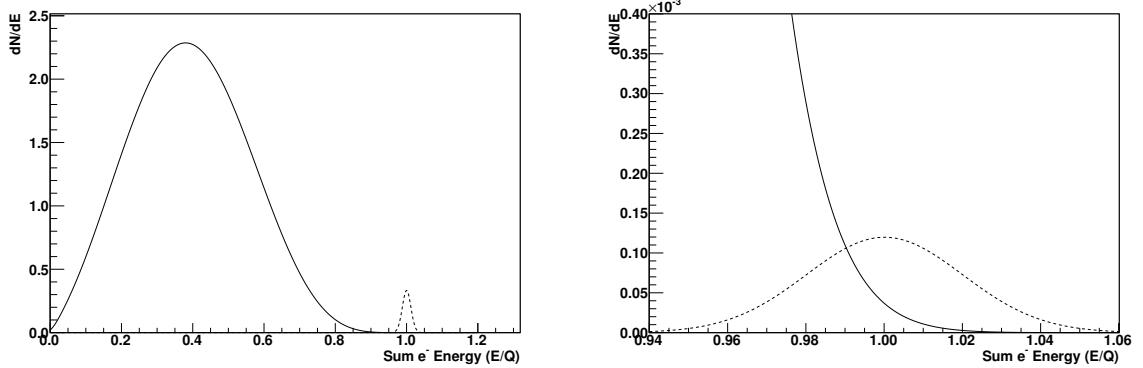


Figure 2.5: The electrons emitted in two-neutrino-emitting double beta decay have a continuous sum energy spectrum (solid line), while the energies of the electrons emitted in neutrinoless double beta decay sum to the Q value (dashed line). Here the spectra have been smeared with a 2% σ/E energy resolution. On the left (right), the rate of $0\nu\beta\beta$ is 100 (1×10^6) times less than that of $2\nu\beta\beta$. If the rate of $0\nu\beta\beta$ is small compared to $2\nu\beta\beta$, good energy resolution is needed to distinguish the two modes.

Table 2.1: Results of searches for neutrinoless double beta decay. Effective masses are as reported by the experiments, and no attempt has been made to standardize the nuclear matrix element calculations used. Where a range has been reported, it is due to variation between different nuclear matrix element calculations.

isotope	half-life ($\times 10^{22}$ yr)	$ \langle m_\nu \rangle $ (eV)	experiment
^{48}Ca	> 5.8	$< (3.5 - 22)$	$\text{CaF}_2(\text{Eu})$ scintillators [24]
^{76}Ge	> 1900	< 0.35	Heidelberg-Moscow [25]
^{76}Ge	2230^{+440}_{-310}	0.32 ± 0.03	Subset of H-M [7]
^{82}Se	> 36	$< (0.89 - 2.43)$	NEMO-3 [26]
^{96}Zr	> 0.92	$< (7.2 - 19.5)$	NEMO-3 [26]
^{100}Mo	> 110	$< (0.45 - 0.93)$	NEMO-3 [26]
^{116}Cd	> 17	< 1.7	Solotvina [27]
^{130}Te	> 280	$< (0.30 - 0.71)$	CUORICINO [28]
^{136}Xe	> 1600	$< (0.14 - 0.38)$	EXO-200 [6]
^{136}Xe	> 1900	$< (0.12 - 0.25)$	KamLAND-Zen [29]
^{150}Nd	> 1.8	$< (4.0 - 6.3)$	NEMO-3 [26]

mode should be observable in a detector with good energy resolution. Figure 2.5 illustrates the difference in the sum electron spectrum for the two modes. If $0\nu\beta\beta$ does indeed proceed through light neutrino exchange, then the rate is given by

$$[T_{1/2}^{0\nu}]^{-1} = G^{0\nu}(E, Z) |\langle m_\nu \rangle|^2 |\mathcal{M}^{0\nu}|^2 \quad (2.11)$$

where $G^{0\nu}(E, Z)$ is a calculable phase-space factor, $\mathcal{M}^{0\nu}$ is the nuclear matrix element, which can be estimated with models, and

$$\langle m_\nu \rangle = \sum_i U_{ei}^2 m_i \quad (2.12)$$

is known as the “effective mass”. A number of experiments have attempted to look for neutrinoless double beta decay, but it has not yet been observed. A controversial claim by a subset of the Heidelberg-Moscow collaboration report discovery of $0\nu\beta\beta$ in ^{76}Ge [7], but this is in tension with results from ^{136}Xe experiments [6, 29]. Table 2.1 summarizes the present status of experiments.

Majoron-Emitting Double Beta Decay

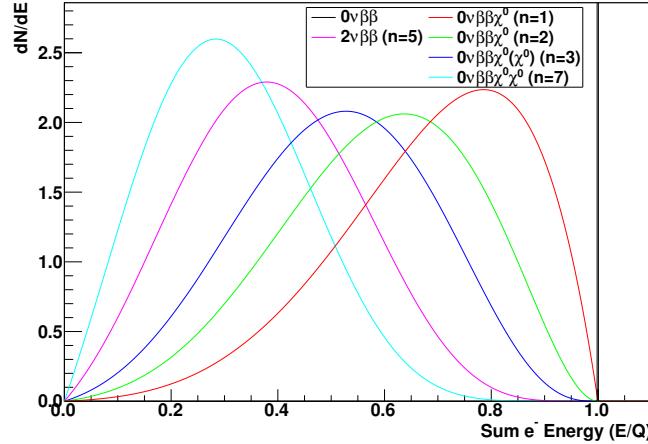


Figure 2.6: The sum electron energy spectra for different Majoron-emitting double beta decays. Decays with different spectral indices are clearly distinguishable. $2\nu\beta\beta$ has a spectral index of 5 and is plotted for comparison. All spectra have been normalized.

Table 2.2: Limits on the half-lives of Majoron-emitting double beta decay modes in ^{136}Xe from the KamLAND-Zen experiment [30].

mode	spectral index (n)	KamLAND-Zen 90% C.L. $T_{1/2}$ (yr)
$0\nu\beta\beta\chi^0$	1	$> 2.6 \times 10^{24}$
$0\nu\beta\beta\chi^0$	2	$> 1.0 \times 10^{24}$
$0\nu\beta\beta\chi^0(\chi^0)$	3	$> 4.5 \times 10^{23}$
$0\nu\beta\beta\chi^0\chi^0$	7	$> 1.1 \times 10^{22}$

There is also a possibility for double beta decay with the emission of one or more Majorons ($0\nu\beta\beta\chi^0(\chi^0)$), illustrated in fig. 2.4b. While the sum energy spectrum is not a sharp peak, it is still distinguishable from the two-neutrino-emitting mode by its shape, due to the different number of degrees of freedom for the invisible particle(s). As mentioned in section 2.2.3, a number of different theories predict Majoron emission. They differ in the nature of the Majoron (chiefly whether it carries lepton number, and whether it is a Goldstone boson) and in the number emitted. Regardless, the energy

spectrum for the observable electrons is characterized by a single integer known as the “spectral index”. $2\nu\beta\beta$, which similarly has invisible particles carrying away energy, has a spectral index of 5. Bamert et al. [17] provide a thorough discussion of the different Majoron-emitting modes of double beta decay. Figure 2.6 shows the spectra for proposed decay modes. The current best limits for Majoron-emitting modes in ^{136}Xe come from the KamLAND-Zen collaboration [30] and are shown in table 2.2. The rate of a Majoron-emitting decay is

$$\left[T_{1/2}^{\chi^0} \right]^{-1} = G^{\chi^0} (E, Z) |\langle g_{ee} \rangle|^{2N} |\mathcal{M}^{\chi^0}|^2 \quad (2.13)$$

where again G is a phase space factor and \mathcal{M} is a nuclear matrix element. N is the number of Majorons emitted. Theories predict $N = 1$ or $N = 2$. $\langle g_{ee} \rangle$ is the effective coupling constant of the Majoron to the neutrino.

Neutrinoless Double Beta Decay as a Black Box

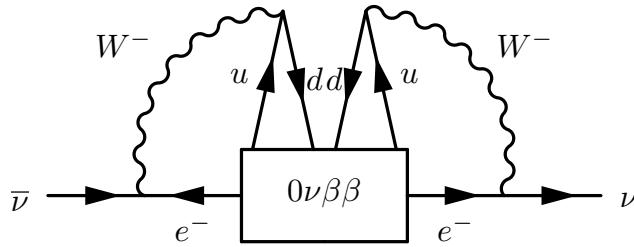


Figure 2.7: Regardless of the mechanism by which $0\nu\beta\beta$ occurs, observation of neutrinoless double beta decay would imply Majorana neutrinos. This diagram corresponds to a Majorana mass term, treating $0\nu\beta\beta$ as a black box.

Even if neutrinoless double beta decay does not proceed through light neutrino exchange, any observation of neutrinoless double beta decay would imply that neutrinos are Majorana particles. This is illustrated in fig. 2.7. Treating $0\nu\beta\beta$ as a black box, a Majorana mass term can be constructed for neutrinos.

2.4.4 Summary

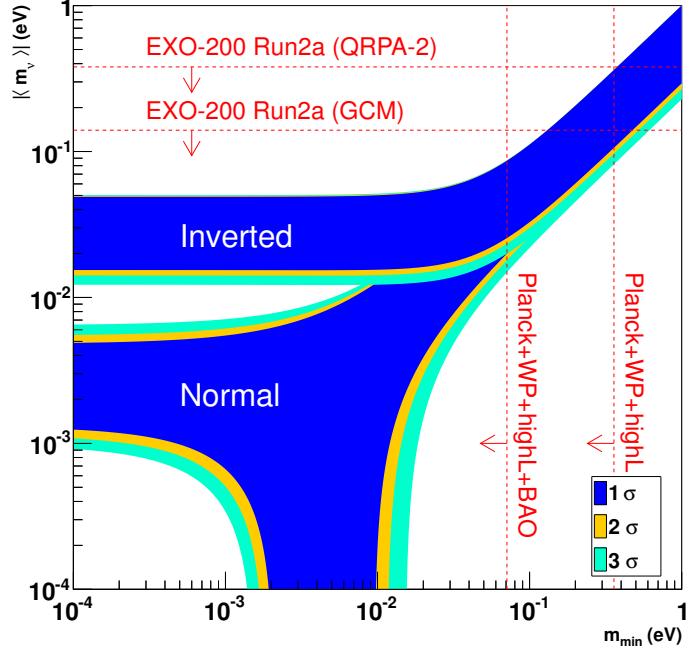


Figure 2.8: The allowed values for the effective Majorana mass as a function of the smallest neutrino mass for both the inverted and normal mass hierarchies. The widths of the bands depend on the Majorana phases. Uncertainties on the mixing angles and mass splittings widen these bands. The mixing angles, mass splittings, and uncertainties used are from a global fit by Forero et al. [31]. The double beta decay limit comes from EXO-200 [6], with the spread due to different nuclear models. The cosmological limit comes from the Planck 2013 data [21], with the spread due to choice of model and which datasets are combined in the analysis. Constraints on the sum of the neutrino masses combined with the mass splittings yield an upper bound for the smallest neutrino mass.

The different experiments presented above are all complimentary. Double beta decay experiments, for example, may be able to determine the nature of neutrinos and, if they are Majorana particles, their mass scale. Oscillation experiments measure the mixing angles and mass splittings and should be able to determine the hierarchy, but as eq. (2.8) showed, they cannot determine the Majorana nature. Figure 2.8 shows one way in which these results can be combined. If, for example, the hierarchy

is determined to be inverted, but double beta decay experiments rule out the corresponding mass range, then neutrinos are not Majorana particles. Meanwhile, tritium endpoint experiments and cosmology can potentially determine the mass scale, but cannot determine the Majorana nature of neutrinos.

Chapter 3

Liquid Xenon for Double Beta Decay Experiments

3.1 Double Beta Decay Isotopes

For nuclei with even mass number A , elements with odd atomic number Z have unpaired nucleons, which is less energetically favorable than elements with even Z . For some even values of A , then, there exist two nuclei that differ by $\Delta Z = 2$ that are each stable against β^\pm decay. This is depicted in fig. 3.1. However, one of these nuclei will have a lower mass than the other, and so double beta decay, as described in section 2.4.3, or the analogous double electron capture or double positron decay can proceed. Thirty-five naturally occurring isotopes have this property and can undergo double beta decay.

An ideal isotope for a double beta decay experiment would be abundant. It would also have a high Q value. Higher Q value decays proceed at a faster rate. Furthermore, higher Q values place the decay energy above that of some natural sources of radioactivity, reducing backgrounds. Finally, it would be easy to use in a low radioactive background detector. Unfortunately, no single isotope is ideal, though several still have many of these favorable qualities. Having several choices is good, because if neutrinoless double beta decay is detected, it will be important to confirm this detection in other isotopes. Furthermore, the decay rate (eq. (2.11)) depends

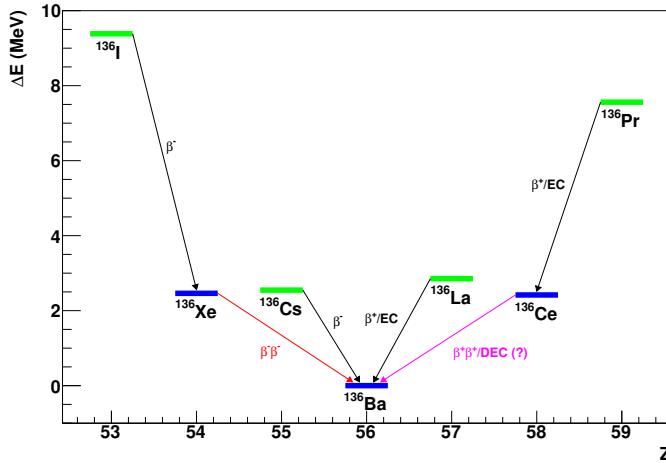


Figure 3.1: Nuclear decay structure for nuclei with $A = 136$. ^{136}Xe is an even-even nucleus and is energetically forbidden to beta decay to ^{136}Cs . Instead, it double beta decays to ^{136}Ba . ^{136}Ce should also decay to ^{136}Ba by double β^+ decay and double electron capture. However, this has not yet been observed.

on nuclear matrix elements which must be estimated and have roughly a factor of 2 uncertainty for most isotopes. This means observation in multiple isotopes would be needed to precisely infer neutrino properties.

3.2 The Choice of Xenon

Ultimately, of the available isotopes, EXO has chosen to focus on xenon. Both ^{134}Xe and ^{136}Xe should double beta decay, though ^{136}Xe has a higher Q value of 2457.8 keV [32]. Xenon has many properties that make it desirable for a double beta decay experiment:

- First and foremost, xenon can serve as both as the source and detector of double beta decay events. Noble gases and liquids have been deployed as radiation detectors for decades. A monolithic detector minimizes the other materials needed to build the detector that might be sources of radioactive backgrounds. Electrons from the double beta decay do not have to pass through other media before reaching the detector, allowing fewer energy losses and better energy resolution.

- The Q value for ^{136}Xe decays is higher than most γ rays from common radioactive nuclides. ^{208}Tl , which occurs on the thorium chain and emits a 2615 keV gamma ray, and ^{214}Bi , which occurs on the uranium chain and emits a 2469 keV gamma ray are notable exceptions. γ rays with higher energies than the Q value can potentially deposit part of their energy in the detector before scattering out, creating an event with energy close to the Q value.
- The natural abundance of ^{136}Xe is 8.9%. Furthermore, xenon is a gas at standard temperatures and pressures, making it simple to process and enrich in ^{136}Xe using ultracentrifugation.
- Xenon is a noble element, and so it is relatively easy to purify of all chemically active contaminants. Furthermore, this purification can be done continuously by recirculating the xenon.
- The isotopes formed in xenon by cosmogenic activation are short-lived, so the xenon only needs a short period underground and a chemical purification before it is ready to be used.
- Xenon can be easily reused and transferred between experiments. This allows the opportunity to use xenon in complimentary or novel detector designs. Smaller experiments can help amortize the cost of larger experiments.
- The barium daughter ion could potentially be tagged, reducing backgrounds immensely. (This technique, however, is not used in EXO-200.)

3.3 Measuring Radiation

3.3.1 Ionization

Radiation that deposits energy in xenon creates electron-ion pairs. The average energy needed to create an electron-ion pair is called the W -value, and is (15.6 ± 0.3) eV [33] for liquid xenon. The number of pairs does not follow a Poisson distribution, as

might be expected. Rather

$$(\Delta N_i)^2 = F \bar{N}_i \quad (3.1)$$

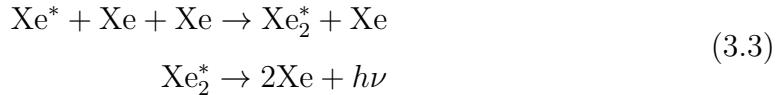
where \bar{N}_i denotes the mean number of electron-ion pairs and ΔN_i denotes the standard deviation. F is known as the Fano factor [34] and is estimated to be 0.059 in liquid xenon [35]. This implies the energy resolution of a perfect liquid xenon detector should be

$$\frac{\sigma}{E} = \sqrt{\frac{FW}{E}} = \frac{0.92}{\sqrt{E \text{ (eV)}}} \quad (3.2)$$

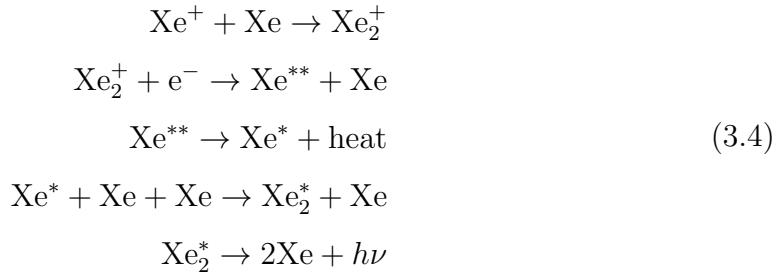
which would be comparable to the energy resolution of Ge detectors. However, no liquid xenon experiments to date have come close to this energy resolution, or even that for the Poisson limit. The reason for this discrepancy remains unclear.

3.3.2 Scintillation

Radiation in xenon can directly excite the atoms, in addition to creating electron-ion pairs. These excited atoms pair with other xenon atoms to create excited dimers, which then emit light when they de-excite:



Additionally, ions and electrons can recombine, which again produces light:



The light produced in xenon is in the vacuum ultraviolet (VUV) spectrum, and in liquid xenon, the emission peaks at 177.6 nm. Xenon is transparent at this wavelength. The scintillation is produced soon after the interaction, with the excited

dimers decaying with a decay time of 2 ns to 4 ns for the singlet state and 20 ns to 25 ns for the triplet state [36].

3.3.3 Combining Ionization and Scintillation

Electron-ion pairs can recombine to produce one scintillation photon. The probability for this to occur depends both on the ionization density and the strength of the electric field applied. A stronger field will cause more electrons to drift away from the ions before they can recombine. Particles such as α particles that are highly ionizing create a higher ionization density, and so an increased ratio of scintillation to ionization can be used to distinguish them from β and γ radiation. While the recombination of ions to produce scintillation can be modeled [37, 38], collecting both ionization and scintillation signals allows an unambiguous measurement of the total energy deposited in the detector.

A toy example of combining the ionization and scintillation signals is shown in fig. 3.2. Even assuming a less-than-perfect correlation between ionization and scintillation due to effects such as imperfect collection and detector noise, this example shows how combining the signals can yield a dramatic improvement in resolution. This has been demonstrated to work in practice [39, 40]. However, the energy resolution for the combined signals in such experiments still falls short of the Poisson limit, suggesting other factors contribute to worsening the energy resolution in liquid xenon.

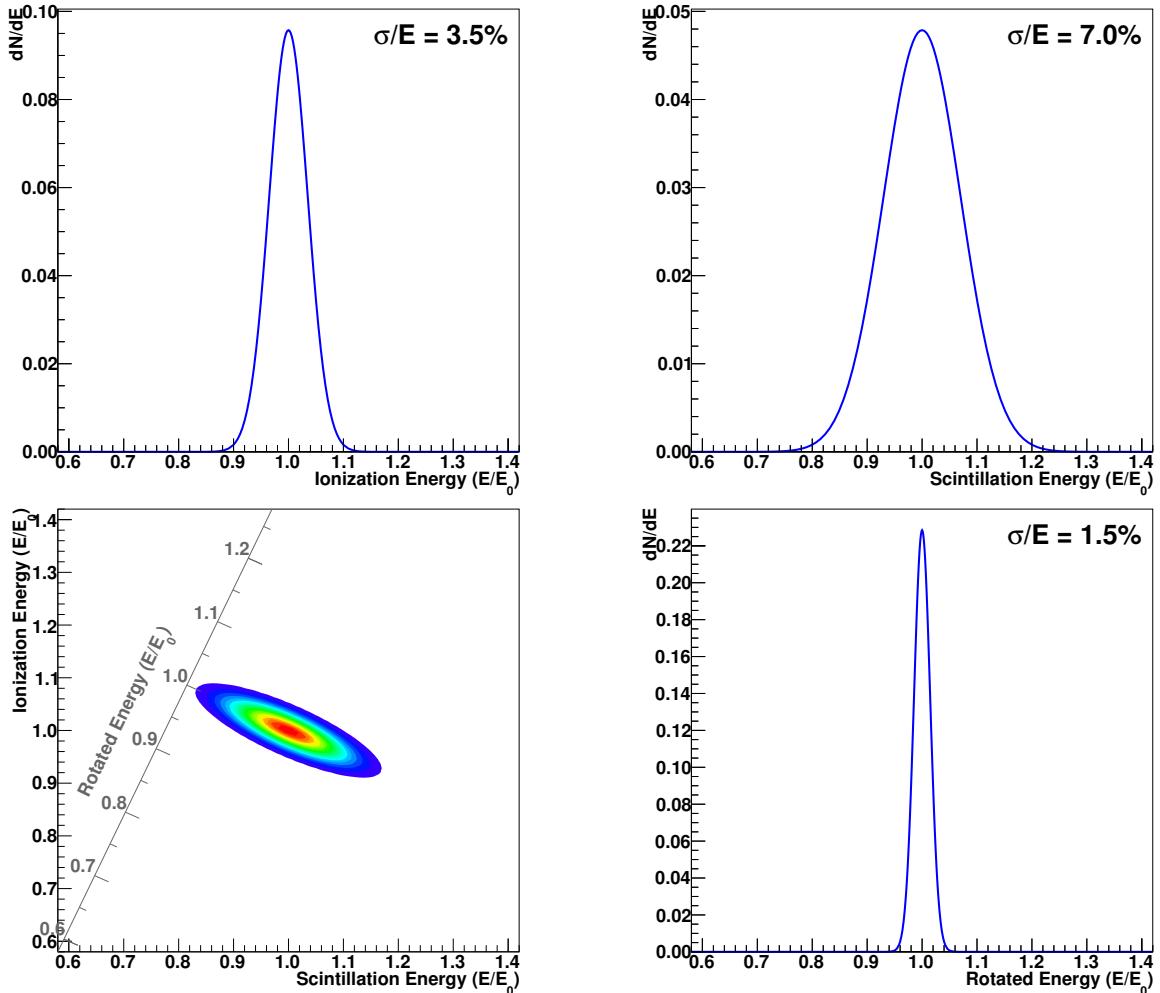


Figure 3.2: A toy example of how the energy resolution of a monoenergetic line can be improved by combining signals in liquid xenon. The ionization signal (top left) and the scintillation signal (top right) show poor energy resolution individually. When the signals are anticorrelated (shown bottom left, with coefficient -0.8), they can be projected onto a rotated axis to yield an improved energy resolution (bottom right).

3.4 Radiopurity

Liquid xenon is dense ($\sim 3 \text{ g cm}^{-3}$) and has a high atomic number ($Z = 54$). It provides good shielding against low-energy γ rays from external sources. The attenuation length of a 100 keV photon is less than 2 mm [41]. However, this shielding grows worse for higher-energy photon. A 1 MeV photon has an attenuation length of nearly 6 cm, and so the self-shielding begins to become less useful for a double beta decay search.

Since xenon is a noble gas, nearly all internal radioactive contaminants can be removed chemically. However isotopes of other noble elements like ^{85}Kr (a β^- emitter primarily produced as a byproduct to nuclear fission), ^{222}Rn , and ^{220}Rn (naturally occurring due to U and Th decays) are potential trace backgrounds. Krypton can be removed through distillation or ultracentrifugation. Since many decays occur along the full radon chains, there is some potential to tag radon daughter decays, especially making use of the α particle identification described above in section 3.3.3. Cosmogenic ^{135}Xe and ^{137}Xe can be mitigated by a rock overburden to provide shielding from cosmic rays, and are in any case short-lived.

Standard double beta decays will always be a background, of course, to searches for other modes. Precise measurements of the standard mode's half-life will make excesses due to nonstandard modes more significant. Good energy resolution, meanwhile, can improve the discrimination between the modes.

Chapter 4

The EXO-200 Detector

4.1 Creating a Sensitive Detector

The number of decays of a radioactive element in a given period of time follow a Poisson distribution, provided the half-life is much longer than the observation time. This is clearly the case for a double beta decay experiment. The expected signal in such a case is

$$N_s = \epsilon \frac{aM}{m_a} \frac{t \ln 2}{T_{1/2}} \quad (4.1)$$

events, where ϵ is the efficiency for detection, a is the isotopic abundance, M is the total mass of the element, m_a is the mass of a single atom, t is the total observation time, and $T_{1/2}$ is the half-life.

If an experiment observes no decays, then it sets a lower limit on the half-life substituting for N_s a number that corresponds to a desired confidence level for the Poisson distribution. The sensitivity to the half-life, then, goes like

$$S(T_{1/2}) \propto \epsilon t \frac{aM}{m_a} \quad (4.2)$$

Now suppose that an experiment has some number N_b of background events in the region of interest. As an approximation, assume that the rate of background events b (in units of counts per unit energy per unit exposure) is flat over an energy-based region of interest, whose width Γ increases with the detector resolution. The number

of background events then increases with increasing exposure Mt , or $N_b \propto bMt\Gamma$. The confidence limit on N_s must take into account the Poisson counting uncertainty on $(N_s + N_b)$, which will be dominated by the uncertainty on N_b (denoted ΔN_b), and so the sensitivity instead goes like

$$S(T_{1/2}) \propto t\epsilon \frac{aM}{m_a} \frac{1}{\Delta N_b} \approx t\epsilon \frac{aM}{m_a} \frac{1}{\sqrt{N_b}} \propto \epsilon \frac{a}{m_a} \sqrt{\frac{Mt}{b\Gamma}} \quad (4.3)$$

where the approximation that the uncertainty $\Delta N_b = \sqrt{N_b}$ is valid when N_b is large enough that Gaussian statistics apply. Poisson statistics on N_b create a transition between eqs. (4.2) and (4.3).

The design of a detector for a double beta decay experiment is guided by eqs. (4.2) and (4.3). It must be able to contain a large mass M of material highly-enriched in the isotope of interest (large a). It must also have a good energy resolution (small Γ) and be constructed to minimize backgrounds (small b). In a search for $0\nu\beta\beta$, recalling eq. (2.11), an experiment with large backgrounds will decrease the smallest $\langle m_\nu \rangle$ it is sensitive to as $t^{-1/4}$, which is much slower than an experiment with small backgrounds, which goes more like $t^{-1/2}$. EXO-200 is designed to meet these goals to achieve good sensitivity.

4.2 Time Projection Chamber

The EXO-200 detector uses liquid xenon as both the source and the detector of double beta decays. It takes the form of a Time Projection Chamber (TPC) that collects both the ionization and scintillation signals so they can be combined to improve energy resolution as discussed in chapter 3. The TPC also collects information about event topology, which is useful for background rejection. Figure 4.1 illustrates the TPC concept.

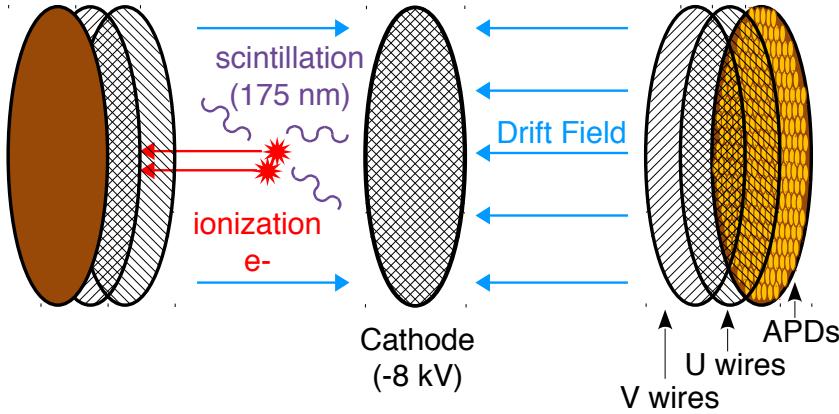


Figure 4.1: The concept for EXO-200. The TPC is divided into two identical halves by a cathode. When energy is deposited in the liquid xenon, it creates both scintillation light and ionization. APDs behind the anode wire planes collect the scintillation signal. Ionization drifts to the anode and is read out on two crossed wire planes.

4.2.1 Scintillation Readout

The scintillation light is collected by Large Area Avalanche PhotoDiodes (LAAPDs, or APDs). APDs can be manufactured with significantly lower radioactivity than traditional photomultiplier tubes. They operate well at cryogenic temperatures, and have a higher quantum efficiency than PMTs at 178 nm. However, APDs do suffer from increased noise and reduced gain compared to PMTs, but this is less of a concern at cryogenic temperatures and for the relatively high energies associated with double beta decay.

The 468 APDs are housed in two identical platters on either end of the TPC. The APDs are electrically connected to the platters, which provide ~ 1.4 kV bias. The platters are plated with Al and MgF_2 for improved reflectivity. PTFE tiles line the walls of the TPC and reflect light to improve the collection efficiency. One device is missing from each platter in order to make room for a diffuser that can deliver light from an external laser through an optical fiber for calibration purposes.

The APDs are electrically grouped together in gangs of 5 to 7. A phosphor bronze “spider” provides a common electrical connection to a gang, and mechanically holds them in the platter. Individual APDs vary significantly in the voltage needed to achieve a desired product of quantum efficiency and gain, and so the gangs are chosen

in order to match electrically-similar devices [42]. A trim voltage is then applied so that all gangs have similar performance.

4.2.2 Ionization Readout

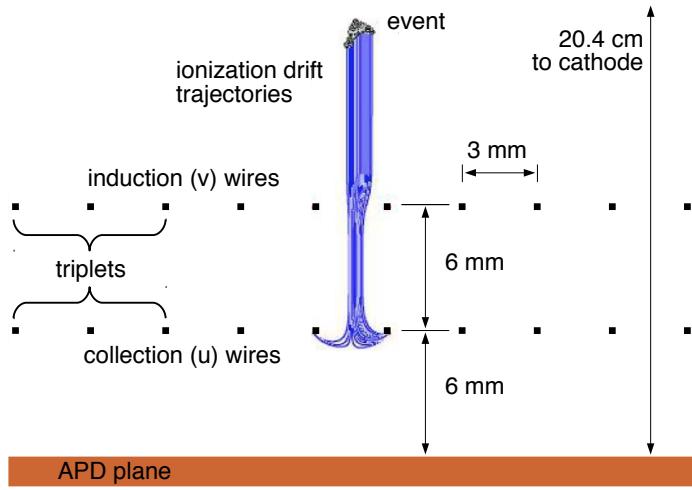


Figure 4.2: The geometry of the ionization readout. An ionization cloud drifts past the v wires, inducing a signal in them. The cloud is then collected on the u wires. For simplicity, the wires are shown collinear here. In reality, they are angled 60° from each other.

A central cathode divides EXO-200 into two drift regions. This cathode creates a drift field that drifts ionization electrons to the anodes at either end of the TPC. The anodes consist of two crossed wire planes, angled 60° to each other. The plane farther from the cathode, denoted u , is held at virtual ground and collects the ionization. The plane closer to the cathode is denoted v and is biased to be fully transparent to drifting ionization. The v wire plane shields the u wire plane from induction effects. The induced signal on the v wires also provides a transverse coordinate. This, along with knowledge of which u wire collected the charge, provides a location for the event in the transverse plane. It is possible to reject backgrounds by using the multiplicity of signals (see section 5.2.4) and location information to determine a fiducial volume (see section 8.1.3).

Each wire plane consists of 114 wires, spaced 3 mm apart. These wires were formed

as triplets by photoetching sheets of phosphor bronze. This gives 38 electrical readouts per plane. Each wire plane is 95.8% transparent to normally incident scintillation light. The cathode is also formed from two pieces of etched phosphor bronze and is 90% transparent to normally incident light.

The wire planes are spaced 6 mm apart, and the collection plane is 6 mm from the face of the platter that holds the APDs. The cathode, biased to -8 kV , is 19.2 cm from the induction wire plane. The anode, cathode and cylindrical field shaping rings along the sides of the drift region create a 374 V cm^{-1} field in the main drift region. A chain of resistors divides the voltage along the 10 field shaping rings in each drift region so that the drift field is uniform. Charge in this region drifts at $1.71\text{ mm }\mu\text{s}^{-1}$. The field is stronger between the u and v wires, and charge drifts about $2.2\text{ mm }\mu\text{s}^{-1}$ in this region.

4.2.3 Construction



Figure 4.3: On the left, a cutaway view of the EXO-200 detector and liquid xenon vessel. On the right is a photo of EXO-200 under construction. The circular APDs are visible in their platter and the flex cables can be seen around the edges. The anode wire planes are above the APD plane and cross each other at 60° . PTFE tiles line the inside of the field-shaping rings. The cathode grid has not been installed in this photo, but would be at the top.

EXO-200 is made out of low-radioactivity copper, and was constructed under

a modest concrete overburden to limit cosmogenic activation. The walls are only 1.37 mm thick to further reduce backgrounds due to radioactivity in the copper. All materials used in and around the detector have been thoroughly screened for radioactivity [43].

The electrical cables for the anode wires and the APDs reach the detector through six ~ 1 m legs, which also support the detector horizontally from the cryostat door. Long cables allow the TPC to be shielded from any radioactivity in the electronics, at the expense of some electronic noise. These legs also allow liquid xenon to circulate through the detector. The high voltage reaches the cathode through a separate connection, located in the middle of the cylinder.

4.3 Calibration

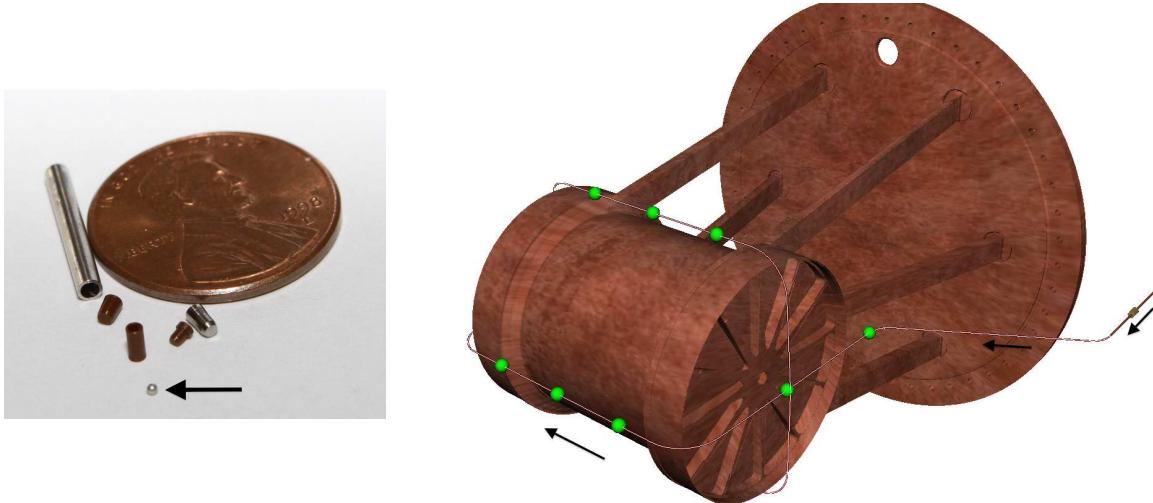


Figure 4.4: To calibrate the detector, a tiny source capsule (left) containing a radioisotope housed in the small sphere indicated by the arrow can be pushed with a cable through a guide tube system to many positions just outside the detector (right).

Routine calibrations help monitor the performance and characterize the energy response of EXO-200. A guide tube system runs just outside the TPC, providing access to both anodes and three positions around the circumference of the cathode. In

a typical calibration run, a miniature radioactive source is pushed to one of these positions by a long cable. After enough time passes to accumulate sufficient calibration data, the source is retracted and normal physics operations resume.

Currently, EXO-200 is calibrated with three different isotopes. ^{137}Cs emits a 662 keV gamma ray when it decays. ^{60}Co decays emit both 1173 keV and 1333 keV gamma rays simultaneously. ^{228}Th eventually decays to ^{208}Tl , which emits a 2615 keV gamma ray when it decays. Typically, short (2 hr) ^{228}Th calibrations are taken a few times per week, chiefly to measure the electron lifetime (see chapter 6) and to track any other time variation in the detector response. Several times per year, longer campaigns with the full suite of sources establish the energy scale of the detector (see section 5.4).

4.4 Infrastructure

4.4.1 Cryostat and Clean Rooms

As shown in fig. 4.5, the TPC vessel that contains the liquid xenon is housed inside a double-walled cryostat. The inner chamber is filled with HFE-7000 [44] fluid that acts a thermal bath. The HFE fluid is radiologically clean and additionally provides some shielding from external γ radiation. The outer chamber is held at 10^{-6} Torr for insulation. The cryostat, made of copper low in radioactivity, is in turn shielded by 25 cm of lead on all sides. The front end electronics lie just outside this lead wall. On the front of the cryostat, a second 25 cm wall sits just outside the electronics. This second wall prevents external radiation from having line-of-site to the detector through the penetrations in the inner lead wall, used for xenon and electrical connections.

A class 1000 clean room houses the cryostat and its shielding in order to reduce trace radioactivity from dust being tracked in and introduced to the detector. This clean room also houses the xenon handling system, the cryogenic refrigerators used to chill the HFE, and other necessary infrastructure.

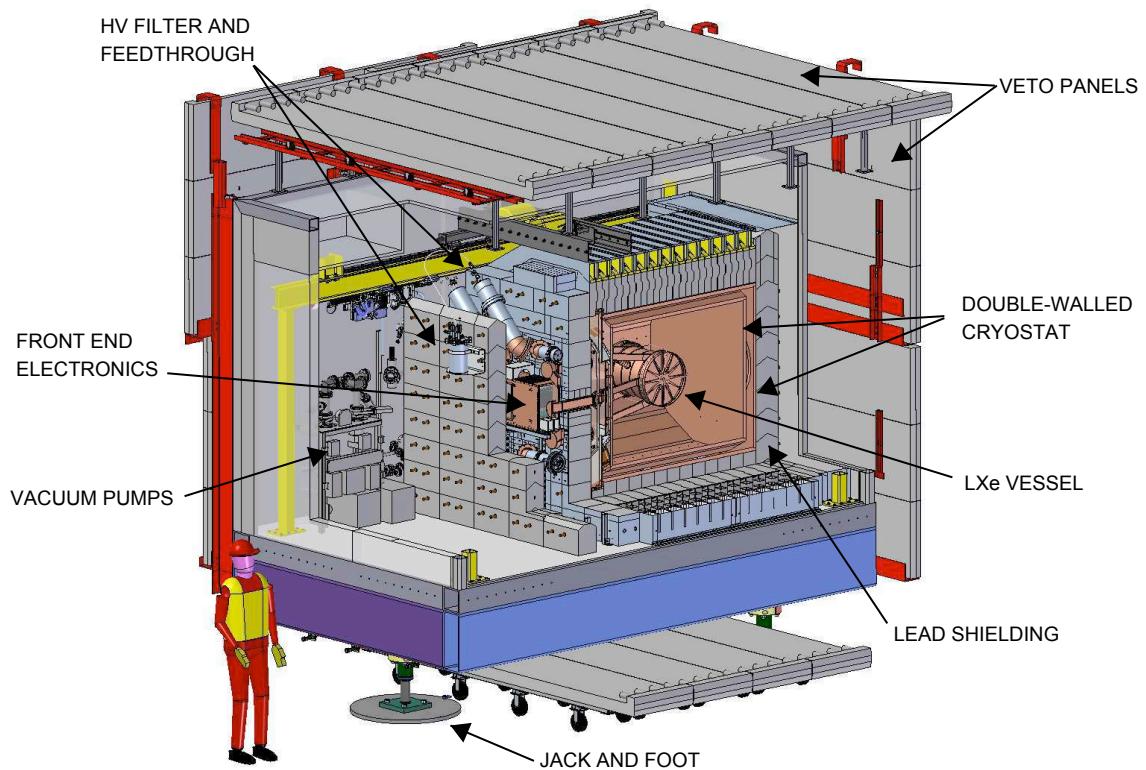


Figure 4.5: A cutaway view of EXO-200 as installed at WIPP. The TPC vessel is surrounded by cold HFE. A cryostat insulates the HFE and liquid xenon. Lead surrounds the cryostat to shield from external radiation. All of this is located inside a clean room, which is surrounded with an active muon veto.

4.4.2 Muon Veto

While muons passing through the TPC can be readily identified (see chapter 7), muons passing near the detector may still cause background events due to radiative processes and gamma rays from nuclear processes due to spallation neutrons. In order to reduce these backgrounds, the clean room housing the detector is surrounded by an active muon veto as shown in fig. 4.5. Thirty-one plastic scintillator panels cover four sides of the clean room and detect $(96.0 \pm 0.5)\%$ of muons traversing the TPC.

4.4.3 WIPP

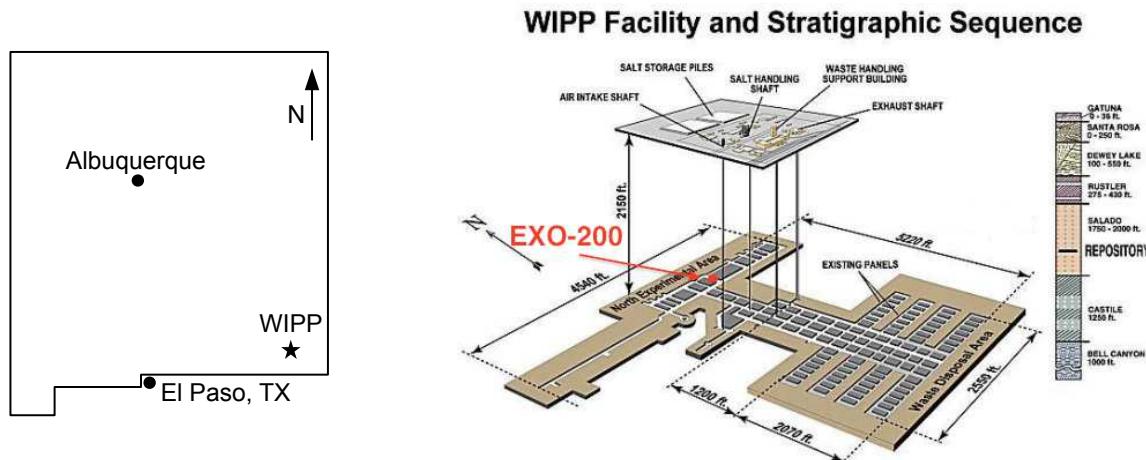


Figure 4.6: The left shows the location of the WIPP site on a map of New Mexico. The right shows a detailed view of the WIPP site. EXO-200 is located in the North Experimental Area, approximately 655 m underground. (Right image courtesy of the US Department of Energy.)

In order to attenuate the cosmic ray muons, EXO-200 is located \sim 655 m underground at the Department of Energy's Waste Isolation Pilot Plant (WIPP) in southeastern New Mexico. WIPP is a salt mine, and its primary purpose is the permanent disposal of transuranic waste. The north end of the mine, far from and upwind of the waste, serves as a suitable site for low-background experiments.

The rock overburden provides 1480 hg cm^{-2} shielding from cosmic rays (see chapter 7), and the salt walls are naturally low in radionuclides. Direct counting finds

a contamination of $(27 \pm 2) \times 10^{-9}$ g/g of ^{238}U , $(66 \pm 2) \times 10^{-9}$ g/g of ^{232}Th , and $(124 \pm 2) \times 10^{-9}$ g/g of ^{40}K [45]. Little radon emanates from the walls; the radon concentration in the air is 7 Bq m^{-3} , similar to that found at the surface.

4.5 Xenon and Recirculation

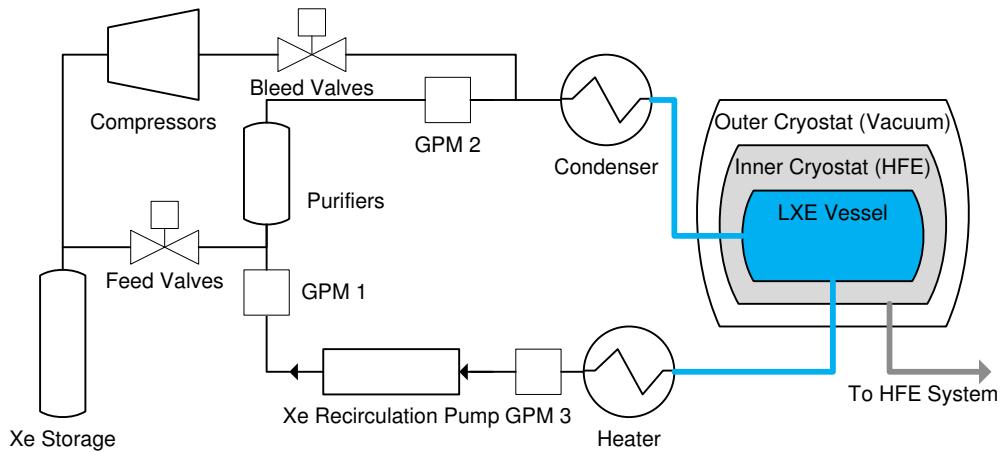


Figure 4.7: A simplified schematic of the EXO-200 xenon recirculation system. A heater vaporizes xenon returning from the TPC. A gas-phase pump recirculates xenon through purifiers. The xenon is condensed and then returns to the TPC. Gas purity monitors (GPMs) look for changes in the electron lifetime of the circulating xenon. The TPC vessel must track the HFE pressure closely. A sophisticated slow controls system controls valves to feed and bleed xenon as needed from storage bottles. Compressors allow the xenon to be placed back into the bottles when bleeding xenon.

As the name suggests, EXO-200 makes use of 200 kg of xenon enriched to 80.6% in ^{136}Xe . Of the remaining fraction of the xenon, isotope 134 comprises 19.1%, and lighter natural isotopes are present in trace amounts. Roughly 175 kg of xenon fills the 58 l vessel, with 110 kg of that in the active region bounded by the anode wires and PTFE reflectors. The remaining xenon is in the recirculation and storage system.

The HFE is cooled to 168 K by cryogenic refrigerators through heat exchangers in the cryostat. Because the TPC vessel has very thin walls, the pressure inside must closely match the pressure of the HFE surrounding it. To keep the xenon in the liquid phase, the HFE is held at ~ 140 kPa absolute pressure. A sophisticated slow control

system can feed xenon gas into the system or bleed it off as needed to match this pressure. The xenon plumbing consists of a high pressure side where gas is stored, and a low pressure recirculation loop. Compressors place gas back into the storage bottles when bleeding. Under normal operations, feeds and bleeds do not happen. Instead, xenon returns from the TPC and vaporizes in a heater. A gas-phase pump [46] recirculates the xenon through purifiers, after which the xenon is recondensed. This is shown in fig. 4.7.

The purifiers are heated Zr getters [47] that chemically remove electronegative impurities from the xenon so that ionization can drift in the TPC. Gas purity monitors [48] measure the quality of the gas returning from the TPC to look for large quantities of impurities, which might indicate purifier failure or a leak. However, in normal operation, continuous recirculation leads to electron lifetimes of > 2 ms, which must be measured with the TPC itself (see chapter 6).

4.6 Summary

As this chapter and the even more thorough description in [45] detail, EXO-200 has been designed to be a sensitive double beta decay experiment. A large sample of xenon enriched in isotope 136 provides good exposure to signal events. Low background materials have been used throughout the construction, and layers of shielding provide additional protection. The detector itself is located deep underground to protect from cosmogenic activity. The TPC design allows for further background rejection through position-based fiducialization and by identifying Compton scattering gamma rays. The ionization and scintillation signals are combined to yield good energy resolution. All of these ensure a low background rate near the double beta decay endpoint. This in turn means good sensitivity to the double beta decay half-life, and that the sensitivity should scale more like eq. (4.2) than eq. (4.3).

Chapter 5

Data Collection and Processing

5.1 Signal Readout

Figure 5.1 shows a simplified schematic for the EXO-200 electronics. The scintillation signals from the APDs and the induction and collection signals from ionization are read out through long flex cables. The long cables separate the electronics from the TPC, removing the need for cryogenic and low-radioactivity components. After passing through a preamplifier, the signals are fed through shaping circuits and are then digitized at a rate of 1 MHz. The shaping circuits consist of a differentiator on the preamplifier, followed by two identical CR-RC stages of differentiation followed by integration. Table 5.1 summarizes the time constants. The transfer functions of

Table 5.1: The shaping times for the signal readouts in EXO-200. The stages are listed in order. ‘D’ denotes a differentiation stage, while ‘I’ denotes an integration stage.

Channel Type	Stage Type				
	D (μs)	D (μs)	I (μs)	D (μs)	I (μs)
APDs	300	10	3	10	3
<i>u</i> wires	measured (nominally 60)	40	1.5	40	1.5
<i>v</i> wires	60	10	3	10	3

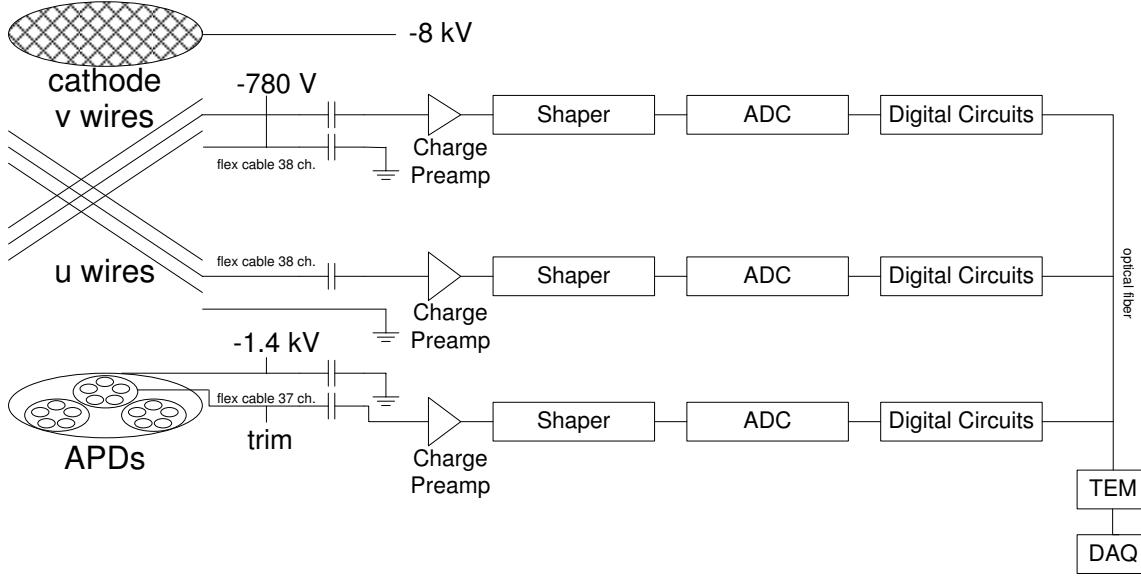


Figure 5.1: A simplified schematic of the EXO-200 electronics. The ionization and scintillation signals are read out through long flex cables. The signals are shaped, digitized, and passed to the Trigger Event Module (TEM), which passes the data to the DAQ computers for recording when signals are detected.

these shapers determine the waveform of a physical signal (ionization or scintillation). In order to improve energy resolution, the preamp stage differentiation times for the *u* wire channels are measured by analyzing the response to charge injected from a capacitor.

After digitization, the Trigger Event Module (TEM) monitors the signals to select interesting events. Due to the low radioactive backgrounds in EXO-200, and the relatively slow rate of $2\nu\beta\beta$, this trigger is extremely permissive. The triggering thresholds have not been thoroughly studied at low energies, but the trigger is 100% efficient at triggering on events that deposit more than 700 keV in the liquid xenon. The average trigger rate during routine data taking is approximately 0.8 Hz. Additionally, the trigger is forced to fire every 10 s in order to ensure the DAQ system is correctly recording events and to provide a measurement of detector live time. When the trigger fires, the DAQ records a frame consisting of digitized waveforms for all channels $\pm 1024 \mu\text{s}$ around the trigger time.

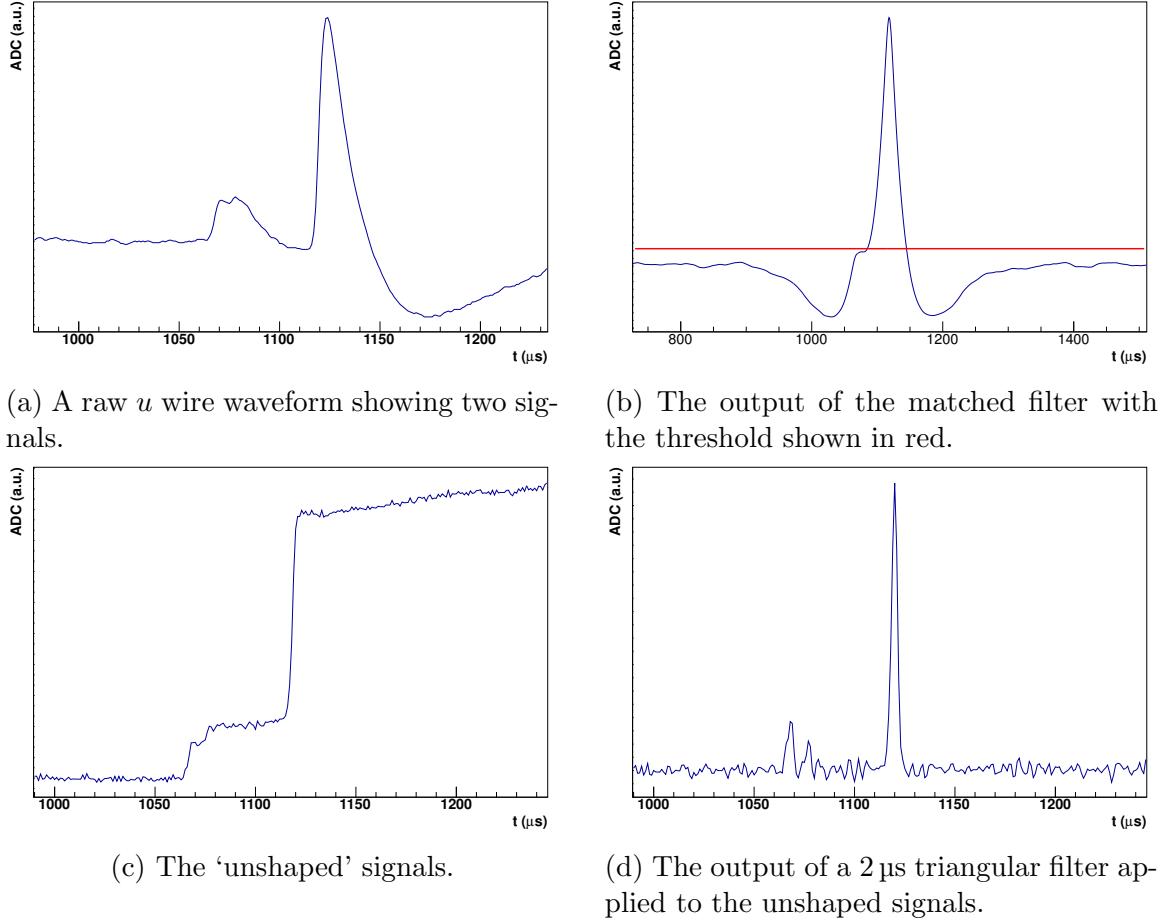


Figure 5.2: An example of finding signals on a u wire waveform.

5.2 Signal Reconstruction

5.2.1 Signal Finding

Either scintillation or ionization signals may cause the trigger to fire. Since there is not a fixed time between the two types of signal, it is necessary to search for signals in the waveforms. This is done in a two stage process: first a matched filter technique searches for signals, and then a waveform unshaping technique refines the found signals. Figure 5.2 shows an example of this process.

The matched filter technique [49] convolves a time-reversed template signal with the waveform. Template signals are produced by passing a step function through

the filters specified in table 5.1 and adding noise sampled from the detector. This is done for individual channels for the u and v wires. For the APDs, the signals on the individual planes are summed and this technique is applied to these two sum signals. A threshold is determined by calculating the mean absolute deviation (MAD) of the waveform from its baseline. Parts of the waveform exceeding $(3\sqrt{\pi/2}) \times \text{MAD}$ are excluded and the MAD is recalculated. A signal is found if the filter output exceeds 5 (4) times this final MAD for the wire (APD) signals. Figures 5.2a and 5.2b show an example of this process.

While the matched filter is good at identifying that there was a signal on a wire, a refined analysis is needed to identify if multiple signals occurred in a short period of time. As fig. 5.2b shows, the matched filter only finds one signal when two are clearly visible. To find other signals, a $256\text{ }\mu\text{s}$ portion of the raw waveform is unshaped (by convolving the signal with the inverse transfer function of the shapers), then reshaped with a $2\text{ }\mu\text{s}$ triangular filter. Peaks in the output correspond to distinct signals. Finding multiple signals aids in discriminating against Compton-scattering γ ray backgrounds.

5.2.2 Signal Parameter Extraction

Once signals have been identified in a waveform, information about their energies and times are extracted by fitting template waveforms (determined by the shaping times in table 5.1) to the waveforms. This fit is done by minimizing the χ^2 statistic formed by comparing the baseline-subtracted waveform to the template waveform, allowing the template's amplitude and time to float. Signals that fit to similar times are combined. Signals that fit to small amplitudes or have large errors on the fit amplitude are removed. Minimization is repeated until no more signals need to be combined or removed. This is done on the individual u and v wire channels, and on the two plane sum APD signals.

Signals that are collected on a u wire channel can induce signals in neighboring channels. These induction signals could be mistaken for collection signals. Induction signals are identified and removed by combining information from a number of metrics:

- The rise times and maximum-to-minimum times of the pulses in the waveforms are shorter for induction signals than collection signals.
- The integrals of the unshaped waveforms are much larger for collection signals than induction signals.
- Induction signals have a better χ^2 statistic when fit to a template induction signal than to a template collection signal.
- Induction signals occur on channels that neighbor channels with very large collection signals (associated with events of $\gtrsim 1$ MeV).

5.2.3 Signal Clustering

Once signals have been identified, signals of the same type may be bundled together:

- u wire signals: signals that are on neighboring channels and are within 3.5 μ s of each other most likely came from the same energy deposition, and so are bundled together. This bundle is assigned the energy-weighted mean time and u position of the constituent signals.
- v wire signals: induction effects from a single energy deposition may affect more than the nearest-neighbor channels. v wire signals are bundled together if $|(t_i + (2.97 \mu\text{s})\Delta v) - t_0| < 4.5 \mu\text{s}$, where t_0 is the time of the largest amplitude signal, t_i is the time of signal i , and Δv is the difference in channel number between the largest signal and signal i . This is based on an empirical measurement that channels farther away from the event fit to earlier signal times. The resulting bundle is assigned the energy-weighted mean v position of the bundled signals and the time of the largest signal.
- The sum APD plane signals are bundled together if they have times within 6 μs (based on the integration times for the shaping circuits). This bundle is assigned the energy-weighted mean time of the signals.

The z position of each u wire bundle is calculated using the time of the most recent scintillation bundle and the known drift velocity for ionization. If no scintillation bundle is found within the maximum physically possible drift time, the bundle is not assigned a z position. Later data quality cuts remove events that have more than one scintillation bundle within this time window, since the z position becomes ambiguous.

Finally, u signal bundles are paired with v signal bundles to form a “cluster”. For each possible pairing, a likelihood is computed using 3 PDFs:

- After correcting for the known drift velocity between the anodes, the distribution of two paired u and v bundle times is a Gaussian centered around 0, with the standard deviation determined by the timing resolution of 1 μ s.
- The relationship between u and v signal magnitudes has been studied. Once the magnitudes of the bundled signals have been appropriately scaled, the distribution of their differences is again Gaussian, smeared by the energy resolution.
- Not all u and v wire channels cross. If a pairing would cause the signal to be outside the hexagon containing the wires, the likelihood is penalized. The perpendicular distance to the edge of this hexagon is calculated. The penalty is the likelihood to be that distance away from the mean of a Gaussian distribution with standard deviation 4.5 mm (half the wire pitch).

The set of pairings that maximizes the likelihood is used. The paired u and v bundles provide a 2D transverse position for the event. Since the wires do not cross at right angles, this position is translated to a (x, y) position. The $+y$ direction is vertically upward, the $+x$ direction is away from the cryostat door, and $+z$ is along the TPC axis and chosen to make the coordinate system right-handed. The position $(x, y, z) = (0, 0, 0)$ is defined to be the center of the detector, which is the center of the cathode. Occasionally, a signal bundle may remain unpaired if no good pairing is found. These still have some well-defined quantities like energy and z position that are saved, but events with such unpaired bundles are typically not used in analyses.

5.2.4 Topology

An event may have zero, one, or many clusters. This proves a useful tool for discriminating beta decays from gamma ray backgrounds. The electron(s) from a beta decay deposit their energy in a small volume, creating mostly events with one cluster. A small fraction create bremsstrahlung photons that create multiple clusters. Meanwhile, gamma ray backgrounds tend to Compton scatter, producing events with multiple clusters. Some, however, do photoionize the xenon or Compton scatter in a small volume, producing only one cluster. When a gamma ray creates an electron-positron pair, this may only produce one cluster if the 511 keV gamma rays from the positron annihilation escape the detector.

Two categories of event are created:

- Single Site (SS) events: events with only one cluster. Furthermore, this cluster must have only occurred on two or fewer adjacent u wires. This selects most beta decay events and few gamma ray events.
- Multiple Site (MS) events: events that are not single site. This selects most gamma ray events, with only a small number of beta decay events.

The clustering algorithm above is limited by the position resolution of the detector. For the u wires, the resolution is dominated by the 9 mm pitch. For the v wires, the magnitude of the induction signal on a channel is correlated with the signal's distance from that channel. Since induction signals occur on many neighboring channels, the 3 mm resolution in v is better than for u . The electronic noise, sampling frequency, and shaping times determine the resolution in z . With the multiple signal finder described above, the resolution is about 3 μ s, which corresponds to about 5 mm.

5.3 Corrections

In order to have good energy resolution, events of the same energy should be reconstructed with the same energy, regardless of where in the detector they originated or on which channels the signals were collected. In order to achieve this, a number

of corrections are applied to correct for both gain and position variation. They are detailed below.

5.3.1 Wire Gain Correction

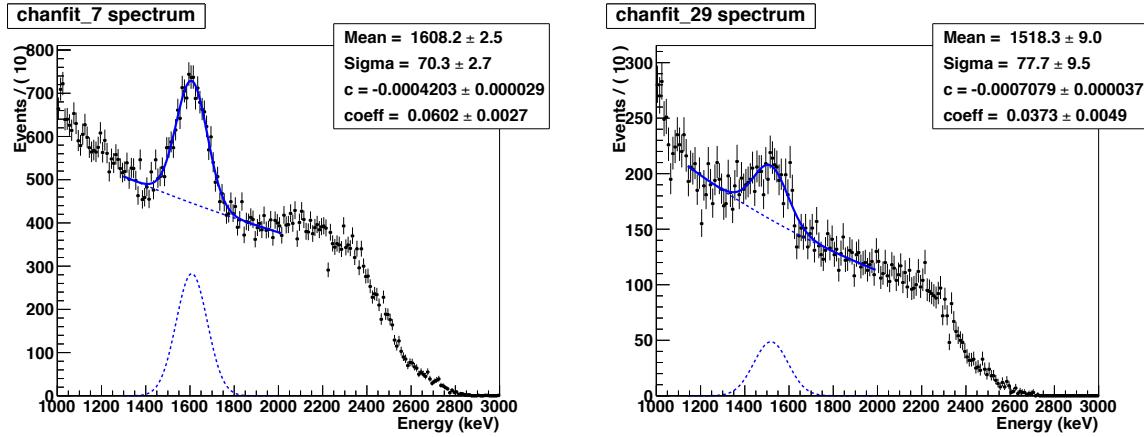


Figure 5.3: The 1593 keV pair-production peak shown on two different, uncorrected u wire channels. A Gaussian peak with an exponential background matches the observed spectrum well, and the mean obtained from the fit is then used to correct for channel-to-channel variations.

Minute differences in the components that read out each u wire channel can cause channel-to-channel variation in their responses. In order to address this, data from ^{228}Th calibration runs is used. The 2615 keV γ ray can produce an electron-positron pair with combined energy 1593 keV. Since the electron and positron deposit their energy in a small volume, this peak is clearly visible on individual wire channels. A scaling factor for each channel is determined by fitting these peaks and scaling so that they occur at the same value for each channel. Figure 5.3 shows examples for two channels.

Environmental conditions can also cause the response to shift over time. To address this, each channel has a capacitor that can inject a known amount of charge. These charge injections are done daily and used to measure the shaping time for the charge preamp (see section 5.1), but can also be used to measure the time variation of each individual channel. The time variation is tracked, and applied as a very small

correction on top of the channel-to-channel correction.

5.3.2 Shielding Grid Correction

When ionization is created in a TPC, the positively charged ions move more slowly than the much less massive electrons. These positive ions can induce a signal on the collection wires that is opposite to the collection signal. The measured signal will be reduced. For events far from the anode, the induction wire grid shields the collection wires from this effect. However, closer to the anode, the shielding grid is not perfect [50], and so this effect must be corrected for.

Monte Carlo simulations using a Maxwell [51] simulation of the electric field inside the TPC suggest that the measured ionization signal E_{measured} is related to the true ionization signal E_{true} by

$$E_{\text{measured}} = E_{\text{true}} \left[1 + p_0 \exp \left(\frac{|z| - 192.5\text{mm}}{p_1} \right) \right]^{-1} \quad (5.1)$$

where 192.5 mm is the distance between the cathode and the v wire plane. The constants p_0 and p_1 are determined by fitting to simulated data, as shown in fig. 5.4. They are found to differ for events that have ionization collected on one u wire and for events that share ionization across 2 u wires. The best-fit values are summarized in table 5.2. The results of applying these corrections are shown in fig. 5.5. There is a lingering $\pm 0.35\%$ variation in the collected signal through the TPC.

Table 5.2: The parameters for the shielding grid correction (eq. (5.1)) obtained from Monte Carlo simulations. The parameters are observed to vary when the signal is collected on a single u wire and when it is shared over two u wires. For comparison, the parameters when 1- and 2- wire signals are not separated are also given.

number of u wires	p_0	p_1
1	0.043	7.02
2	0.064	8.49
1 and 2	0.043	6.60

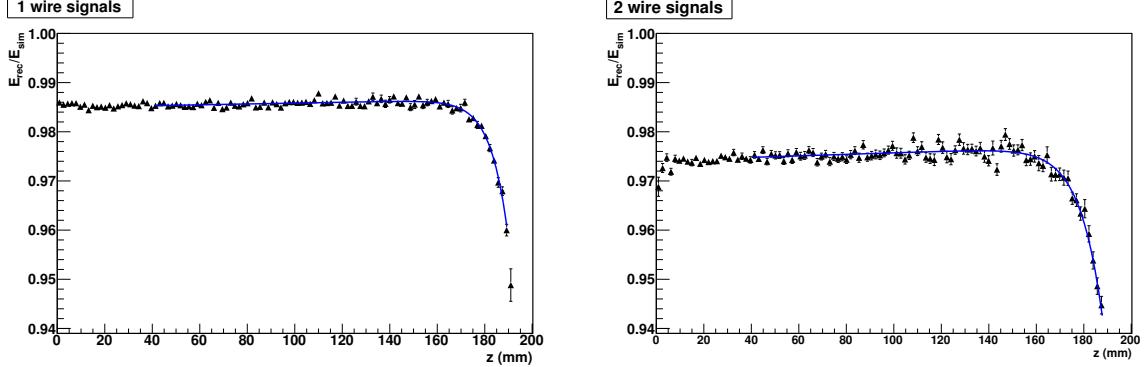


Figure 5.4: Fits of the shielding grid correction function (eq. (5.1)) to simulated ^{228}Th data. Each point represents the ratio of the reconstructed to the simulated full-absorption peak energy for events in a particular z bin. The left shows the ratio for signals collected on 1 u wire, while the right shows the ratio for signals collected on 2 u wires.

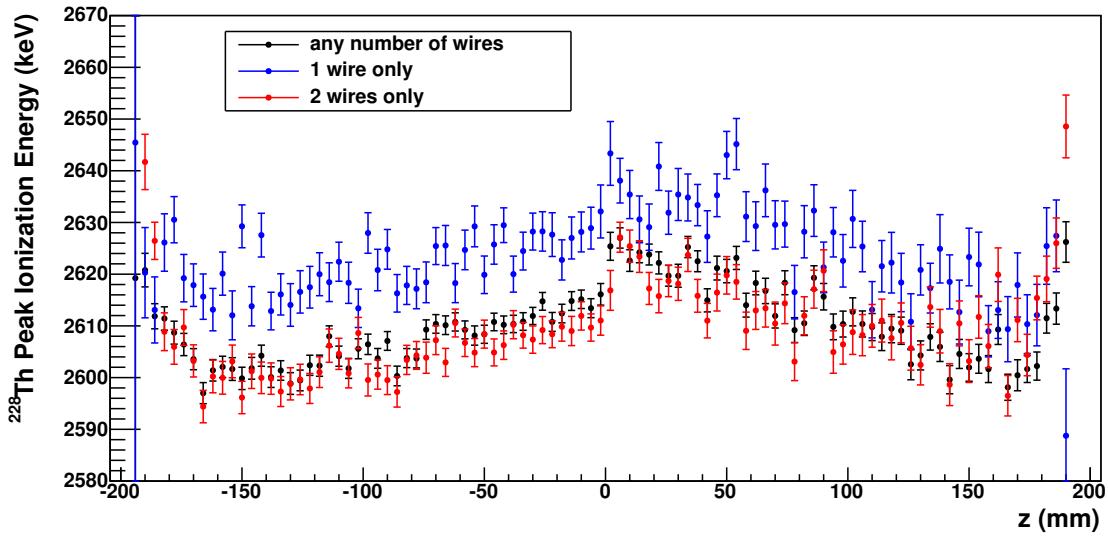


Figure 5.5: The results of applying the shielding grid correction (eq. (5.1)) with the parameters in table 5.2 to data collected with a ^{228}Th source positioned at the cathode. Each point represents the center of the full-absorption peak for events in a particular z bin. A $\pm 0.35\%$ variation is observed throughout the TPC.

5.3.3 Electron Lifetime Correction

Ionization drifting in the TPC is attenuated by attachment to electronegative impurities. For uniformly distributed impurities, this results in an exponentially decreasing signal as a function of drift time, and the corresponding time constant is the “electron lifetime”. This effect can be measured by measuring the attenuation of a known signal (as during a radioactive source calibration). This measurement is then used to correct unknown signals. Chapter 6 details the measurement of the electron lifetime and correction.

5.3.4 Light Collection Correction

While trim voltages are applied to the APDs to try to achieve a uniform response, there is still some variation in the gain from channel to channel. Additionally, there is geometric variation in the amount of scintillation light actually reaching the APDs due to shadowing from the wire planes, imperfect reflection off of the walls, and other effects. Using radioactive source calibration runs, the overall response from both of these factors can be mapped out and corrected. Appendix B details this correction.

5.4 Energy Calibration

The calibration system described in section 4.3 is used to calibrate the energy scale of the detector. This is done separately for single site and multiple site events.

5.4.1 Fit Model

For gamma ray lines, a Gaussian distribution representing the full-absorption peak is combined with a complimentary error function representing Compton scatter events that escape the detector. Explicitly, the model is:

$$\frac{dN}{dE} = a \frac{1}{\sigma \sqrt{2\pi}} \exp\left(\frac{-(E - E_0)^2}{2\sigma^2}\right) + (1 - a) \operatorname{erfc}\left(\frac{E - E_0}{\sqrt{2}\sigma}\right) \quad (5.2)$$

where the peak energy E_0 , width σ , and peak fraction a are all allowed to float.

Table 5.3: The bias for fits using eq. (5.2) obtained from Monte Carlo simulations of the sources. It is defined as $E^{\text{bias}} = (E^{\text{true}} - E^{\text{fit}})$. It differs for single site (SS) and multiple site (MS) events.

$E^{\text{bias}} = (E^{\text{true}} - E^{\text{fit}})$ (keV)			
	^{137}Cs (at 662 keV)	^{60}Co (at 1173 keV)	^{228}Th (at 2615 keV)
SS	3.19	3.19	2.50
MS	8.13	6.19	6.35

This model is fit to spectra collected in ^{137}Cs and ^{228}Th calibrations, and a sum of two such models is fit to data from ^{60}Co calibrations, since it emits two gamma rays similar in energy. The energy spectra are fit using an unbinned maximum likelihood method. This fit is performed twice, first over a broad energy range to roughly identify the peak, then again in the range $(-3.0\sigma, +2.5\sigma)$ from the found peak to more precisely determine the mean energy, and to minimize bias introduced by the choice of fit window.

This model exhibits a slight bias. That is, the mean energy found is different than the true value. This can be measured by applying the fit model to Monte Carlo simulations of the sources, where the true peak energy is known. The biases are summarized in table 5.3. The bias varies as a function of energy resolution, and with the peak-to-Compton ratio, which is why it is not the same in all cases.

5.4.2 Rotation Angle

As described in section 3.3.3, the ionization and scintillation signals can be combined by projection onto a rotated axis in order to achieve good energy resolution. In order to determine the optimal angle, the energy spectra for calibration runs are constructed using a range of different angles using

$$E_R = E_I \cos \theta_R + E_S \sin \theta_R \quad (5.3)$$

where E_I and E_S are the ionization and scintillation signals after all corrections, but no calibrations, have been applied. The spectra are fit with the model in eq. (5.2), and the angle that optimizes the resolution σ/E_0 is selected. Figure 5.6 shows an example of this process. Because E_I and E_S are uncalibrated, the value of θ_R determined in this manner has no physical meaning.

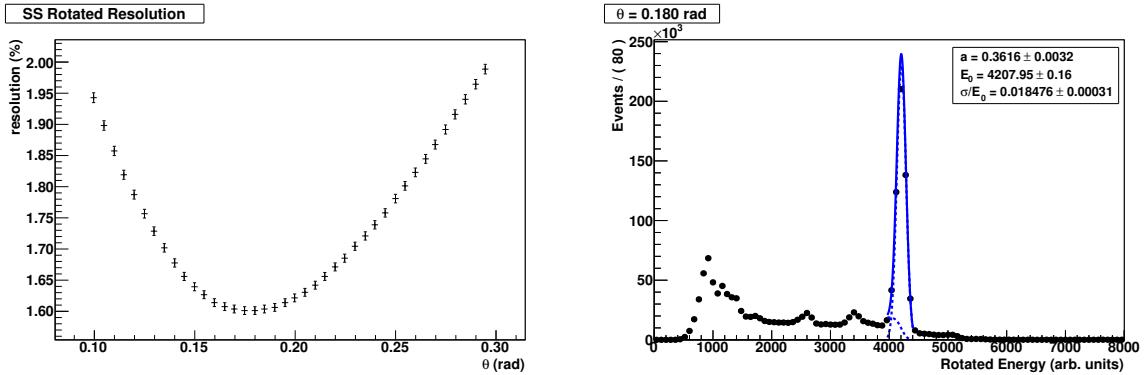


Figure 5.6: The rotation angle is scanned over a range of values (left) in order to find the angle that optimizes resolution. In this case, the optimal angle was $\theta = 0.18$. On the right, the fit of the model in eq. (5.2) to thorium calibration data with the optimal rotation applied (according to eq. (5.3)).

5.4.3 Thallium Peak from Thorium Runs

Thorium source calibrations are taken every 1–2 days. On a week-by-week basis, the optimal rotation angle for the 2615 keV full absorption peak from the decay of ^{208}Tl is determined from these calibrations. This angle is used for all data during that week. 2615 keV is close to the 2458 keV Q value for ^{136}Xe , and so optimizing the resolution for this peak should also give good resolution near the $0\nu\beta\beta$ peak. Figure 5.7 shows the optimal angle for the Run 2a time period analyzed in chapter 8. Some slight variation is seen. This is primarily due to variation in the electronic noise of the APDs. As the noise increases, better energy resolution can be obtained by using less of their signal, corresponding to smaller θ .

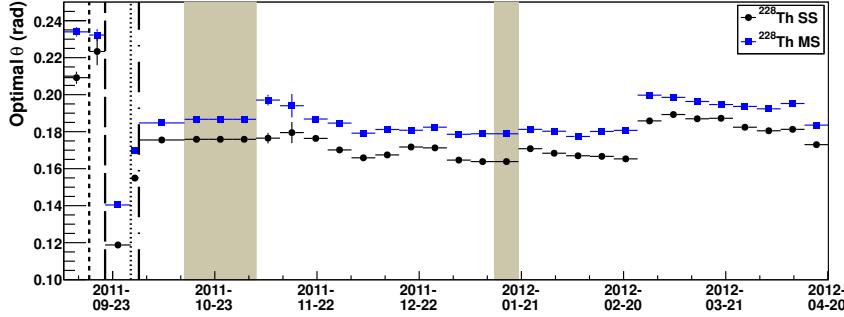


Figure 5.7: The angle that yields the optimal energy resolution for the 2615 keV full-absorption peak is determined on a week-by-week basis for ^{228}Th calibration runs (top). Shaded periods represent times in which the Th source was not used, and so the optimal angle is based on the combined data from the previous and next weeks that the Th source was used. The vertical lines represent electronics changes prior to the start of Run 2a.

5.4.4 Overall Energy Scale

While thorium source calibrations are taken frequently, the cobalt and cesium sources are generally reserved for dedicated campaigns. Throughout these campaigns, it has been observed that the ratio between the peak energies remain relatively stable, even though their values show some drift. This is shown in fig. 5.8. Therefore, the three ratios $E_{\text{Cs}}/E_{\text{Th}}$ and $E_{\text{Co},i}/E_{\text{Th}}$ are used as the basis for the calibration. The process for calibration is then:

1. On a week-by-week basis, use thorium calibration data to determine the optimal rotation angle $\theta(t)$ and to fit the peak $E_{\text{Th}}^{\text{fit}}(t)$ in the spectrum for thorium data obtained using eq. (5.3) and that rotation angle. If no thorium data is available for a given week, data from the adjacent weeks are combined and used to estimate the angle and peak for that week.
2. For each multiple isotope calibration campaign, compute the ratios $E_{\text{Cs}}^{\text{fit}}/E_{\text{Th}}^{\text{fit}}$ and $E_{\text{Co},i}^{\text{fit}}/E_{\text{Th}}^{\text{fit}}$
3. Using the mean values from these campaigns and the true ratios, fit a polynomial function f such that $f([E^{\text{fit}}/E_{\text{Th}}^{\text{fit}}]_{\text{data}}) = [E^{\text{fit}}/E_{\text{Th}}^{\text{fit}}]_{\text{simulation}}$. The ratios used for this fit are shown in table 5.4.

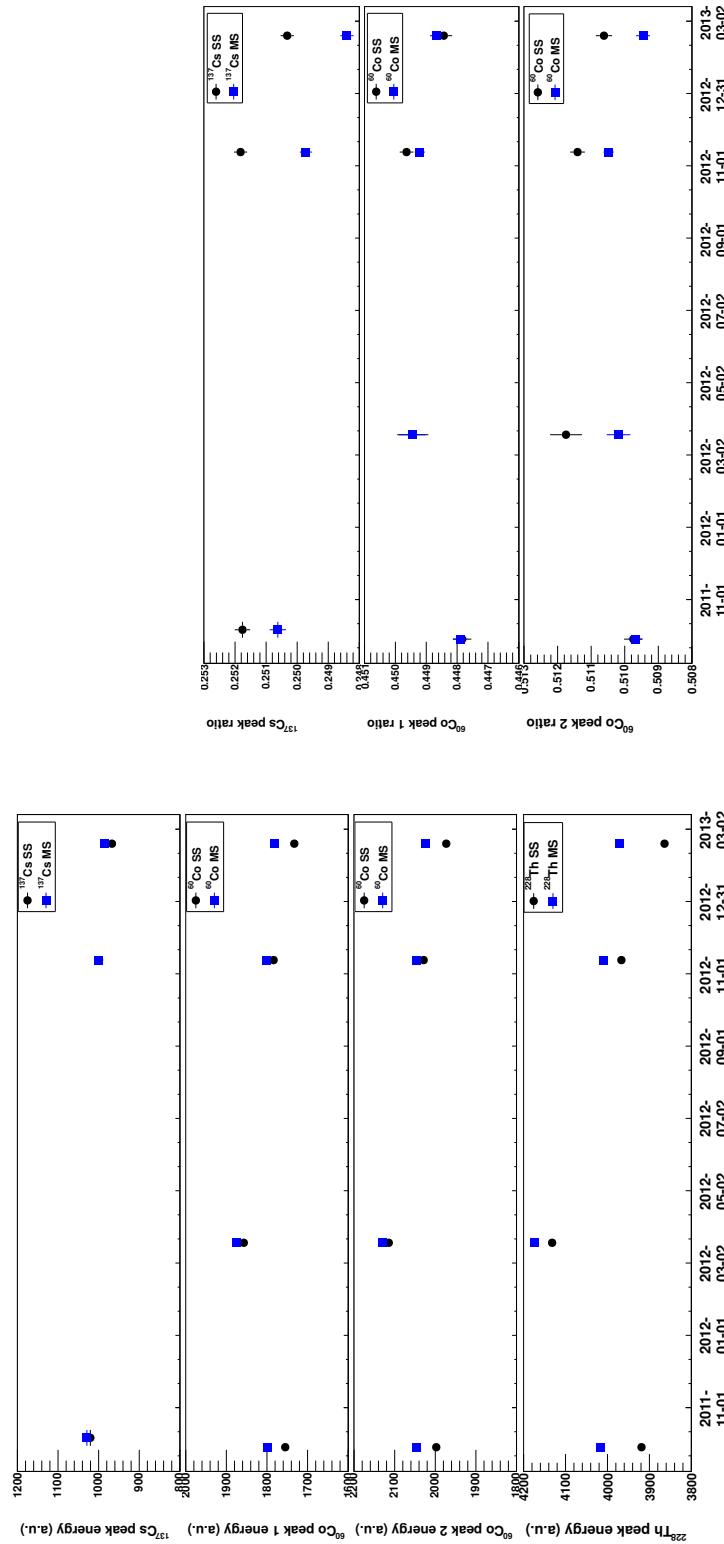


Figure 5.8: Although the rotated energy of the full-absorption peaks for ^{137}Cs , ^{60}Co , and ^{228}Th vary by about $\pm 3\%$ over time (left), the ratios between the full-absorption peaks and the ^{228}Th peak (right) remain much more steady, varying less than 1%.

Table 5.4: The mean ratios of the fitted Cs and Co peaks to the Th peaks taken from calibration campaigns. The error on data is the RMS of the ratios from the campaigns.

isotope	E (keV)	Single Site		Multiple Site	
		data	simulation	data	simulation
¹³⁷ Cs	662	0.2513 ± 0.0007	0.2523	0.2496 ± 0.0009	0.2507
⁶⁰ Co	1173	0.4488 ± 0.0007	0.4480	0.4488 ± 0.0006	0.4474
⁶⁰ Co	1333	0.5109 ± 0.0008	0.5095	0.5100 ± 0.0004	0.5087

The calibrated energy of an event, then, with rotated energy $E_R(\theta(t))$ is (taking into account the fit function bias)

$$E = (E_{\text{Th}}^{\text{true}} - E_{\text{Th}}^{\text{bias}}) f \left(\frac{E_R(\theta(t))}{E_{\text{Th}}^{\text{fit}}(t)} \right) \quad (5.4)$$

where

$$f(x) = \begin{cases} 5.17 \times 10^{-3} + 0.9797 \cdot x + 1.516 \times 10^{-2} \cdot x^2 & (\text{single-site events}) \\ 5.37 \times 10^{-3} + 0.9784 \cdot x + 1.635 \times 10^{-2} \cdot x^2 & (\text{multi-site events}) \end{cases} \quad (5.5)$$

The residuals from the calibration campaigns are shown in fig. 5.9. This ratio-based method allows for some correction of the time variation without requiring additional time to be spent calibrating.

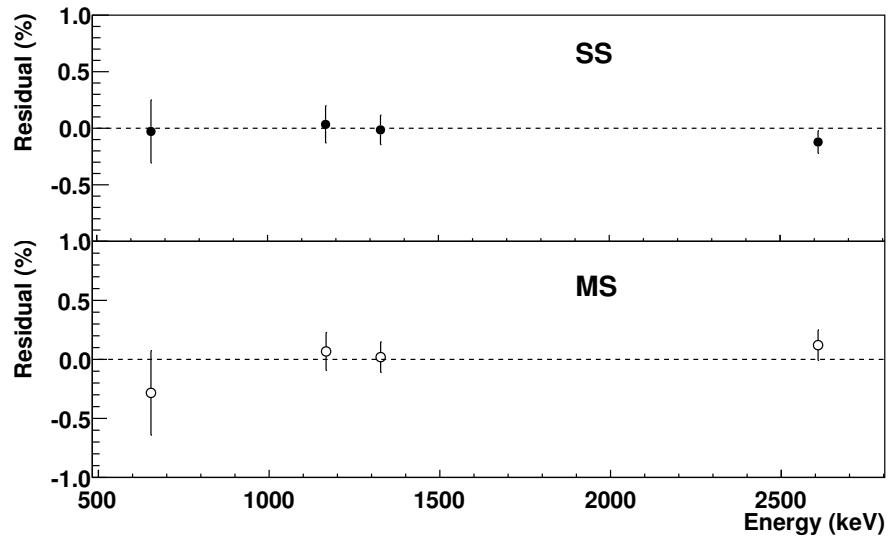


Figure 5.9: The residuals $(E^{\text{fit,data}} - E^{\text{fit,sim}})/E^{\text{sim}}$ for the calibration source peaks after the calibration is applied. The fit values are used because both data and simulation suffer from a fitting function bias. The top shows them for the single site calibration, while the bottom shows them for the multiple site calibration. The points are the mean of all source calibration campaigns, while the error bars represent the RMS. The small values indicate the calibration successfully reconstructs the true energy of an event.

5.5 Data Quality Cuts

Once the data has been processed, a number of data quality cuts are applied before it can be used in analyses. These cuts help ensure that the detector response and behavior are consistent throughout the data set. Summarizing them:

- The DAQ must be operating normally. All voltages must be nominal, and the trigger must be the standard one. Channels must not have an unusual baseline or RMS (indicating high noise).
- The number of forced 0.1 Hz triggers must be consistent with the DAQ recorded start and end time for the data run.
- The rate at which the DAQ records triggers must not be unusually high or low. Individual channels must not be creating abnormally many or few triggers.
- The active muon veto must be operational.
- The electron lifetime in the TPC must be good and not changing rapidly. The considerations behind this cut are explained in section 6.2.3.
- The rate of events that cannot be reconstructed (as in section 5.2) must be small.
- The rate of events with energies greater than 300 keV, 1000 keV, and 2000 keV should not vary significantly from previously observed rates.
- The run must be long enough (1800 s) that the above criteria can be properly evaluated.

Chapter 6

Electron Lifetime

6.1 Electron Capture on Impurities

As described in section 3.3.1, when an electromagnetic process deposits energy in a noble liquid detector, it ionizes the atoms, producing electrons and ions that drift in an electric field. In a detector consisting of perfectly pure noble liquid, the electrons would all reach the anode and would be collected for an energy measurement. In a non-ideal detector, however, electronegative impurities can capture the drifting electrons and form ions. The ions are more massive and drift more slowly, and so they escape inclusion in the signal used for energy measurement.

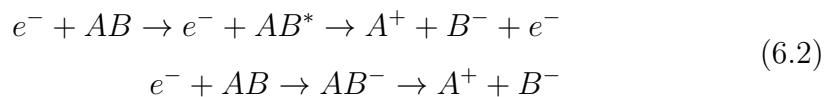
Electronegative impurities may capture electrons in three ways [52]. Denoting the impurities, which may be atoms or molecules, as AB , these are:

1. Radiative attachment



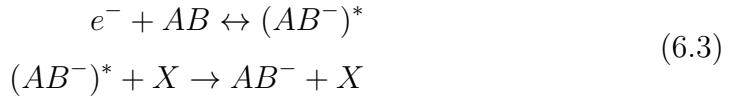
which has a much smaller cross section than the other processes below.

2. Dissociative attachment



which requires the electron's energy to be much higher than typically found for an electron drifting in a liquid or dense gas.

3. Three-body attachment through the two-stage Bloch-Bradbury reaction



where X represents the atom or molecule that make up the majority of the liquid.

The three-body reaction shown in eq. (6.3) releases some amount of energy, given by the “electron affinity” of AB . The electron affinity is positive if AB is electronegative. Noble elements have a negative electron affinity, so the reaction does not take place in a pure detector.

The rate of the reaction shown in eq. (6.3) is given by:

$$\frac{dn_{AB}}{dt} = -k_3 n_{AB} n_X n_{e^-} \tag{6.4}$$

where k_3 is constant for the 3-body reaction, and n_{AB} , n_X , and n_{e^-} are the densities of the impurity, the atoms or molecules of the liquid, and the electrons, respectively. k_3 depends on the species of the impurity, the species of the liquid, and the electric field strength.

According to eq. (6.4), electrons will be captured, forming AB^- at a rate proportional to the density of electrons. Thus, the number of free electrons will decay exponentially over time according to:

$$N_{e^-}(t) = N_0 \exp(-t/\tau_e) \tag{6.5}$$

where N_0 is the original number of electrons, and τ_e is the “electron lifetime”.

In general, there can be several different species of electronegative impurity present.

In that case, they all contribute to the electron lifetime according to:

$$\tau_e^{-1} = \sum_i k_i n_i = \sum_i \tau_i^{-1} \quad (6.6)$$

where n_i is the density of an electronegative impurity and k_i is the cross section for electron capture by that impurity.

6.2 Measuring Electron Lifetime

6.2.1 Method

Equation (6.5) provides a simple recipe for measuring the electron lifetime: measure the exponential attenuation of a known quantity of ionization as a function of drift time. A source of monoenergetic gamma rays can provide such a signal. The ionization produced by the full absorption of a gamma ray forms a peak in the energy spectrum. If some of the ionization is being attenuated, then the peak's mean will shift downward in energy from its true value.

A single calibration source at the cathode of the detector illuminates both TPCs, and the gamma rays interact throughout the entire drift region. After a sufficiently long calibration run, the TPC is divided into drift time bins. The size of the bins is a compromise: a larger bin in drift time will have more events, and thus the measurement of the full absorption peak energy will be better. A smaller bin will have fewer events, but will suffer less energy smearing because events will have similar drift times. In practice, 16 bins per TPC seems a good compromise.

The fit model used for a full-absorption peak is described in section 5.4.1. For each drift time bin, this simple Gaussian + complementary error function model is fit to the energy spectrum for single site events in that bin using an unbinned maximum likelihood fit. This fit is performed twice, first over a broad energy range to find the full-absorption peak, then in the range $(-3.0\sigma, +2.5\sigma)$ of the found peak to precisely determine the peak energy. Figure 6.1 shows an example fit.

Plotting the full absorption peak energy from each drift time bin as a function of

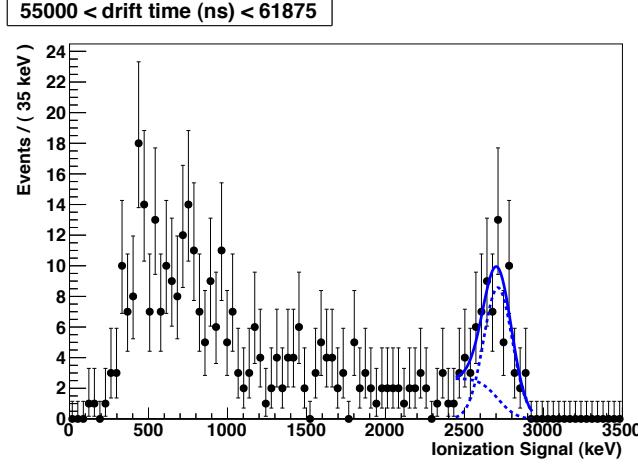


Figure 6.1: A fit of the simple Gaussian + complementary error function model to one single drift time bin. In this example, the full absorption peak is the 2615 keV gamma line from a ^{228}Th calibration source.

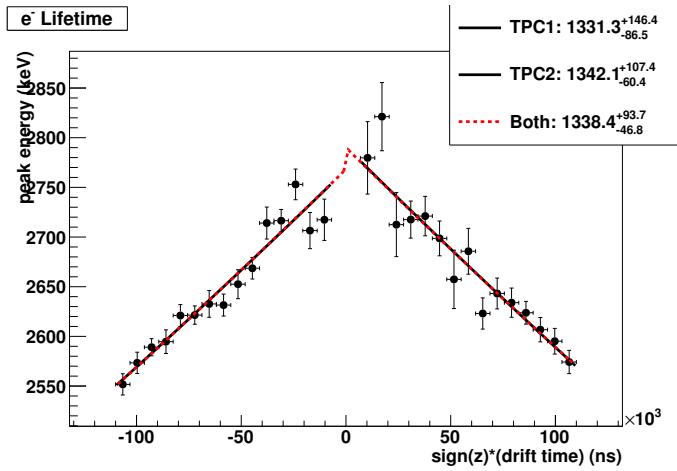


Figure 6.2: Measuring the electron lifetime by fitting a decaying exponential to the full-absorption peak energies binned by drift time. TPC 2 is assigned a negative drift time for convenience in visualization. Both fits to the individual TPCs and to both TPCs together are shown. In this case, there is no large difference between the combined fit and the individual fits.

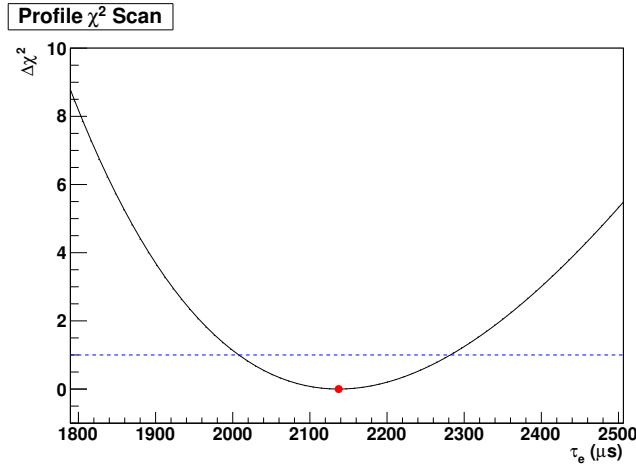


Figure 6.3: Confidence intervals around the best fit electron lifetime come from a profile scan, shown here. As this figure shows, this is superior to simply estimating the 1σ errors from the second derivative at the best fit value, since the profile is asymmetric around the best fit (indicated by the red dot). The blue line indicates $\Delta\chi^2 = 1$, corresponding to a 68% confidence interval.

drift time reveals the exponential decay described in eq. (6.5). Fitting an exponential to each TPC yields a measurement of the electron lifetime for each. Alternatively, a fit with a single electron lifetime to the entire detector uses information from both TPCs. In all cases, the amplitude of the exponential is allowed to float in the fit, since only the relative decay matters when measuring the electron lifetime. Figure 6.2 shows an example. Presently, the separate TPC lifetimes are used when correcting for electron lifetime in EXO-200, while the single measurement is used for detector monitoring and to ensure data quality. In EXO-200, the electron lifetime is typically $> 1 \text{ ms}$ over a $100 \mu\text{s}$ drift time, and so the attenuation begins to look approximately linear.

The electron lifetime measurement comes from minimizing the χ^2 statistic. Confidence intervals for the measurement come from doing a profile scan. Short electron lifetimes are easily distinguished, while longer electron lifetimes are not, and so the intervals will be asymmetric about the minimum. For a profile scan the electron lifetime is set to some fixed value away from the best fit value, and the amplitude

is allowed to vary to minimize χ^2 . All electron lifetime values for which this remimized χ^2 is less than 1 above the minimum value define the 1σ (68%) confidence band. Figure 6.3 shows an example.

6.2.2 Comparison to Simulation

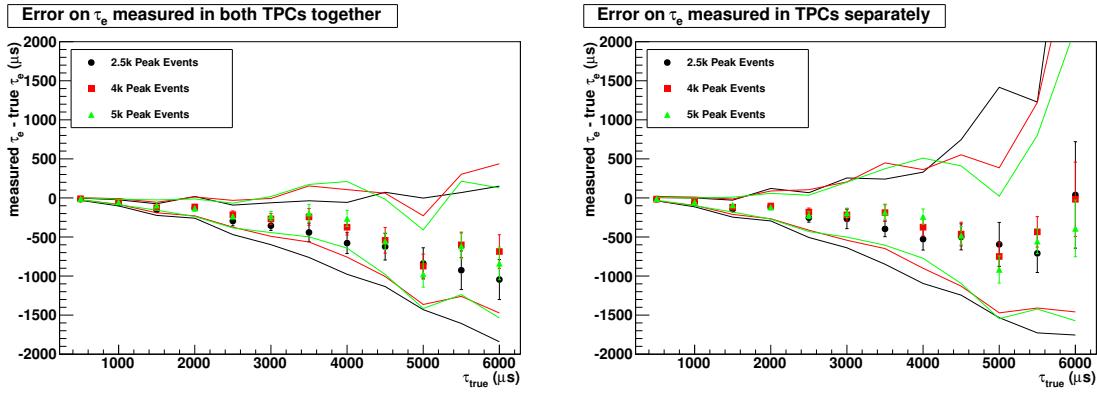


Figure 6.4: The absolute error on the electron lifetime measurement for the electron lifetime measured in both TPCs simultaneously (left) and individually (right). The points show the mean error of 20 simulations. The lines show the size of the 68% confidence bands returned by the measurement for those simulations. The measurement method systematically underestimates the purity, and the error grows as the electron lifetime gets large. Typical calibration runs include 2.5k events in the full absorption peak (black), but using more events would improve the error.

Figure 6.4 shows a comparison between a known simulated electron lifetime and the measurement of that lifetime using the method described above. The error is small for electron lifetimes below 1 ms. For large electron lifetimes, however, the method consistently underestimates the electron lifetime, with the effect getting worse as the electron lifetime improves.

This effect seems to be due to some z -dependence introduced in processing the data. In simulations of a ^{228}Th source at the cathode with infinite electron lifetime, the method reports electron lifetimes of $(3.0 \pm 0.9) \times 10^4$ μs measured in both TPCs simultaneously, and $(4.3 \pm 0.9) \times 10^4$ μs measured in a single TPC for runs with 2.5k events in the full absorption peak. It reports electron lifetimes of $(5.0 \pm 0.7) \times 10^4$ μs measured in both TPCs simultaneously and $(4.9 \pm 0.8) \times 10^4$ μs measured in a single

TPC for runs with 5k events in the full absorption peak. This effect does not appear in simulations that have not had the shielding grid correction (section 5.3.2) applied. The shielding grid correction may introduce a slight z dependence that is measured when attempting to correct for electron lifetime.

The error due to this effect is small, however. A measured electron lifetime of 3000 μs corresponds to a true electron lifetime of roughly $(1/3000 - 1/40000)^{-1} = 3250 \mu\text{s}$. This will only move the corrected ionization signal 0.3% higher than its true value if the ionization drifts over the full 120 μs drift time. The effect on the energy resolution in the ionization channel will be half of this, and even smaller in the rotated spectrum. Since the source of the error is presently not well understood, it is not corrected for, and so is accepted as a potential detrimental effect on the energy resolution.

6.2.3 Practical Considerations

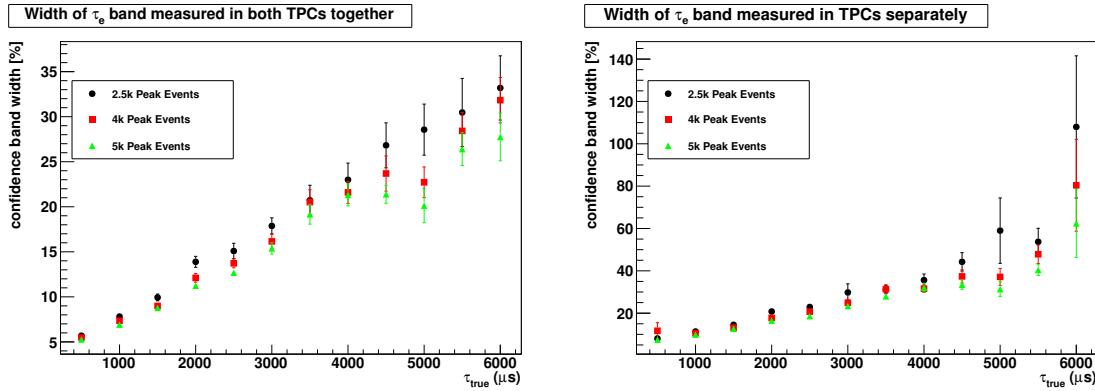


Figure 6.5: The width of the 68% confidence band for the electron lifetime measured in both TPCs simultaneously (left) and individually (right). Typical calibration runs include 2.5k events in the full absorption peak (black dots), but using longer calibration runs to get more events in the peak would improve the uncertainty. However, for long electron lifetimes, it becomes difficult to detect the attenuation and measure it, giving a large uncertainty.

As the electron lifetime grows large, it becomes increasingly difficult to measure.

For a 4000 μs electron lifetime, ionization drifting the full distance will only be attenuated about 3%. This is comparable to the 3.4% energy resolution in the ionization channel at the 2615 keV full absorption peak from ^{228}Th . This effect is shown in fig. 6.5. The width of the confidence band for the measurement grows to nearly 20% of the measured value for large electron lifetimes. Taking more calibration data only partially mitigates this effect, also shown in fig. 6.5. For low electron lifetimes, a simulated calibration run with twice as many events in the full absorption peak (and requiring twice as much time to run) shrinks the confidence band by roughly $\sqrt{2}$. For large electron lifetimes, more events don't shrink the band as much, suggesting systematic effects such as the relatively short drift time begin to dominate.

6.3 Effects of Electron Lifetime on the Energy Resolution

The error on the energy E_{combined} (obtained through the combination of ionization and scintillation as described in section 5.4.2) due to an uncertainty $\Delta\tau$ in the electron lifetime is:

$$\frac{\Delta E_{\text{combined}}}{E_{\text{combined}}} = \cos(\theta) \frac{t_d}{\tau^2} \Delta\tau \quad (6.7)$$

and the error on E_{combined} due to an uncertainty Δt_d in the drift time is:

$$\frac{\Delta E_{\text{combined}}}{E_{\text{combined}}} = \cos(\theta) \frac{\Delta t_d}{\tau} \quad (6.8)$$

6.3.1 Position Uncertainty

True Single-Site Events

For true (point-like) single-site events, the drift time can be measured to within about 0.2 μs thanks to the information provided by including multiple points in the fit that extracts information from waveforms. This uncertainty will cause some smearing of the resolution, since the correction relies on a measurement of the drift time. The effect is easy to calculate using eq. (6.8). Table 6.1 provides some concrete numbers.

Table 6.1: The effect of a $0.2\text{ }\mu\text{s}$ drift time uncertainty on the rotated energy resolution, assuming a $100\text{ }\mu\text{s}$ drift time.

τ (μs)	$\Delta E/E$ (%)
100	0.12
200	0.06
400	0.03
800	0.02
1000	0.01
1500	0.01
2000	0.01
2500	<0.01

Events with Spatial Extent

Most events, however, are not point like, and the ionization arrives over a finite period of time. The resolution of the multiple signal finder described in section 5.2.1 leads to an uncertainty on the true drift time of $\sim 3\text{ }\mu\text{s}$, which will result in more smearing of the energy resolution due to the electron lifetime correction. Table 6.2 shows the spread for this scenario.

6.3.2 Electron Lifetime Uncertainty

The effect of electron lifetime uncertainty on the energy resolution is given by eq. (6.7). With a fixed drift time, it becomes increasingly difficult to measure long electron lifetimes because the attenuation is so slight. The drift distance in EXO-200 is fixed. Varying the electric field to increase the drift time carries some risk to the experiment and so is not done routinely. Furthermore, the attachment rates for electronegative species vary with electric field (see section 6.4.4). Lacking knowledge of the precise nature of the impurities in EXO-200, it is not clear this would provide a good estimate for the electron lifetime at normal electric fields. Since the uncertainty cannot be easily removed, parameterizing the uncertainty as a function of electron lifetime provides a way to estimate its effect on energy resolution.

As shown in fig. 6.6, a polynomial function nicely fits the observed uncertainties.

Table 6.2: The effect of a 3 μ s drift time uncertainty on the rotated energy resolution, assuming a 100 μ s drift time.

τ (μ s)	$\Delta E/E$ (%)
100	1.81
200	0.91
400	0.45
800	0.23
1000	0.18
1500	0.12
2000	0.09
2500	0.07
3000	0.06
3500	0.05

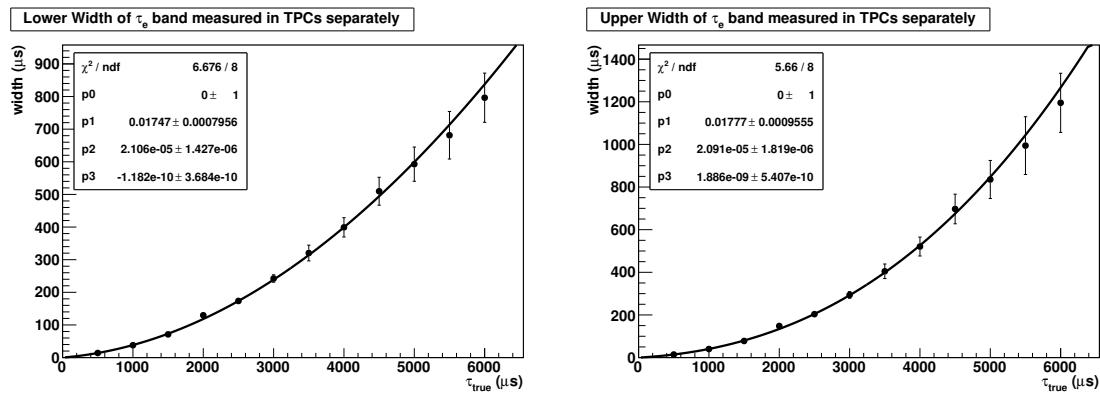


Figure 6.6: The width of the confidence bands on single electron lifetime measurements as a function of the electron lifetime. Since the uncertainties are asymmetric, the negative uncertainty (left) and the positive uncertainty (right) are fit separately. The fit function is a cubic polynomial.

Table 6.3: The effect of the electron lifetime uncertainty on energy resolution. This is based on the parameterization shown in fig. 6.6. Note that the reported resolution assumes only one measurement is used for the correction. In practice, more measurements are used, and so the effect will be smaller.

τ (μs)	$\Delta E/E$ (%)
100	2.38
200	1.32
500	0.68
800	0.53
1000	0.48
1500	0.41
2000	0.38
2500	0.37
3000	0.36
3500	0.35

The effect on energy resolution is tabulated in table 6.3. In practice, however, a number of measurements go into the actual correction function applied to the data, and so the effect will be smaller.

6.3.3 Rate of Change

The electron lifetime can vary with time. As the pump operates, it can lose pumping ability. This slows the recirculation rate and causes a decline in electron lifetime. The most dramatic instances are following pump outages, when recirculation through the purifiers (as described in section 4.5) is completely stopped, allowing xenon to stagnate. Impurities leaching out of the materials in the vessel, or from liquid xenon coming into contact with different sections of plumbing can cause the purity to degrade rapidly. When the pump is turned back on, the purity recovers over several days.

Suppose there are only two measurements of electron lifetime, $\tau_1 \pm \sigma_{\tau_1}$ and $\tau_2 \pm \sigma_{\tau_2}$. The best estimate of the rate of change is simply

$$\frac{d\tau}{dt} = \frac{\tau_2 - \tau_1}{\Delta t} \quad (6.9)$$

Table 6.4: The effect of the rate of change of electron lifetime on energy resolution. This shows the additional effect due to uncertainty in the rate of change, with measurements taken 1 day apart. Longer intervals between measurements will result in more smearing. Note that the reported resolution assumes only two measurements are used for the correction. In practice, more measurements are used, and so the effect will be smaller.

τ (μs)	$d\tau/dt$ (μs d ⁻¹)	$\Delta E/E$ (%)
1000	10	0.03
	50	0.15
	100	0.29
	500	1.04
	1000	1.50
2000	10	0.01
	50	0.04
	100	0.08
	500	0.15
	1000	0.55
3000	10	<0.01
	50	0.02
	100	0.04
	500	0.16
	1000	0.29

However, propagating the uncertainties, the error on the estimate is

$$\Delta \frac{d\tau}{dt} = \frac{d\tau}{dt} \sqrt{\left(\frac{\sigma_{\tau_1}}{\tau_1}\right)^2 + \left(\frac{\sigma_{\tau_2}}{\tau_2}\right)^2} \quad (6.10)$$

If the true rate of change is $+1\sigma$ away from the estimate, events close to the first measurement will be undercorrected, and events close to the second measurement will be overcorrected. This results in an overall smearing to the energy resolution, tabulated in table 6.4. This effect will grow if the time between measurements is increased, since more data will be taken with an estimate far from the true value.

6.3.4 Overall

The effects of electron lifetime on resolution are taken in account when selecting which runs will be used for the analysis. The guidelines are such that the resolution smearing due to the electron lifetime correction is no more than 0.5%. Runs can be used if:

1. Electron lifetime is above 1000 μ s (due to the effects of position uncertainty and electron lifetime uncertainty)
2. Four or more measurements all show similar electron lifetime (to reduce the effect of the electron lifetime uncertainty)
3. The electron lifetime must not be increasing by more than half of its current value per day, nor decreasing by more than a quarter of its current value per day (to reduce the effect of time variation)

6.4 Measurements of Electron Lifetime in EXO-200

Calibration runs taken every 1–2 days serve to measure the electron lifetime in EXO-200. In a typical calibration run, a ^{228}Th source at the cathode creates 2.5×10^5 events in the TPC. Of these, approximately 2500 are single site events within 2σ of the full-absorption peak.

6.4.1 Time Variation and Correction Function

The measured electron lifetime varies in time. Usually, this variation is small and slow. To account for this, a piecewise polynomial is fit to the measured electron lifetimes. This piecewise polynomial can be discontinuous across sudden changes in electron lifetime. The polynomial degree can change when the behavior of the electron lifetime changes, such as when rapidly-increasing lifetime after resuming recirculation becomes a steady-state, slowly-varying value. Figure 6.7 shows the time variation and

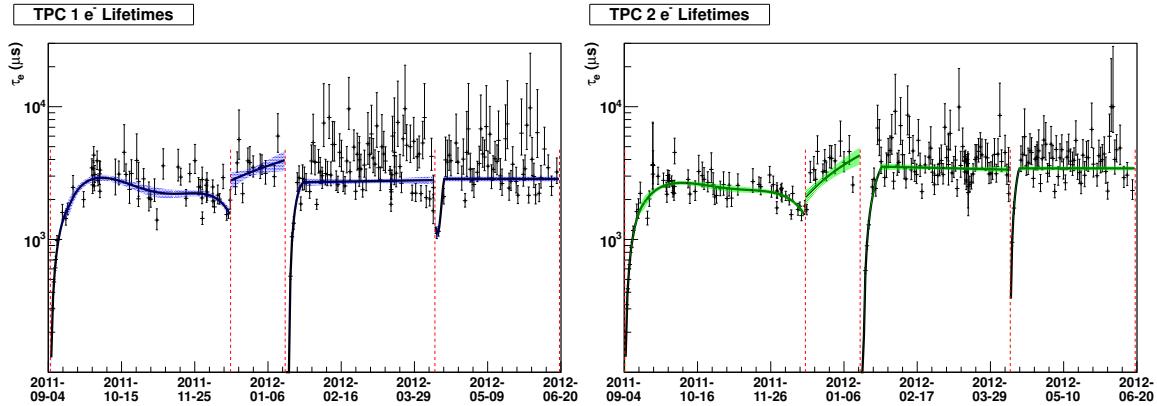


Figure 6.7: The fit of a piecewise polynomial to electron lifetime in TPC 1 (left) and TPC 2 (right) for the time period containing the Run 2a data set. The colored bands show the 68% confidence interval on the fit. The vertical dashed lines indicate discontinuities in the electron lifetime due to sudden changes in recirculation rate or xenon feeds.

the polynomial fit for the separate TPCs. Separate electron lifetimes are used for the different TPCs because there could be some purity gradient in the chamber, and splitting the chamber in half provides a modest approximation. Furthermore, the measured values in the different TPCs are observed to sometimes vary outside of each others' confidence bands.

For a good event in EXO-200, the reconstruction algorithms find both a drift time and an (attenuated) ionization signal. The polynomial fit described above provides an estimate of the electron lifetime at the event time. Equation (6.5) provides a recipe for correcting the attenuated ionization signal to get a corrected signal, using the drift time and measured electron lifetime.

6.4.2 Comparison with Recirculation Rate

The rate at which xenon is recirculated through the purifiers affects electron lifetime. Figure 6.8 shows a clear trend of increasing electron lifetime with increasing recirculation rate. The highest electron lifetimes are achieved with a recirculation rate above 13 slpm, which corresponds to completely recirculating the volume of the chamber in 1.8 days.

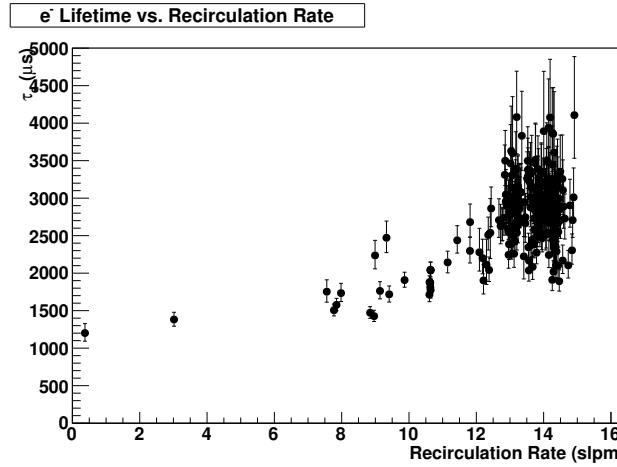


Figure 6.8: The electron lifetime plotted as a function of recirculation rate for the Run 2a time period. Time periods when recirculation was fast, but electron lifetimes were low due to a pump stoppage or feed event have been removed. A clear trend is visible in which electron lifetime increases with recirculation rate. The points for rates below 6 slpm do not reflect the equilibrium values of the electron lifetime for rates; lower values have been observed with similar recirculation rates.

Figure 6.9 shows the time history of the electron lifetime, plotted along with the recirculation rate. Since the electron lifetime in the chamber decreases when the pump is recirculating at a reduced rate, there is most likely a constant source of impurities in the TPC. That the electron lifetime worsens significantly after prolonged recirculation stoppages supports this. However, when the recirculation stops, the pressures throughout the system change. This causes liquid levels to change and may also cause the slow control system to feed in more gas. Newly exposed or submerged plumbing, and new gas (even though it is fed through the purifiers, some may be able to feed backward through the pump) could also cause this decrease in electron lifetime, and it is difficult to disentangle the effects. In any case, once recirculation is resumed, the electron lifetime recovers over the course of a few days.

6.4.3 Comparison with Gas Purity Monitor Readings

The gas purity monitors [48] provide real-time monitoring of the recirculating xenon. GPM 1 samples the gas at the output of the recirculation pump. GPM 2 samples

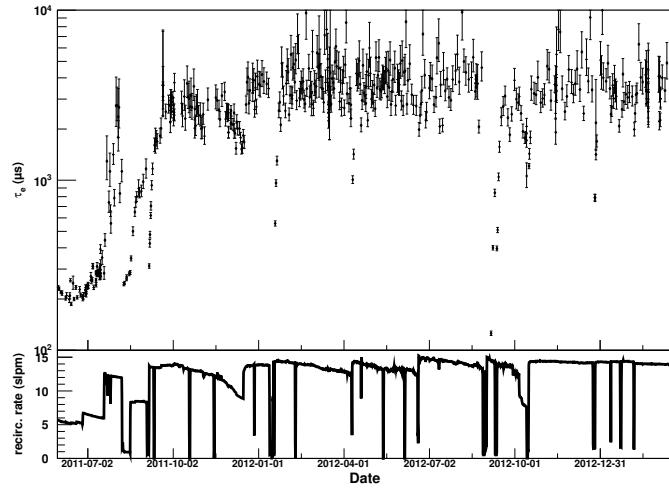


Figure 6.9: A time history of the electron lifetime, with the recirculation rate plotted below. The electron lifetime drops when the pump slows or completely stops, and recovers when recirculation resumes at a fast rate.

the gas coming out of the purifiers. GPM 3 samples the gas returning from the TPC. They are shown in fig. 4.7. Because these instruments sample room temperature gas instead of cryogenic liquid, some impurities may be in different concentrations than in the TPC due to different solubilities. Likewise, the gas purity monitors have a small electric field and short drift distance, and so they will not be able to measure long electron lifetimes. Despite this, GPM 3 may be able to provide some information about the purity of the xenon in the TPC when that purity is poor.

Figure 6.10 shows a plot of the gas purity monitor reading for gas returning from the TPC during periods when the electron lifetime was poor due to a pump stoppage or feed event. A loose correlation is visible, with a correlation coefficient of 0.66 over all points. More interesting is that the correlation is stronger during individual recovery incidents, denoted by different markers and colors. So, while the gas purity monitors do not measure the electron lifetime in the TPC on an absolute scale, they provide a good indication of large changes in the electron lifetime.

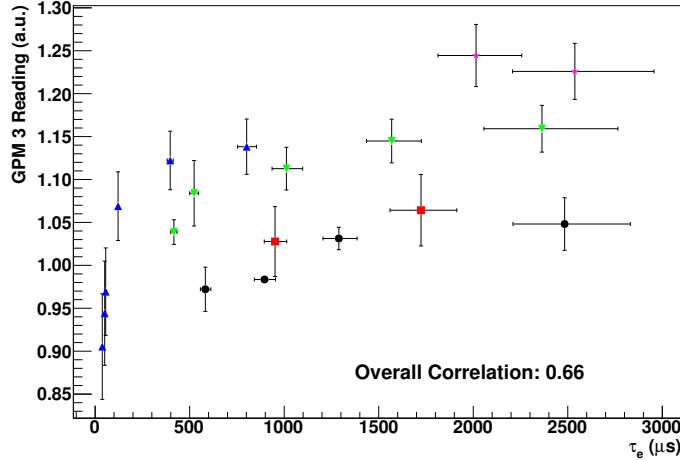


Figure 6.10: The reading from the gas purity monitor sampling the gas returning from the TPC (GPM 3), during periods when the electron lifetime was recovering from a pump stoppage or feed event. Different recovery incidents are denoted with different colors and markers.

6.4.4 Electron Lifetime for Different Electric Fields

The rate constant for electron attachment to impurities varies with the electric field strength, and the species of impurity determines the nature of this variation. As shown in fig. 6.11 [53], O₂ impurities show decreasing attachment with increasing field strength, while N₂O show the opposite.

Calibration runs of the standard length were taken with the cathode at -1.1 kV , -1.45 kV , -2.2 kV and -4.4 kV , corresponding to electric field strengths of 20 V cm^{-1} , 38 V cm^{-1} , 75 V cm^{-1} and 187 V cm^{-1} . Normal operation is at -8 kV , corresponding to 374 V cm^{-1} . Unfortunately, there were not enough events in the full-absorption peak to measure the electron lifetime at the lowest field strength. The results are shown in fig. 6.11. The electron lifetime does not seem to increase with increasing electric field, as would be expected for oxygen. There is not enough information to identify the exact species of impurity present. For species with similar molecular weight and electron attachment as N₂O and O₂, the required concentration of impurities for a 3 ms electron lifetime is on the order of 10^{-11} g/g .

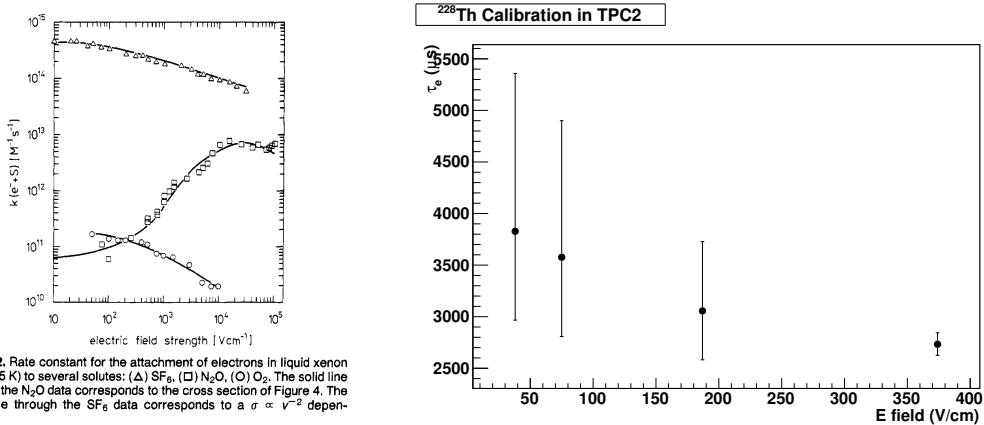


Figure 2. Rate constant for the attachment of electrons in liquid xenon ($T = 165 \text{ K}$) to several solutes: (Δ) SF_6 , (\square) N_2O , (\circ) O_2 . The solid line through the N_2O data corresponds to the cross section of Figure 4. The solid line through the SF_6 data corresponds to a $\sigma \propto V^{-2}$ dependence.

Figure 6.11: On the left is a figure by Bakale et al. [53], showing the electron attachment constant as a function of electric field strength. On the right is the measured electron lifetime, which is inversely proportional to the attachment constant, for a number of electric field strengths. For these runs, the calibration source was located slightly on the TPC 2 side of the cathode, giving more events and allowing a better measurement in this TPC. The nominal field for normal operation is 374 V cm^{-1} . Although a trend is visible, the results are not incompatible with a constant electron lifetime ($\chi^2/\text{n.d.f.} = 3.15/3$). (Left figure reprinted with permission. Copyright 1976 American Chemical Society.)

Chapter 7

Muons

7.1 Motivation

Cosmic rays interact in the upper atmosphere, forming pions and kaons that decay to muons. Some of these muons have very high energies, and so can travel through lots of matter, reaching deep underground. These muons can produce backgrounds in EXO-200. For example, some muons pass through EXO-200, depositing energy. Furthermore, radiative losses create gamma rays or charged particles that can deposit energy in EXO-200. Finally, muons can produce neutrons through spallation. Spallation neutrons can capture on detector materials or the HFE, producing high-energy gamma rays that might interact in the detector. Any of these processes could leave behind energy close to the Q value for ^{136}Xe . These processes occur promptly or on millisecond timescales, and so these events can be vetoed if the muon can be tagged.

More troublesome is when the neutron capture creates an unstable isotope. For example, a neutron capture on ^{136}Xe produces ^{137}Xe , an isotope that beta decays with a Q value of 3.8 MeV, which means the beta particle produced in the decay could have an energy close to the Q value of ^{136}Xe . Likewise, the $(n, 2n)$ reaction on ^{136}Xe can produce ^{135}Xe , which beta decays with a Q value of 1.2 MeV and is a background for the $2\nu\beta\beta$ mode. The half-lives of ^{137}Xe and ^{135}Xe are 3.8 min and 9.1 h respectively, and so vetoing becomes more difficult. However, a good measurement of the muon flux can constrain the rate at which these and other isotopes are formed.

7.2 Identifying Muons

Muons that reach underground have typical energies of 1 GeV to 1000 GeV, a range in which they are minimally ionizing. Those that pass through the detector typically leave a straight line of ionization along their path. Like any energy deposition in EXO-200, some of the energy deposited becomes scintillation light, while the remaining ionization is drifted and collected. Thus, a muon passing through EXO-200 will show a bright flash of light, followed by ionization across many wire channels, linearly spread in time. Figure 7.1 shows a typical example. This distinct linear trail provides a means to tag the muon.

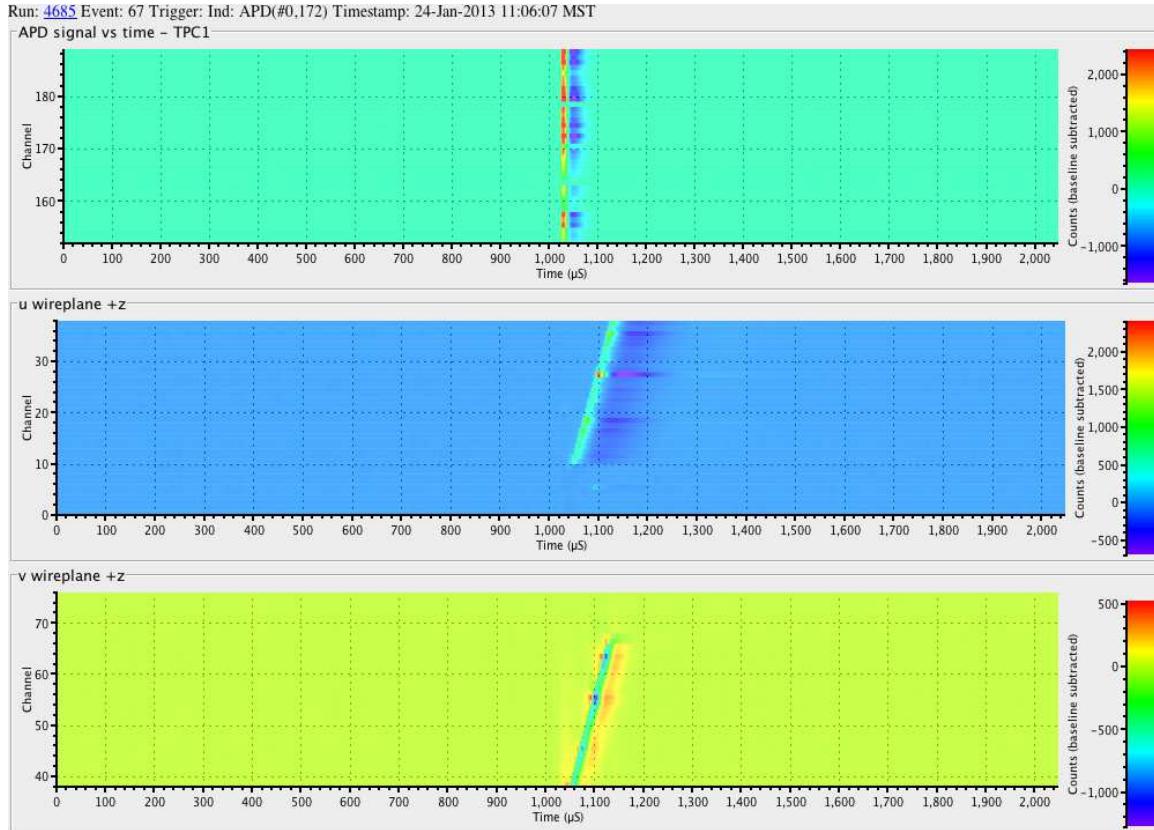


Figure 7.1: The event display for a muon passing through EXO-200. The upper panel shows the APD channel waveforms versus time, displaying the flash of light when the muon passes through. The middle (lower) panel shows u (v) wire channel waveforms versus time, showing the characteristic linear trail of ionization. The trail makes a positive signal in u , and a negative signal in v .

7.2.1 Identifying Muons with the Hough Transform

The Hough transform [54, 55] is an algorithm invented to look for tracks in bubble chamber photographs. A line can be completely parameterized by the perpendicular distance r from the line to the origin and by the angle η that perpendicular line makes with an axis. The Hough transform maps a point (x, y) in an image to all possible lines that can pass through that point:

$$r = x \cos \eta + y \sin \eta \quad (7.1)$$

Thus, a point (x, y) is transformed into a sinusoidal curve in (η, r) space. If some of the points form a line, these curves will converge at the coordinates corresponding to that line.

This can be applied to EXO-200 to look for muons. First, an event must pass a simple scintillation cut. The sum of the individual APD waveforms must sum to at least 10000 ADC units, which is above the typical signal associated with γ and β events under 3 MeV. Next, ionization deposits are identified by simply looking at the collection wire channel waveforms. The first 256 samples are averaged to find the baseline. A “hot spot” is associated with the time value at which the waveform peaks, once it has climbed 80 ADC units above the baseline. The induction wires are analyzed the same way, except looking for the waveform dipping 80 ADC units below its baseline. The hot spots undergo a Hough transform, and the curves are placed in a histogram. The bin with the most entries corresponds to the line that passes through the most points. This line is a projection of the muon’s track onto the wire planes. In the case when multiple bins have the most entries, meaning the best track does not quite line up with the binning, the median of the bins (sorted by η) is used. Figure 7.2 shows an example of a reconstructed muon track.

After passing a check for noise (described in Appendix A) and a check for a large amount of scintillation light, an event will be tagged as a muon if it has good tracks in both the induction and collection wire planes in at least one TPC. A “good” track must be reconstructed from at least 5 hot spots that lie along the reconstructed track, or at least 3 hot spots along the track if fewer than 5 total spots were found.

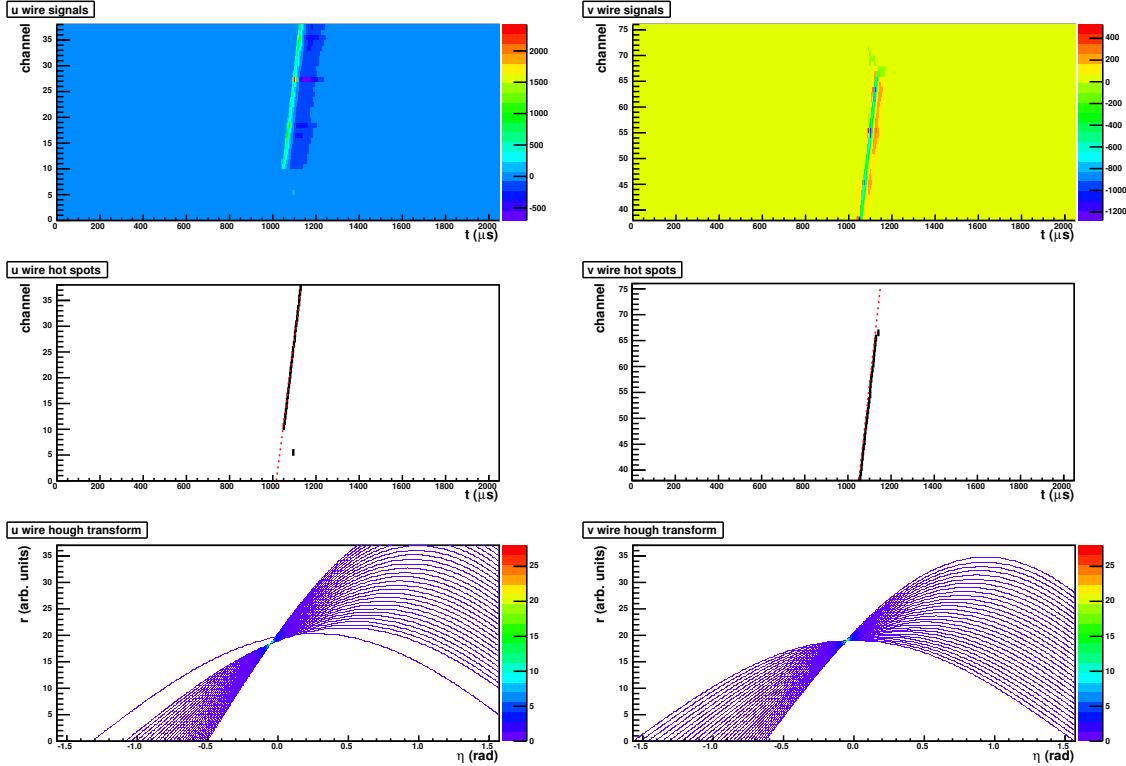


Figure 7.2: The Hough transform used to identify and reconstruct a muon passing through EXO-200. The left side shows collection wire channels, and the right side shows induction wire channels. The upper panels show the raw waveforms, baseline subtracted. The middle panels show the identified hot spots in the waveforms in black, and the reconstructed tracks in red. The lower panels show the Hough transforms of the hot spots above, which converge on the points corresponding to the red lines in the middle panels.

The Hough transform reconstructs the projection of the muon's path onto the wire planes. Since the drift velocity and wire spacings are well known, these projections can be translated back to incident zenith (θ) and azimuthal (ϕ) angles for the muon. However, some information is lost in the time projection. For example, consider muons that pass parallel to the wire plane. All their charge arrives simultaneously on the wires. While the signal is distinctive and will be tagged as a muon, there is no way to determine the zenith angle. Likewise, a muon may pass parallel to a wire. In this case, it may not create enough hot spots to be tagged as a muon, or some ionization may reach perpendicular wires, leading to an incorrect reconstruction of the angle.

7.2.2 Validation with Monte Carlo Simulations

In order to validate this muon tagging algorithm, 2 million muons were generated in EXOSim, a GEANT4 [56] simulation of the EXO-200 detector. The overall efficiency of the tagging depends on the angular distribution of the incident muons for the reasons described above. The angular distribution underground is approximated [57] by

$$\frac{dN}{d\Omega} = \Phi(h) (\cos \theta)^{1.53} e^{-8 \times 10^{-4} h (\sec \theta - 1)} \quad (7.2)$$

where θ is the zenith angle and h is the vertical depth in hg cm^{-2} (1 hg cm^{-2} is equivalent shielding to one meter of water). For WIPP, previous experiments have measured h to be $1585_{-6}^{+11} \text{ hg cm}^{-2}$ [58], so this value was used for the simulations. This distribution is a good approximation for $\theta < \pi/3$.

The energy distribution for muons underground can also be approximated [59]. The distribution for muons at the surface is given by

$$\frac{dN}{dE} \propto E^{-2.7} \left(\frac{1}{1 + \frac{1.1E \cos \theta}{115}} + \frac{0.054}{1 + \frac{1.1E \cos \theta}{850}} \right) \quad (7.3)$$

for E in GeV. The flux underground is then

$$\frac{dN}{dE} = \frac{dN}{dE_0} e^{bh \sec \theta} \quad (7.4)$$

where bE defines the rate of continuous energy loss for muons. For standard rock, b is about $4 \times 10^{-6} \text{ g cm}^{-2}$. This makes the substitution

$$E_0 = e^{bh \sec \theta} (E + \epsilon) - \epsilon \quad (7.5)$$

which is the average energy of surface muons that pass through $h \sec \theta$ of material and emerge with energy E . The parameter ϵ is a critical energy above which discrete, radiative energy losses dominate, rather than continuous losses described by b . For standard rock, ϵ is about 693 GeV [60].

Muon events were simulated by first picking an azimuthal angle ϕ uniformly between $-\pi$ and $+\pi$ and a zenith angle θ from eq. (7.2) multiplied by a factor of $\cos \theta$ (explained below). The energy was selected in a range of 1 GeV to 500 GeV from eq. (7.4) for the selected zenith angle. A point was chosen uniformly randomly from a horizontal disk of radius 3 m just below the TPC. This is the reason for the $\cos \theta$ factor in the angular distribution, since the flux through that disk will be $\cos \theta$ less than a disk of the same area normal to the incident angle. The 3 m radius ensures the flux will be accurately simulated for $\theta \lesssim 80^\circ$.

A muon was then generated at a height of 3 m vertically above the plane of the disk with the selected energy and incident angles, and passing through the selected point. To save computation time, only muons that would actually pass through the TPC volume (taken to be a cylinder of radius 22.74 cm, the interior radius of the TPC, and length 2×20.44 cm, twice the distance from the cathode to the APD plane face) were simulated fully. Positively charged muons were generated in a 1.25 ratio to negatively charged muons, similar to the typical value in nature. The “standard rock” values for b and ϵ were used, and the depth h was varied 5% around 1585 hg cm^{-2} to account for systematic effects of the angular distribution on efficiency.

7.2.3 Reconstruction Accuracy

For most incident muon angles, the algorithm correctly reconstructs their incident angle. Figure 7.3 shows the rate for muons to be misreconstructed more than 5° from their true angle. Several pathologies are visible. As discussed above, muons

parallel to the wire plane are poorly reconstructed. Likewise, the angular resolution for near-vertical muons is poor, since hot spots are only identified with a resolution of $1\text{ }\mu\text{s}$. This analysis concentrates on the region with good angular reconstruction. This region is bounded by the polygon with vertices $(\theta, \phi) = \{(24^\circ, 0^\circ), (24^\circ, 48^\circ), (28^\circ, 48^\circ), (28^\circ, 60^\circ), (40^\circ, 60^\circ), (40^\circ, 72^\circ), (72^\circ, 56^\circ), (60^\circ, 56^\circ), (60^\circ, 60^\circ)\}$ and their reflection across $\phi = 0$, as well as the symmetric region on the opposite side of the detector. These specific values were chosen to align with the binning used for the analysis.

7.2.4 Efficiency

Overall, of two million simulated muons that hit the TPC, $1125723 \pm 1061(\text{stat})$ were tagged as muons by the algorithm described above. This is an overall efficiency of $56.3 \pm 0.1(\text{stat})\%$. However, the geometry of the detector means that the efficiency will be a function of the incident muon angle. Figure 7.4 shows the efficiency of detecting muons based on the ratio of muons reconstructed in an angular bin to muons simulated in that bin. This is useful for estimating the true flux in a bin.

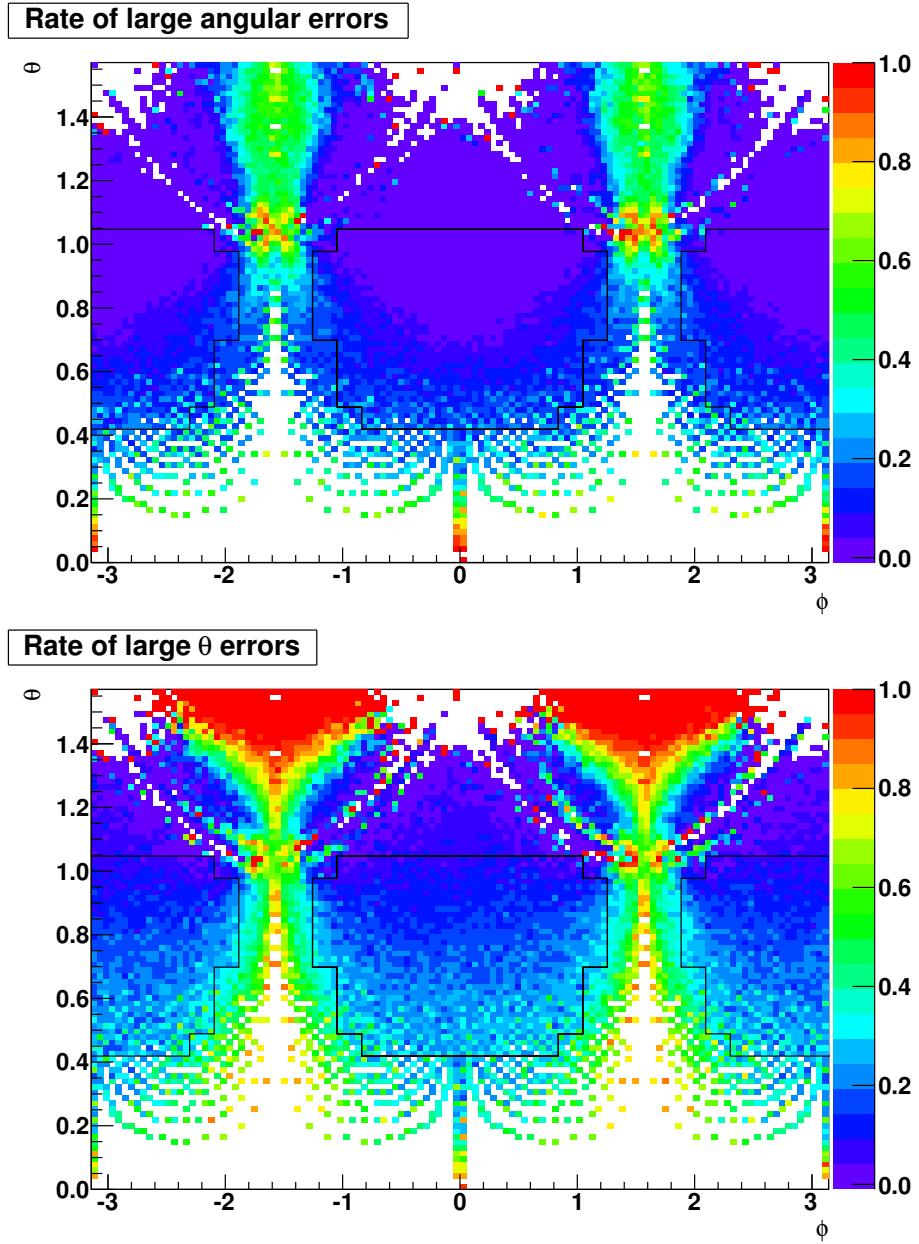


Figure 7.3: The rate for muons to be reconstructed far from their true direction. The top figure shows the rate for errors greater than 5° in total separation. The bottom figure shows the rate for errors greater than 5° in θ only, of concern because the flux varies in θ , but is isotropic in ϕ . The black polygon indicates the region with good reconstruction used for this analysis. The azimuthal angle ϕ has 3° bins, and the zenith angle θ has 1° bins. Blank bins had no muons reconstructed in that bin. Poor angular resolution for near-vertical muons leads to the many blank bins at small θ .

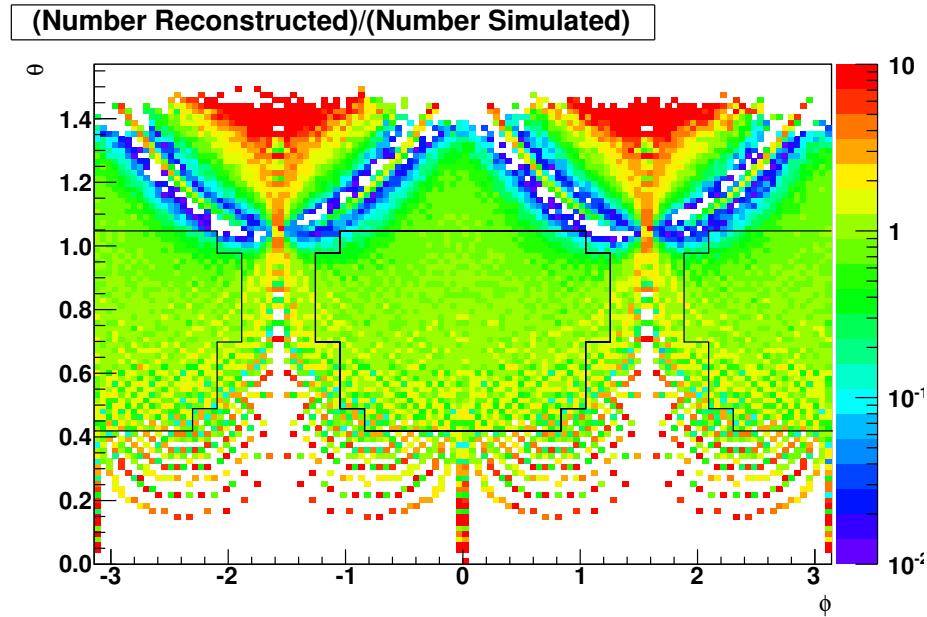


Figure 7.4: The ratio of muons reconstructed in an angular bin to the muons simulated in that bin. This includes muons incorrectly reconstructed, so that this can be used to estimate a total flux. The region inside the black polygon, used for this analysis, shows reasonable efficiencies. The bin at $(\theta, \phi) = (0, 0)$ is off scale high, with a ratio of 5.2×10^3 , owing to poor angular reconstruction for near-vertical muons, and the degeneracy for muons passing parallel to the wire plane.

7.3 The Muon Flux at WIPP

The muon flux through the detector divided into discrete solid angle bins is simply

$$\Phi = \frac{\sum_i N_\mu}{\Delta T \sum_i \Delta\Omega_i \epsilon(\theta_i, \phi_i) A(\theta_i, \phi_i)} \quad (7.6)$$

where N_μ is the number of muons observed in time ΔT , ϵ is the efficiency for that bin, $\Delta\Omega_i$ is the solid angle subtended by bin i , and A is the projected area of the detector for a particle incident from (θ_i, ϕ_i) . The projected area of a cylinder on its side is

$$A(\theta, \phi) = \pi r^2 |\cos(\phi)| \sin(\theta) + 2r\ell \sqrt{1 - \sin(\theta)^2 \cos(\phi)^2} \quad (7.7)$$

where r is the cylinder's radius and ℓ is its height. The dimensions of the cylinder used were the same as for the simulation ($r = 22.74\text{ cm}$, $\ell = 2 \times 20.44\text{ cm}$). If the muons were simulated through a larger cylinder, the efficiency would be lower, but the product of area and efficiency would be the same.

The pions and kaons in the atmosphere traveling with large θ are more likely to decay to muons than interact as they spend more time in the low-density upper atmosphere. This provides a $\sec\theta$ enhancement in the flux underground. Therefore, most cosmic ray measurements quote a vertical muon flux:

$$\Phi_v(\theta, \phi) = \Phi(\theta, \phi) \cos\theta \quad (7.8)$$

Integrating eq. (7.2) over the polygonal region of interest, the ratio of the total flux in the region of interest to the vertical flux is 0.553 ± 0.020 where the error is found by varying the depth h by 5% around 1585 hg cm^{-2} .

To calculate the flux, EXO-200 was divided into 4° bins in θ and 12° bins in ϕ to ensure good statistics in each bin. Only bins that had centers in the region of interest were considered. In data passing the quality cuts described in section 5.5 taken between October 14, 2011 and January 9, 2013, (11209 ± 105.9) muons were observed in $2.140 \times 10^7\text{ s}$. The integrated product of area, efficiency, and solid angle for bins in the region of interest was $(2.35 \pm 0.01) \times 10^3\text{ cm}^2\text{ sr}$. Applying the ratio

above to convert this into a total flux yields a vertical flux of

$$\Phi_v = (4.01 \pm 0.04(\text{stat})^{+0.04}_{-0.05}(\text{sys})) \times 10^{-7} \frac{\text{Hz}}{\text{cm}^2\text{sr}} \quad (7.9)$$

Comparing the vertical muon flux to a global fit [61] suggests the overburden at WIPP is $1481^{+8}_{-6} \text{ hg cm}^{-2}$. Other global fits [57, 62] yield consistent results.

7.3.1 Comparison with Previous Results

Esch et al. [58] report a vertical muon flux of $(3.10^{+0.05}_{-0.07}) \times 10^{-7} \text{ Hz cm}^{-2} \text{ sr}^{-1}$ underground at WIPP. To do so, they deployed at WIPP two parallel scintillator panels, oriented horizontally. Using these panels and subtracting off backgrounds, they measure a flux and then convert it to a vertical flux using the integral of eq. (7.2). They report the ratio to be $\Phi_v = (0.65 \pm 0.04)\Phi$. However, computing the ratio yields

$$\frac{\Phi_v}{\Phi} = \frac{2\pi \int_0^{\pi/2} \varphi(h, \theta) \cos \theta \sin \theta d\theta}{2\pi \int_0^{\pi/2} \varphi(h, \theta) \sin \theta d\theta} = 0.814^{+0.006}_{-0.005} \quad (7.10)$$

where $\varphi(h, \theta)$ is the distribution in eq. (7.2) and $h = 1526 \text{ hg cm}^{-2}$ (with the error due to varying this $\pm 10\%$) as in their paper. The $\sin \theta$ and 2π factors come from the integral over solid angle.

Furthermore, in their paper they use an efficiency of $(88.5 \pm 0.2)\%$ to calculate their flux, while later quoting an efficiency of $(84.25 \pm 0.2)\%$. A simple toy Monte Carlo simulation of their geometry with two parallel panels $76.2 \text{ cm} \times 305 \text{ cm}$ separated by 30.5 cm with muons generated according to eq. (7.2) yields a geometric efficiency of $(80.6^{+1.5}_{-1.6})\%$, with the main error due to uncertainty about which two surfaces are separated by 30.5 cm .

Using their background-subtracted muon count of 5224^{+89}_{-123} , their published area of $23\,225 \text{ cm}^2$, their published running time of $532\,800 \text{ s}$, and the corrected efficiency, their flux becomes $(5.24 \pm 0.10) \times 10^{-7} \text{ Hz cm}^{-2}$. Applying the correct vertical flux ratio, the vertical flux becomes $(4.26 \pm 0.09) \times 10^{-7} \text{ Hz cm}^{-2} \text{ sr}^{-1}$ which is consistent with the result above in eq. (7.9).

Chapter 8

Measuring Double Beta Decay

This is a reanalysis of the EXO-200 Run 2a data, taken between October 2011 and April 2012. This data has previously been used to establish a limit on the half-life of $0\nu\beta\beta$ of $T_{1/2}^{0\nu} > 1.5 \times 10^{25}$ yr and measure the half-life of $2\nu\beta\beta$ to be $T_{1/2}^{2\nu} = (2.23 \pm 0.017(\text{stat}) \pm 0.22(\text{sys})) \times 10^{21}$ yr [6]. This analysis attempts to make a more precise measurement of $T_{1/2}^{2\nu}$, and also to set a limit on the half-lives of Majoron-emitting modes $0\nu\beta\beta\chi^0(\chi^0)$.

There are numerous improvements over the previous analysis:

- Improved efficiency through identifying induction signals (section 5.2.2)
- Improved detector homogeneity through improvements to the electron lifetime and shielding grid corrections (section 5.3)
- A calibration that takes time variation into account (section 5.4)
- Improved simulations of backgrounds in detector materials (section 8.2)
- Incorporating position information into the signal and background models (section 8.1.4)
- An improved understanding of the fiducial volume (section 8.1.3)

The data is processed as described in chapter 5. A treatment of this processed data in order to measure $2\nu\beta\beta$ and attempt to measure $0\nu\beta\beta\chi^0(\chi^0)$ is described below.

8.1 Event Selection

8.1.1 Timing-based Veto

In order to reduce backgrounds due to cosmic-ray muons (described in section 7.1), two cuts are applied to the data:

- Events that occur within the 25 ms following a hit in a muon veto system panel are cut.
- Events that occur within the 60 s following a muon passing through the TPC (identified by the method in chapter 7) are cut. The muon events are cut themselves, as well.

These cuts are designed to remove as many cosmogenic background events as possible while not cutting so much detector live time that the trade-off in signal-to-background ratio is not worth it.

Events occurring near each other in time are likely to be due to a decay of some radioactive contaminant, followed by another decay of a short-lived daughter. For example, the ^{222}Rn daughter ^{214}Bi β^- decays to ^{214}Po , which then α decays with a 164 μs half-life. An event is vetoed if it occurs within $\pm 1\text{s}$ of another event.

Finally, conditions external to the detector may make the data unusable. Data from these time periods is vetoed. The most common cause is the mine evacuation siren, which is tested periodically. It is loud enough to create a large amount electronic noise through microphonic pickup. Time periods in which the siren is going off are vetoed.

The effects of all timing-based vetoes are summarized in table 8.1.

Table 8.1: A breakdown of the impact of the muon vetoes, TPC event-event coincidence veto, and the external conditions veto on the detector live time. Since multiple vetoes can be in effect at the same time, the combined impact of some vetoes is not simply their sum.

	Time (hr)	%
Vetoed Time	400.3	12.5
External Conditions	18.6	0.6
Physics Vetoes	381.8	11.9
Muons	163.5	5.1
TPC Muon	144.8	4.5
Panel Muon	19.6	0.6
Event-Event Coincidence	236.5	7.4
Live Time	2796.0	87.5
Total	3196.3	100.0

8.1.2 Other Vetoes

Scintillation-to-Ionization Ratio

As mentioned in section 3.3.3, α particle interactions produce more scintillation and less ionization in the liquid xenon than β and γ interactions. Therefore, events with a large scintillation-to-ionization ratio are removed from the data used for double beta decay physics. Since ^{222}Rn is a common source of α particles, these decays are later used in order to constrain backgrounds due to other isotopes in the ^{222}Rn decay chain.

Scintillation Timing

Events are vetoed if they contain more than one scintillation signal. This eliminates ^{214}Bi - ^{214}Po coincidences, as described above, and also eliminates events in which there could be ambiguity in associating ionization signals with multiple scintillation signals. Additionally, events are vetoed if the scintillation comes too late in the waveform for ionization to have time to drift the full TPC drift length. Since the waveforms are centered around the trigger time, this does not happen in typical data, but may happen in source calibration runs.

Full Reconstruction

Events are vetoed if they cannot be fully reconstructed. That is, all u wire signal bundles must be paired with v wire signal bundles as described in section 5.2.3. The resulting clusters must also be paired with a scintillation signal. This is necessary in order to know the 3D position of a cluster.

8.1.3 Fiducial Volume

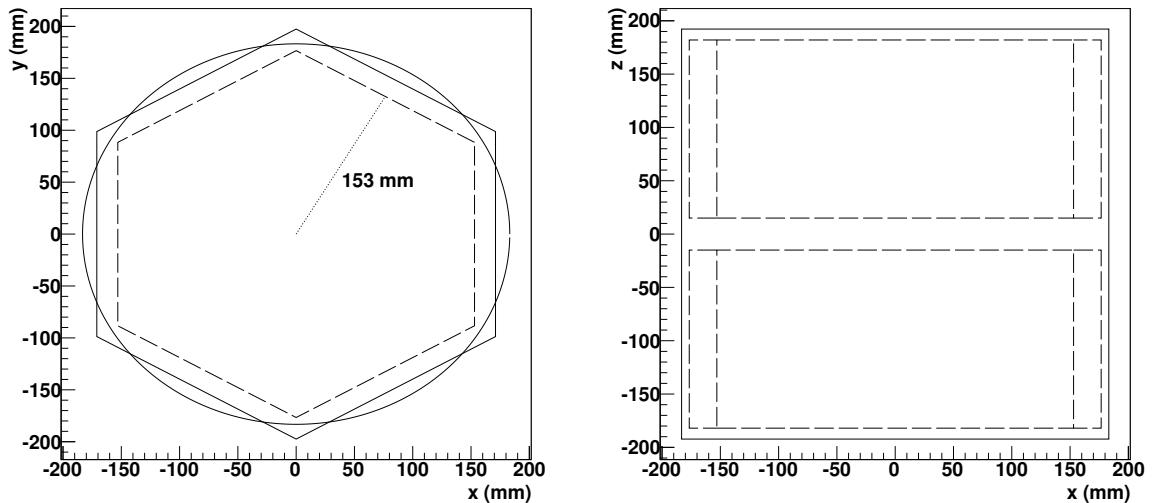


Figure 8.1: Projections of the hexagonal-prism-shaped fiducial volume. The dashed lines represent the fiducial volume. The solid lines are the anode wire plane and the teflon reflector. For the xz projection on the right, the two sets of vertical dashed lines are the maximum and minimum radial boundaries of the fiducial volume.

All the liquid xenon within the teflon reflectors and between the anodes is “active”. That is, ionization and scintillation in the active region are collected. However, only events within a fiducial volume are used in the analysis. This fiducial volume is a right hexagonal prism. The hexagon is coaxial with the detector and has apothem 153 mm. In each TPC, it begins 5 mm away from the cathode and extends to 182 mm (which is 10.2 mm from the v wire plane). Figure 8.1 provides an illustration. The total volume is 27.1 l, which contains 81.9 kg of liquid xenon.

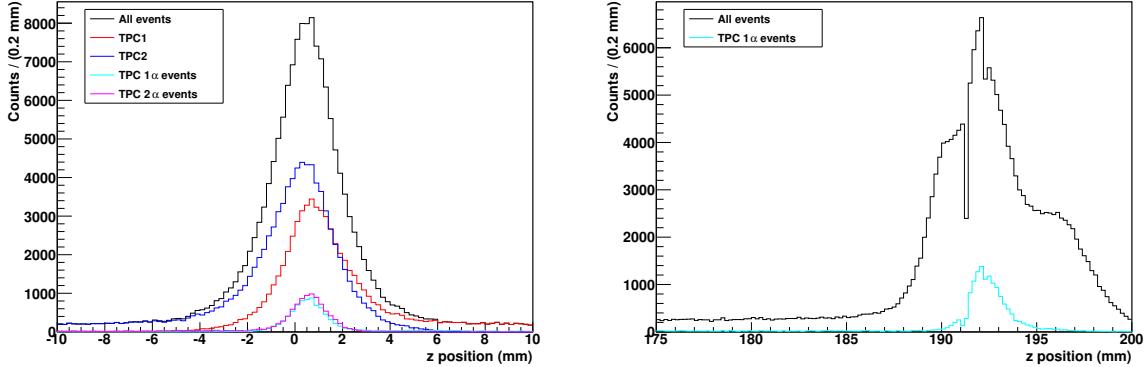


Figure 8.2: The choice of fiducial volume is partially motivated by the increased event rate seen near the cathode (at $z = 0$ mm) (shown left) and near the anodes (shown right; the v plane is at $|z| = 192$ mm and the u plane is at $|z| = 198$ mm). An increase is seen for both the total event rate, and for the rate of α particle interactions, which suggests the increased rate is due to backgrounds. Requiring $5 \text{ mm} < |z| < 182 \text{ mm}$ puts the fiducial volume well away from the volumes that show an increased rate.

This particular choice of fiducial volume has two chief motivations. The first is to reduce radioactive backgrounds from the detector materials. Figure 8.2 shows event rates near the cathode and anodes. There is nothing special about the xenon near these planes, and so the increased event rate must be due to backgrounds. The z cuts above eliminate the volumes that show higher-than-typical event rates.

The second motivation is to ensure a homogeneous detector response inside the fiducial volume. As shown in the previous discussion of the shielding grid correction in section 5.3.2, the correction becomes significant near $|z| = 182$ mm (see fig. 5.4). Even with the correction applied, the detector response begins to diverge near that position (see fig. 5.5), and so it is used as the boundary of the fiducial region. In the transverse dimension, a hexagonal cross-section is used. This matches the anode wire geometry, and so makes applying the fiducial volume cut straightforward. Figure 8.3 shows that requiring the apothem of this hexagon to be less than 153 mm ensures uniform energy response.

8.1.4 Quantities of Interest

Once an event passes the data quality cuts, the following quantities are calculated:

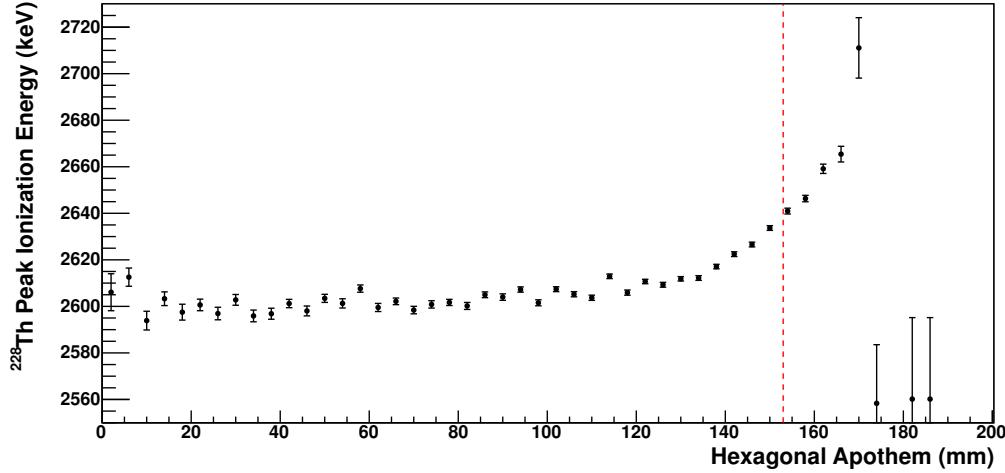


Figure 8.3: The reconstructed energy of the 2615 keV peak from a ^{228}Th source calibration. This is found by fitting a Gaussian + complimentary error function model to data with an increasingly large allowed fiducial volume. A larger allowed hexagonal apothem extends the fiducial volume closer and closer to the edges of the anode wire planes. The detector response begins to deviate significantly when the cut is extended above 153 mm, and so this is used to define the fiducial volume.

1. The multiplicity (single site or multiple site), using the criteria described in section 5.2.4.
2. The event energy using the calibration obtained as described in section 5.4.
3. The “standoff distance”. The radial distance of the cluster to the teflon reflector is calculated. So is the z distance to the v wire plane. The standoff distance is the minimum of these two distances. If multiple clusters are involved in the event, the minimum standoff distance of all clusters is used.

The single site and multiple site information helps separate γ interactions from β interactions. The energy spectra of different signals and backgrounds are different. The standoff distance describes the spatial distribution of events. For example, $2\nu\beta\beta$ events occur uniformly in the detector, and so their standoff distance distribution should extend to high values. External gamma ray backgrounds, meanwhile, are attenuated by the xenon, and so their standoff distance distributions fall off more quickly.

8.2 Monte Carlo Simulations of Signals and Backgrounds

8.2.1 Simulations

A detailed simulation of the EXO-200 detector and its surrounding (see fig. 8.4) known as EXOSim has been created using the GEANT4 simulation toolkit [56]. EXOSim is used to determine both the energy spectra and spatial distributions of various double beta decay signals and radioactive backgrounds in EXO-200. The following are simulated:

- Double beta decay of ^{136}Xe : $2\nu\beta\beta$, $0\nu\beta\beta$, and $0\nu\beta\beta\chi^0(\chi^0)$ for spectral indices 1, 2, 3, and 7
- Potential backgrounds in the liquid xenon: ^{135}Xe , and ^{222}Rn
- Potential backgrounds in the detector copper: ^{40}K , ^{54}Mn , ^{60}Co , ^{65}Zn , ^{238}U , and ^{232}Th
- ^{222}Rn in the air gap between the cryostat and lead wall

For double beta decay and backgrounds in the liquid xenon, the events are simulated uniformly in the volume of active xenon. All double beta decay modes are simulated using the Fermi function by Schenter and Vogel [63]. For ^{222}Rn , there are additional complications, however. Decays of radon and daughters in the inactive xenon can emit γ rays that also deposit energy in the active volume. Therefore, these are simulated as well. Additionally, charged radon daughters can drift to the cathode. ^{214}Bi emits γ rays when it decays. The daughter ^{214}Po has a $164\text{ }\mu\text{s}$ half-life, and so both are typically eliminated by the event-event coincidence and multiple scintillation cuts. However, on the cathode, the α particle from the ^{214}Po decay may go into the cathode and not be seen. Therefore, ^{214}Bi on the cathode is simulated separately.

For backgrounds in the detector copper, the TPC vessel, all internal components, the legs, and the high-voltage feedthrough are simulated as a whole. These are all modeled in the simulation, as shown in fig. 8.4. It is assumed that the impurities are

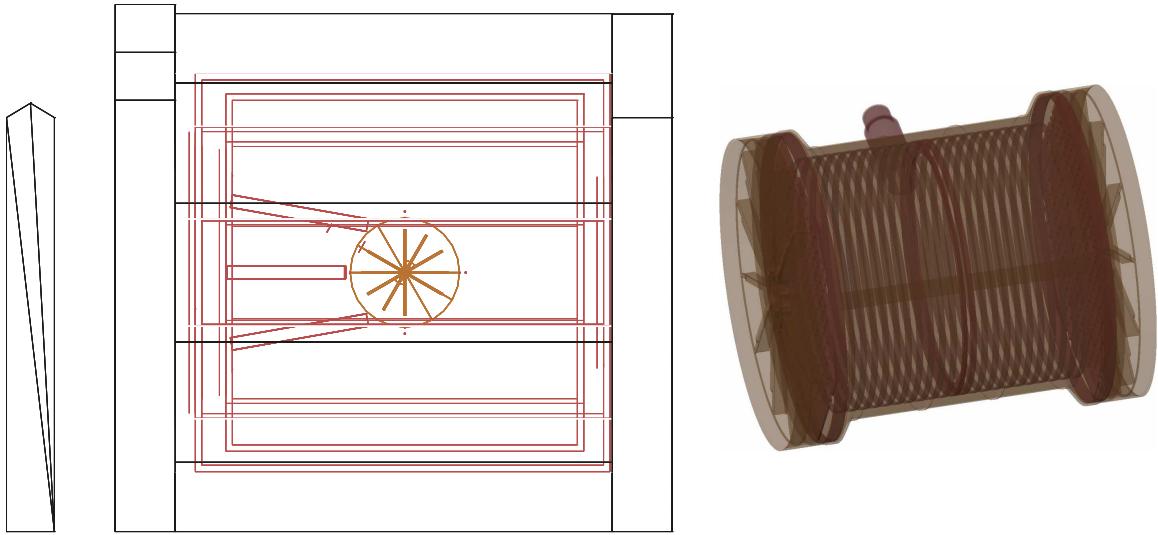


Figure 8.4: EXO-200 as implemented in GEANT4 for simulations. The geometry incorporates the lead shielding, cryostat, HFE, and detector (shown left). The detector itself (shown right) is quite detailed in order to accurately simulate backgrounds from the detector materials.

uniformly distributed according to copper mass. Even if some of the contamination is distributed according to surface area, the thin-walled construction of EXO-200 makes this a close approximation.

8.2.2 PDF Generation

For a single event, the GEANT4 simulation provides a list of energy deposits in the liquid xenon, their positions, and their times. These must then be converted into simulations of the signals actually seen by the detector. This process begins by estimating the amount of drifting ionization by using the 15.6 eV W -value (see section 3.3.1) and the observation that $\sim 80\%$ of ionization drifts instead of recombining in strong electric fields [64]. α particles have their ionization yield quenched by a factor of 0.055 [65]. This ionization is drifted to the anodes in small time steps, and the signals on the wires are determined using the Shockley-Ramo theorem using a 2D Maxwell [51] simulation of the electric fields in the detector.

The energy that doesn't go into drifting ionization goes into scintillation. The

anticorrelation between ionization and scintillation is not modeled, and so the fraction of energy in scintillation is constant. The scintillation photons are distributed across the two planes in a parameterized way. All photoelectrons from the APDs are assumed to arrive instantaneously since the scintillation signal arrives in ns, compared to the μ s digitization time.

White noise is added to the ionization and scintillation signals to yield a similar signal-to-noise ratio as in the detector. The signals are then run through simulations of the shapers found in the detector electronics, and are digitized at 1 MHz, just as the real signals. The matched filter and multiple signal finder are applied to identify signals, and the signals are bundled and clustered. The process is identical to the one described in section 5.2. This simulates the efficiency and energy threshold effects that come from signal-finding and reconstruction, as well as the position resolution.

The final step is to form distributions in multiplicity, energy, and standoff distance that can be fit to the data collected by the detector. The multiplicity and standoff distance can be calculated directly from the reconstruction described above. For the energy, however, the true energy from the Monte Carlo simulation is used. For each cluster found by the reconstruction process, the energy deposits that arrived on the wires associated with the cluster at the appropriate time are summed to get the cluster's true energy.

In order to simulate the detector energy resolution, the true energies are smeared by a parameterized resolution function given by

$$\frac{\sigma(E)}{E} = \sqrt{\frac{p_0^2}{E^2} + \frac{p_1^2}{E} + p_2^2} \quad (8.1)$$

where p_0 is a constant, noise-like smearing, p_1 is a smearing due to counting statistics, and p_2 is an energy-dependent smearing that can be due to drifting gains or other inhomogeneities. These parameters are determined for the entire data set by fitting a simulated calibration source spectrum to the calibration data collected by the detector during the entire time period. The values of the parameters are shown in table 8.2, and the resolution as a function of energy is shown in fig. 8.5.

The true energies from the signal and background Monte Carlo simulations are

Table 8.2: The values for the energy resolution parameters used in this analysis for both single site (SS) and multiple site (MS) events. The parameters are defined in eq. (8.1).

Multiplicity	p_0 (keV)	p_1 ($\sqrt{\text{keV}}$)	p_2
SS	39.2	7.0×10^{-4}	0.01
MS	40.9	2.8×10^{-4}	0.01

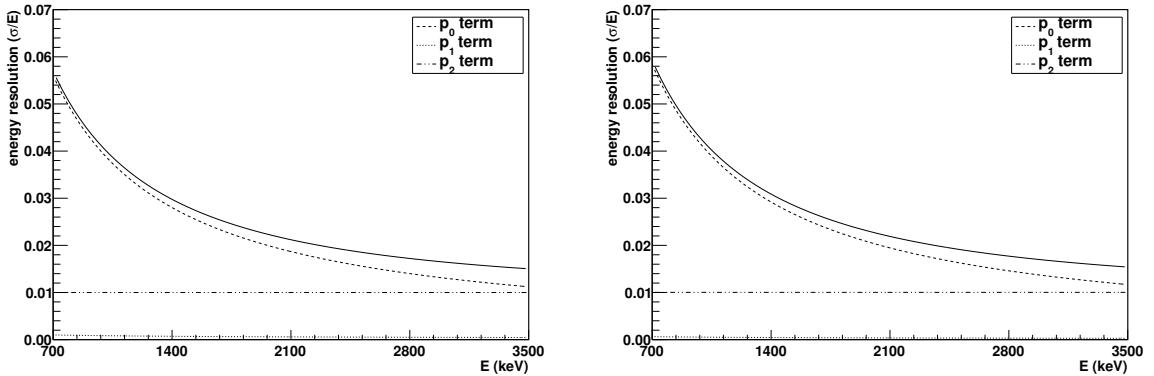


Figure 8.5: The parameterized energy resolution as a function of energy for single site (left) and multiple site (right) events. The dark curve is the resolution, while the various dashed lines show the contributions of the individual terms. This is determined by finding the set of parameters that best match a simulated spectrum to the spectrum obtained in calibration runs.

smeared according to the parameterization in eq. (8.1). For each simulation, two probability density functions (PDFs) are formed in energy and standoff distance: one for single site events, and one for multiple site events. Figure 8.6 shows an example of two smeared PDFs. For each simulation, the efficiency for a simulated event to make it into the final PDFs is recorded, as well as the fractional split between single site and multiple site events.

8.2.3 Agreement between Simulation and Data

In order to verify that the simulations accurately predict multiplicity, energy, and standoff distance distributions, they are cross-checked by simulating the calibration sources and comparing the results with calibration data. This is done for both the ^{60}Co source and the ^{228}Th source.

Single Site Fraction Agreement

Since the distinction between single site and multiple site events plays a strong role in estimating backgrounds, it is important to confirm that the simulations accurately reflect the fraction of events of each multiplicity. Figure 8.7 shows that the maximum observed deviation in the single site fraction is less than 10 % and varies with energy. The systematic error due to the energy-dependent deviation is evaluated with the shape agreement below. The remaining energy independent error on the overall fraction is 5.9 %. This is determined by taking the mean of the measured residuals, weighted by the fraction of the $2\nu\beta\beta$ spectrum in each energy bin. The weighting is done so that the error reflects the portion of the spectrum where the most events are collected.

Shape Agreement

One simple check is to compare the shapes of the energy and standoff distance distributions in data and simulation. Figure 8.8 shows this comparison. The overall agreement in both shape and standoff distance looks reasonable. However, this agreement is not perfect. As fig. 8.9 shows, the simulations underpredict the fraction of single

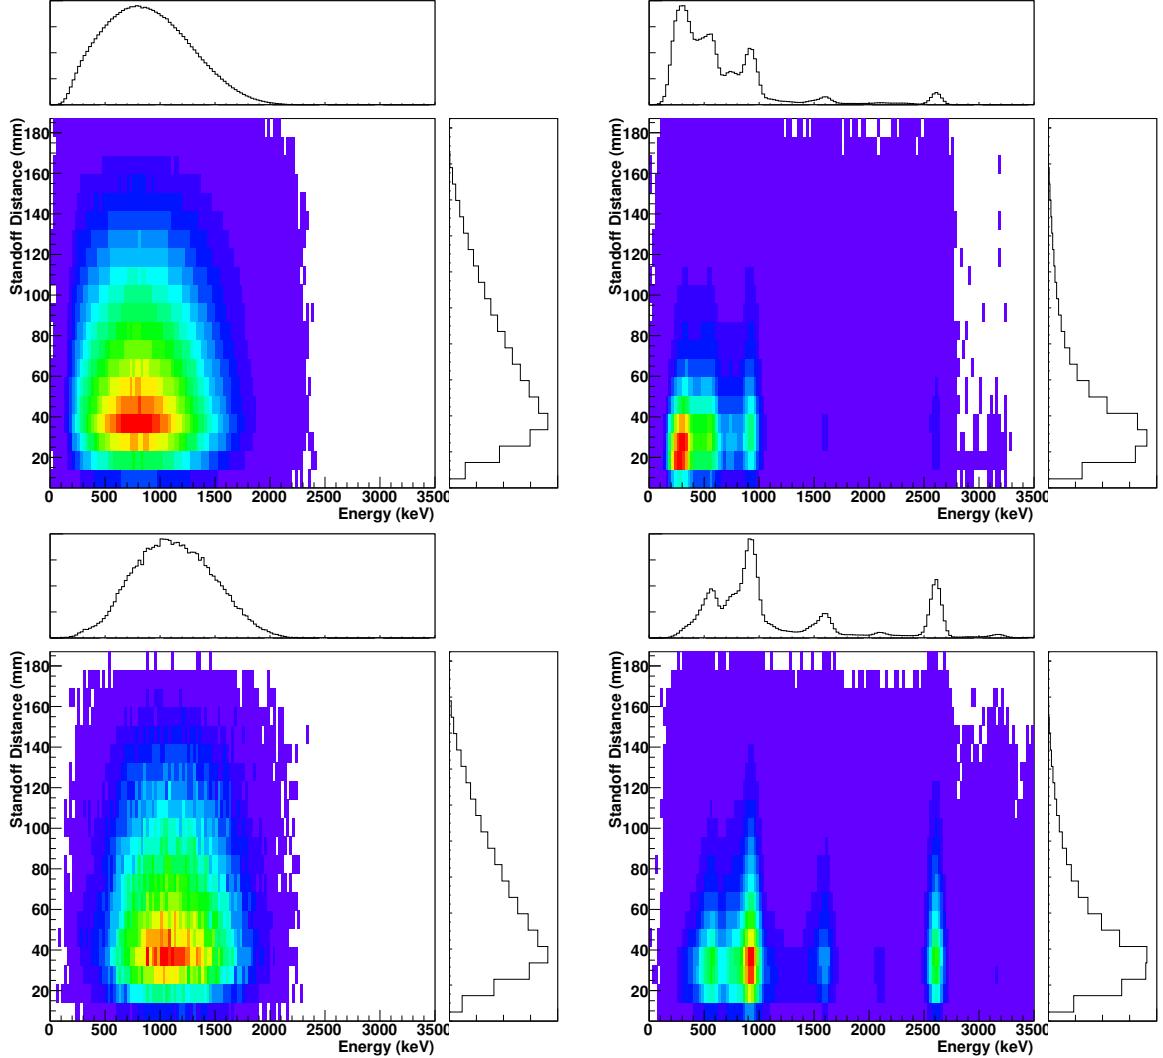


Figure 8.6: Two-dimensional PDFs generated by the Monte Carlo simulations described in the text. On the left is the spectrum for $2\nu\beta\beta$, and on the right is the spectrum for ^{232}Th in the copper of the TPC vessel. On top are the single site spectra, while multiple site spectra are on the bottom. The vertical axis in all plots is the standoff distance, while the horizontal axis is energy. 1D projections in energy and standoff distance are also shown.

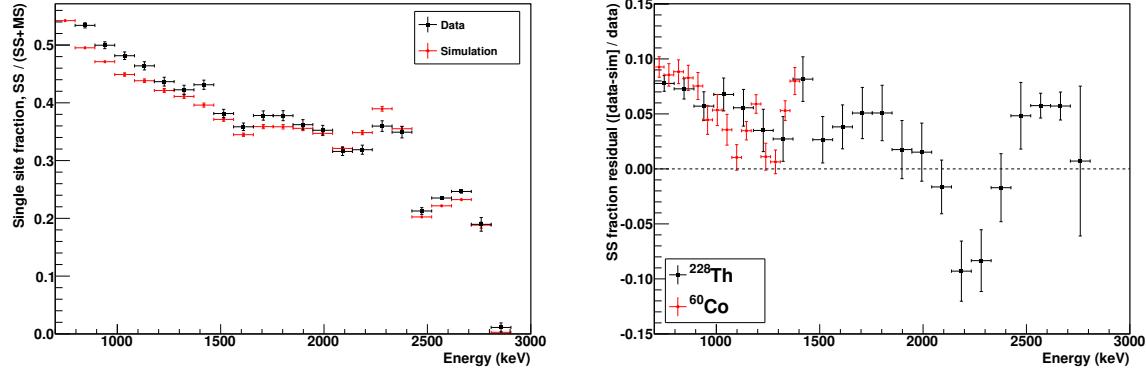


Figure 8.7: The left plot shows the agreement between the fraction of events that are single site in simulations and data for a ^{228}Th calibration. The right plot shows the fractional residual error between simulation and data for ^{60}Co and ^{228}Th data.

site events at low energies. This is due to a slight difference in efficiency for detecting low-energy clusters between data and simulations. This results in some events being shifted from single site to multiple site in simulations. This is not corrected for, but the systematic effect is taken into account, as described in section 8.3.1.

Rate Agreement

Table 8.3 shows the agreement between the true activity of calibration sources (based on their activity at time of manufacture and known half-life) and the rate estimated using data from the detector and Monte Carlo simulations of the sources in order to determine efficiency. For most isotope and position combinations, the error is less than 5%. Taking a mean (weighted by the uncertainties) of all the residuals yields $(\text{True} - \text{Estimated})/\text{True} = (1.7 \pm 2.3)\%$. Contributions to this error include several effects that may not be present in non-calibration data. The source activities are only known to within 1.2%, and the rate estimate is quite sensitive to the actual location of the source in the simulations.

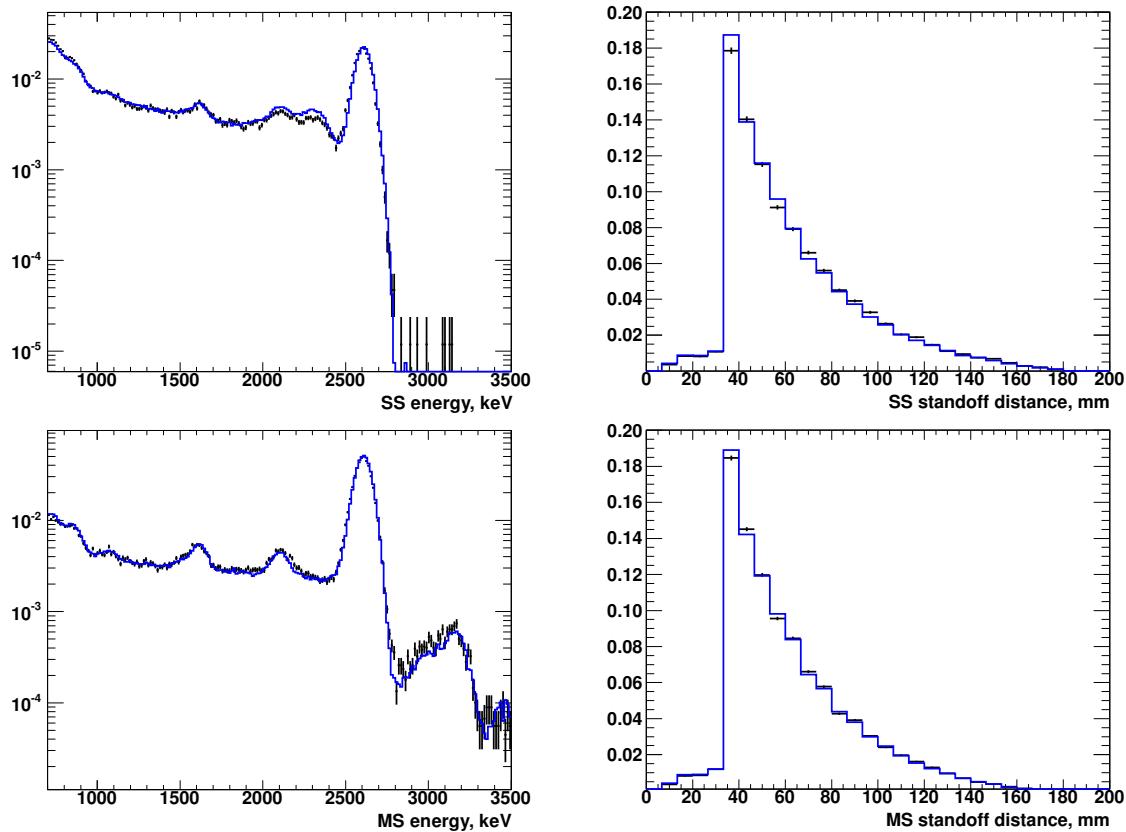


Figure 8.8: The agreement between simulations (in blue) and data (in black) for a ^{228}Th calibration source located at the cathode. The projection of the energy spectra is shown on the left, and the projection of the standoff distance distribution is shown on the right. The top row is for single site events, while the bottom row is for multiple site events.

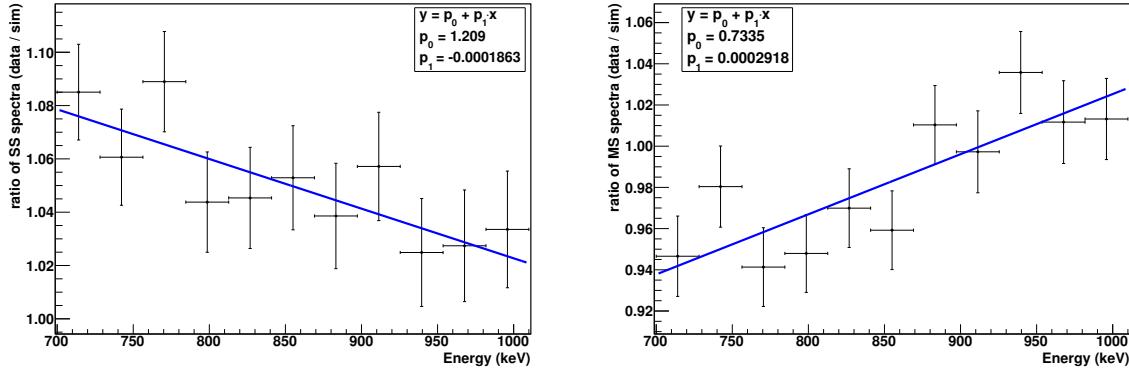


Figure 8.9: The ratio of the simulated energy spectrum to data for a ^{60}Co source. Single site events are shown on left, and multiple site events are shown on right. This indicates that the simulation underpredicts the fraction of single site events at low energies.

Table 8.3: Comparison of the estimated calibration source activities based on measurements in the EXO-200 detector and Monte Carlo simulations of efficiencies to their true values, based on the activity at the time of manufacture and known half-life.

		Absolute rate agreement	
Location	Isotope	(True - Estimated)/True (%)	
Anode	$-z$	^{228}Th	$6.2^{+0.2}_{-1.7}$
		^{60}Co	$6.0^{+0.4}_{-1.7}$
Anode	$+z$	^{228}Th	$-0.4^{+0.6}_{-0.5}$
		^{60}Co	$2.3^{+0.5}_{-0.8}$
Cathode	$+x$	^{228}Th	$3.6^{+0.7}_{-1.4}$
		^{60}Co	$0.9^{+1.0}_{-1.5}$
Cathode	$+y$	^{228}Th	$4.5^{+3.3}_{-3.2}$
		^{60}Co	$5.1^{+3.2}_{-4.0}$

8.3 Systematic Errors

8.3.1 Rate Uncertainty

While the rate agreement with calibration sources provides an estimate of the accuracy to which the detector can measure decay rates, it is possible to estimate this uncertainty directly. Table 8.4 summarizes these estimates, which are described below.

Noise Events

Noise events (as identified by the noise tagger described in appendix A) may occur in coincidence with real events. By simply comparing the 75 mHz rate for real events and the 11 mHz rates for noise events, the probability for one to occur in the 1024 μ s following a trigger by the other, thus removing a real event from the data set, is 8.9×10^{-5} . Likewise, hand-scanning events tagged as noise puts an upper limit on the rate at which legitimate physics events are tagged as noise at $< 5.8 \times 10^{-4}$.

Only One Scintillation Signal Requirement

Noise fluctuations in the APD waveforms might be reconstructed as scintillation signals. These events would be rejected by the requirement that events have only one scintillation signal. About 0.75 % of events in real data are removed by this cut, while $\sim 4\%$ of events in Monte Carlo simulations of $2\nu\beta\beta$ events are removed. Since the real data consists chiefly of $2\nu\beta\beta$ events, and since there should be some genuine coincides in that data, too many simulated events are failing this cut. This is attributed to the matched filter signal finding technique (section 5.2.1) being optimized to work on APD signals with real noise, when the Monte Carlo simulations only have simulated white noise. When determining the efficiency for double beta events to survive all cuts, all of these false coincidences could survive the other cuts, or all could fail. Therefore, the efficiency used when computing a half-life is taken midway between these extremes, and half the difference between the extremes is propagated as a systematic error. Meanwhile, compared to this, the systematic error due to the

Table 8.4: Estimated contributions to the uncertainty for rates measured in the detector. Some contributions are energy independent, and so are calculated separately for the different modes.

Component	Contribution to rate uncertainty (%)				
	$2\nu\beta\beta$	$0\nu\beta\beta\chi^0(\chi^0)$	$n = 1$	$n = 2$	$n = 3$
Tagged noise in coincidence with physics events				1×10^{-3}	
Physics events mis-identified as noise					< 0.06
Falsely finding a second scintillation signal	2.2	1.6	1.9	2.1	2.1
Fiducial volume due to bias in position reconstruction					1.73
Fiducial volume due to position reconstruction resolution					0.36
Energy calibration and resolution uncertainty	0.36	0.04	0.09	0.36	0.56
Energy scale difference between β and γ	0.17	0.09	0.20	0.37	1.40
Induction tagging	< 0.2	< 0.66	< 0.44	< 0.21	< 0.28
Triggering efficiency (above 700 keV)				≈ 0	
Signal-finding efficiency (above 700 keV)				≈ 0	
Clustering efficiency				2.0	
Shape distortion in 700 keV to 1000 keV				0.75	
Xenon isotopic fraction				0.6	
Total	3.6	3.3	3.4	3.5	3.8

one scintillation signal cut removing truly coincident events is negligible due to the overall low rate of events.

Fiducial Volume

Monte Carlo simulations of $2\nu\beta\beta$ events suggest that the RMS error on the u , v , and x coordinates that are used to define hexagonal cross-section of the fiducial volume are 2.4 mm, 1.2 mm, and 2.7 mm, respectively. The error on the z coordinate due to slight deviations from a uniform electric field is 0.42 mm. Meanwhile, studies of events that occur on the cathode and so create signals in both TPCs suggest there might be a 1.5 mm offset in u and v and a 0.5 mm offset in z between the two TPCs. These might be physical, due to slight asymmetries in the construction, but are taken as systematic errors. Combining these uncertainties gives a total 1.73 % uncertainty on the fiducial volume. Furthermore, position resolution considerations mean some events inside the fiducial volume can be reconstructed outside, and vice versa. Assuming uniform events, this should be a +0.36 % effect on the rate. It is, however, taken as a symmetric systematic error and is not corrected.

Energy Scale

The energy range used for the fit has a 700 keV lower limit, since the trigger and reconstruction are known to be 100 % efficient above that energy. Uncertainty on the energy calibration means some events might mistakenly be left out of the data set used in the fit. At 700 keV, the energy uncertainty is 2.5 keV. Comparing the number of events in the spectrum within (700.0 ± 2.5) keV to the total number of events above 700 keV gives an estimate of this uncertainty's effect. Since it depends on the shape of the spectrum under consideration, it is computed separately for the different double beta decay modes.

As fig. 8.10 suggests, there may be a systematic shift in the energy scale for β events compared to γ events. This has been studied. The 2615 keV γ ray from ^{228}Th may produce an electron-positron pair that should be topologically similar to a double beta decay event. Using the 511 keV γ rays produced by positron annihilation

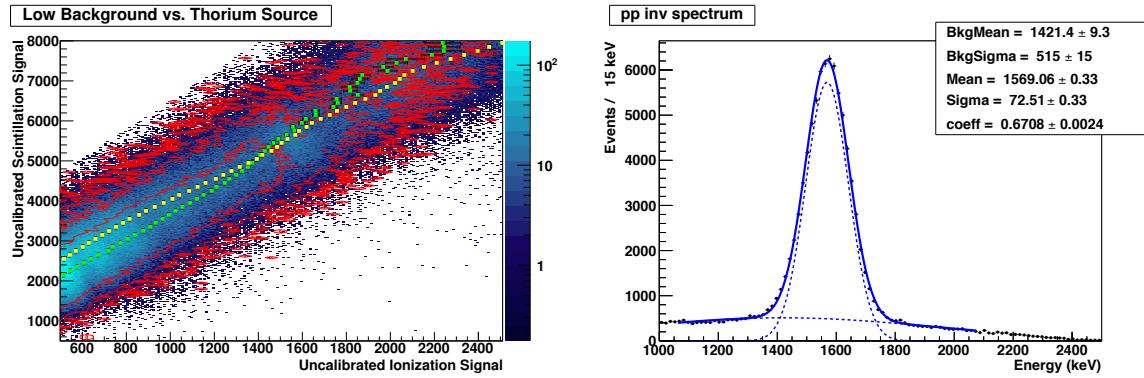


Figure 8.10: (Left) A comparison of the mean ionization signal for a given amount of scintillation in low background (green dots and red contours) data, compared to that for thorium calibration (yellow dots and blue shading). The low background data, consisting chiefly of β events, shows larger ionization signals than the thorium data, consisting chiefly of γ events. (Right) After applying the energy calibration (based on γ ray emitting sources) the pair production peak of a 2615 keV, which is a β type event, reconstructs below its expected 1593 keV energy.

as a tag gives a very clean sample of these events. As shown in fig. 8.10, there appears to be a systematic offset in the amount of ionization produced by these events compared to what calibrations from γ sources suggest. Propagating this difference through all rotations and calibrations suggests that β events should have their energies multiplied by a factor 1.0115 ± 0.0052 in order to match the calibration from γ events. Meanwhile, generating PDFs for different energy scale differences and determining which best fit the data (as described below) suggests the best value for this factor is 1.001 ± 0.003 . Since these are inconsistent, the uncertainties are enlarged. For $2\nu\beta\beta$ the effect on the rate uncertainty is done by performing the fit at 1.001 ± 0.011 and observing the change in the number of events found. For the Majoron-emitting modes, this is instead determined similarly to the energy scale uncertainty, by calculating how many events pass above and below the 700 keV threshold as the factor is varied between 1.001 ± 0.011 .

Cluster Reconstruction

The clustering efficiency uncertainty is determined by comparing the number of events that pass all cuts except the one requiring all charge bundles to be fully clustered to the number of events that pass all cuts. This is done in both simulations of $2\nu\beta\beta$ and real data. Since there are backgrounds in the real data which may have different rates, background subtraction is done using the best fit to the data (described below in section 8.6). The difference between background-subtracted data and simulation is energy dependent, and so integrating over the $2\nu\beta\beta$ spectrum gives the uncertainty listed in table 8.3.

Shape Disagreement

Simulations predict a shape for the spectrum between about 700 keV to 1000 keV that disagrees with calibration data. This shape error is similar for ^{60}Co and ^{228}Th , and can be parameterized as shown in fig. 8.9. This could be used as a correction, but instead it is taken as a systematic uncertainty. All PDFs are “skewed” to effectively undo this error, and then used to generate toy Monte Carlo simulations with similar amounts of signal and background as in the Run 2a data. The resulting toy sets are then fit with the original, unskewed PDFs. The result is that the number of $2\nu\beta\beta$ events returned by the fit is systematically 0.75 % higher than the number generated in the simulation. Therefore, this is the contribution to the $2\nu\beta\beta$ rate uncertainty. The Majoron-emitting modes with spectral indices 1, 2, and 3 have less of their spectrum at these low energies, and so using the $2\nu\beta\beta$ uncertainty is a conservative estimate. For spectral index 7, the sign of the distortion would tend to cause the number of these decays to fit to higher than their true values. Thus, if a limit is obtained for spectral index 7, it will still be conservative, even using the $2\nu\beta\beta$ uncertainty.

Number of ^{136}Xe Atoms

Finally, there is uncertainty as to precisely how many ^{136}Xe atoms are in the fiducial volume. The uncertainty on the enrichment fraction of $(80.6 \pm 0.2)\%$, for example, contributes a 0.3 % uncertainty. Isotope fractionation between the liquid and gas

phases is small, contributing less than 0.01 % uncertainty. Some natural xenon from a commissioning run may have remained dissolved in plastic TPC parts, possibly diluting ^{136}Xe by up to 0.04 %. The dominant contribution comes from uncertainty about the density. Allowing for a ± 2.5 % uncertainty in xenon temperature changes the density by ± 0.6 %.

8.3.2 Other PDFs

Other Backgrounds

A number of radioactive isotopes are only simulated in the TPC vessel. However, since the cryostat, for example, is made of the same copper, it seems reasonable to include simulations for it as well. And, likewise, other sources of background can be imagined. In the Run 2a data set, however, there are not enough events to statistically distinguish backgrounds in different components. Figure 8.11 shows projections of the energy and standoff distance spectra for ^{238}U in the TPC vessel and the inner cryostat, normalized to the number of ^{238}U decays returned by the fit. Statistical errors are shown. The two spectra are statistically indistinguishable.

Neutron-induced backgrounds are difficult to model. The most likely candidate background, due the large volume of HFE and its close proximity to the TPC, is the 2.2 MeV γ ray produced by neutron capture on hydrogen. Using the muon flux measured in chapter 7 and FLUKA [66] simulations of the EXO-200 detector, a conservative estimate for the number of 2.2 MeV γ rays seen in the data is at most around 40, split evenly between the single-site and multiple-site spectra. Attempting to include the γ line in the fit (described below) yields a best fit of 23 events, but not significant at even the 1σ level. Therefore, this potential background is not included in the fit, and will only be a small effect if present.

Missing Wire Channel

For a period of roughly 1 month, the data collected from the detector did not include one collection wire (numbered channel 16). In order to determine the effect of this period on the overall spectrum, PDFs were constructed for data with and without

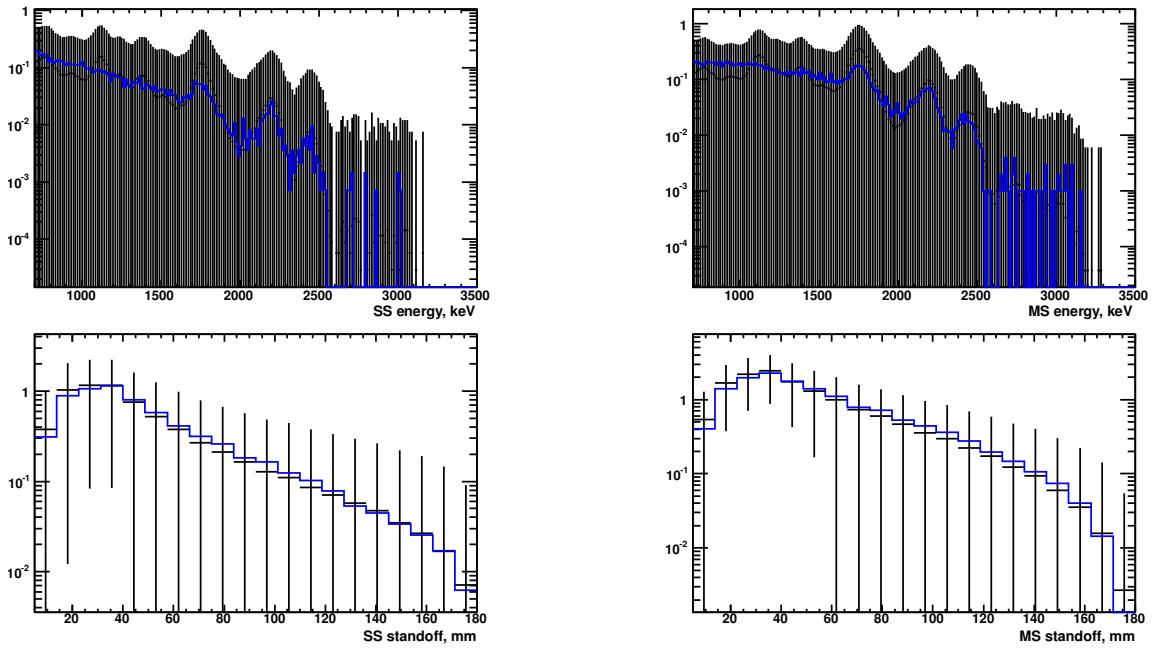


Figure 8.11: Projections of the energy (top) and standoff distance (bottom) spectra for ^{238}U in the TPC vessel (black bands) and in the inner cryostat (blue line). Single site events are on the left, and multiple site events are on the right. The error bars are statistical, normalized to the number of ^{238}U decays returned by the fit. The two spectra are statistically indistinguishable without more data.

this channel included. The PDFs were combined, weighted by their fraction of total live time in Run 2a. The period with the missing channel constitutes 5.2 % of the total time. These weighted PDFs were compared with the default PDFs with all channels working. The reduced χ^2 of 0.1 shows the spectra are quite compatible. A Kolomgorov-Smirnov test reports 89.7 % probability for compatibility. Therefore, the default spectra are used, since they closely match the weighted spectra, which are more complicated to produce.

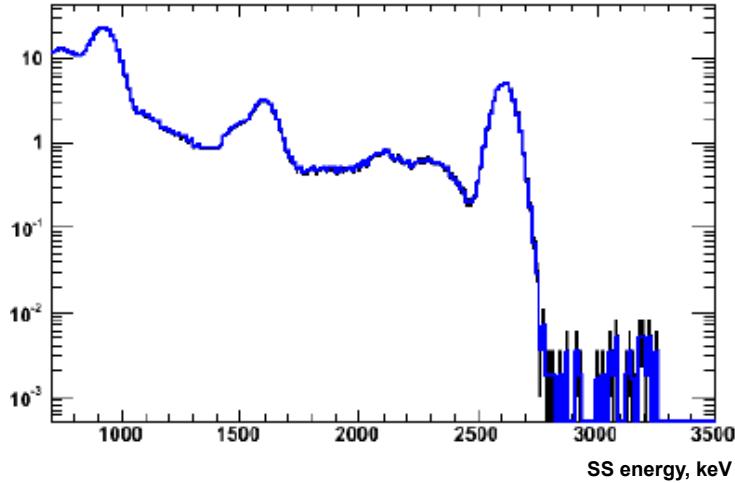


Figure 8.12: A comparison of the default PDF for ^{232}Th in the TPC vessel (black) to one that accounts for the time period in which one collection wire was not taking data (blue). The two spectra are quite compatible, with a reduced χ^2 of 0.1, and a Kolmogorov-Smirnov test indicating 89.7 % probability for compatibility.

8.4 Maximum Likelihood Method

8.4.1 Theory

When attempting to measure a decay rate (as for double beta decay), the parameter measured is the number of decays in a given amount of time. This is complicated

by the presence of backgrounds. Ultimately, the collected data is a combination of many backgrounds and possibly signals, each parameterized by a number of decays. A maximum likelihood analysis provides a way to estimate these parameters and their uncertainties, or to place limits on the parameters.

In a maximum likelihood analysis, first a likelihood function \mathcal{L} is established:

$$\mathcal{L} = \prod_i f(x_i|\boldsymbol{\theta}) \quad (8.2)$$

where $f(x_i|\boldsymbol{\theta})$ is a probability density function with parameters $\boldsymbol{\theta}$ for the observed data x_i . The data are fixed, while the parameters can vary. Thus \mathcal{L} is the likelihood that the parameters $\boldsymbol{\theta}$ describe the probability distribution underlying the data.

Finding the set of parameters that best describe the data, then, becomes a matter of maximizing \mathcal{L} . In practice, since \mathcal{L} is always less than one and can be quite small for some choices of parameters, it is easier to work with the negative logarithm of the likelihood function, or “negative log-likelihood”. Maximizing \mathcal{L} becomes minimizing the negative log-likelihood. This can be done analytically for simple cases, but for most problems a numerical minimization package such as MINUIT [67] is used.

The maximum likelihood method can be used to construct confidence intervals. Define

$$\lambda(\theta_0) = \ln \mathcal{L}_{max}(\boldsymbol{\theta}) - \ln \mathcal{L}_{max}(\theta_0, \boldsymbol{\theta}') \quad (8.3)$$

where $\mathcal{L}_{max}(\boldsymbol{\theta})$ denotes the likelihood maximized over all parameters. $\mathcal{L}_{max}(\theta_0, \boldsymbol{\theta}')$ denotes the likelihood maximized over all parameters except for θ_0 , which is held fixed. Wilks’ theorem [68] says that the distribution of λ can be approximated by

$$2\lambda(\theta_0) \sim \chi^2_{\dim(\theta_0)} \quad (8.4)$$

where $\chi^2_{\dim(\theta_0)}$ is the χ^2 distribution with degrees of freedom equal to the dimensionality of θ_0 . Then the $\alpha\%$ level confidence interval on θ_0 is bounded by θ_0^- and θ_0^+ such that $2\lambda(\theta_0^\pm)$ is equal to the χ^2 quantile for $\dim(\theta_0)$ degrees of freedom at the $\alpha\%$ level.

It may happen, however, that when measuring physical quantities, the best fit

parameters or confidence intervals may be unphysical. A radioactive isotope can not have a negative number of decays, for example. Rolke et al. [69] propose a method to extract limits in these situations. Suppose that θ_0 cannot be negative, but the best fit has $\theta_0 < 0$. Their method of bounded likelihood sets $\theta_0 = 0$, and the limit on θ_0 is taken using

$$\lambda'(\theta_0) = \ln \mathcal{L}(\theta_0 = 0, \boldsymbol{\theta}) - \ln \mathcal{L}_{max}(\theta_0, \boldsymbol{\theta}') \quad (8.5)$$

and the relationship in eq. (8.4).

When employing the method of bounded likelihood, it is necessary to verify that the method provides coverage. A nominal CL% confidence interval for a parameter should contain the true value of the parameter at least CL% of the time. This can be studied through Monte Carlo simulations. Namely, the experiment is simulated some large number of times for a range of parameter values. For each simulation, the confidence interval is obtained using the method of bounded likelihood. The fraction of the time the true parameter value lies within the confidence interval can then be computed. The method is said to “overcover” if a CL% confidence interval contains the true value more than CL% of the time. The opposite is to “undercover”. In this case, the approximation in eq. (8.4) does not hold well, and so the value of λ' used to form the confidence interval must be adjusted.

8.4.2 Application to EXO-200

For each signal and each background, two dimensional PDFs (in energy and standoff distance) are generated for single site and multiple site multiplicities as described in section 8.2. These PDFs are generated for $700 \text{ keV} \leq E < 3500 \text{ keV}$ and $0 \text{ mm} \leq d_s < 200 \text{ mm}$, where d_s is the standoff distance. There are 200 evenly-spaced (14 keV) bins in energy, and 20 evenly-spaced (10 mm) bins in standoff distance.

For each PDF, there are two parameters: the total number of events N_i associated with that PDF, and the fraction a_i of those events that are single site events. There is also one overall normalization parameter c_{norm} , explained below. Denote the number of events in the bin with multiplicity (SS or MS) m , energy E , and standoff distance

d_s for PDF i as $n_i(m, E, d_s)$. The combined PDF predicts

$$n(m, E, d_s) = c_{\text{norm}} \sum_i N_i a_i n_i(m, E, d_s) \quad (8.6)$$

events in that bin.

Suppose a bin contains k events in the data. The combination PDF predicts $n(c_{\text{norm}}, \mathbf{N}, \mathbf{a})$ events in that bin. The likelihood function for that individual bin is then given by the Poisson distribution:

$$f(k|c_{\text{norm}}, \mathbf{N}, \mathbf{a}) = \frac{[n(c_{\text{norm}}, \mathbf{N}, \mathbf{a})]^k \exp[-n(c_{\text{norm}}, \mathbf{N}, \mathbf{a})]}{k!} \quad (8.7)$$

The total likelihood function that is maximized is the product of the likelihood functions for all bins.

8.5 Constraints

8.5.1 Theory

The maximum likelihood method also admits a simple way to constrain parameters. Suppose a parameter θ is measured to have value $\theta_0 \pm \sigma$. Then one way to constrain θ is to multiply the likelihood function by

$$f(\theta) = \exp \left[\frac{-(\theta - \theta_0)^2}{2\sigma^2} \right] \quad (8.8)$$

That is, the likelihood for a given θ is Gaussian distributed around the measurement θ_0 , with the width given by the uncertainty σ in the measurement. In this case, the normalization factors for the Gaussian have been omitted, since they just scale the likelihood function without affecting the location of the maximum. This is just an example, and any likelihood function can in principle be used to constrain a parameter or parameters.

8.5.2 Single Site Fraction

As described in section 8.2.3, the overall systematic error on the single site fraction is estimated to be 5.9 %, based on comparisons of calibration data to simulations. Therefore, the single site fraction for each component in the fit is assigned a Gaussian constraint (as in eq. (8.8)). The mean of this Gaussian is the single site fraction from the simulation, and the width is 5.9 % of the mean. The constraint on each component is assumed independent of every other component. While some correlation should be expected, it has not been studied, and independence provides a conservative assumption.

8.5.3 Radon in the Xenon

Since radon is noble, it cannot be chemically purified from the xenon. ^{222}Rn and some daughters emit α particles when decaying, and these can be identified through their increased scintillation and decreased ionization signals. Counting these decays reveals there is $(3.65 \pm 0.37) \mu\text{Bq kg}^{-1}$ of ^{222}Rn dissolved in the liquid xenon in the TPC. This is used to constrain the contributions of radon in the xenon in the fit to the spectrum:

- The active volume of xenon has a total mass of 126.1 kg, so there is a total activity of $(460.1 \pm 46.0) \mu\text{Bq}$ in the active xenon.
- Of the ^{214}Po α decays observed, roughly 75% occur at the cathode. Since half of the decays on the cathode will have the α particle emitted invisibly into the cathode material, this means 83% of all radon daughters decay at the cathode. Since ^{214}Bi is the only one that can deposit energy in the fiducial volume, its activity is constrained to $(381.9 \pm 38.2) \mu\text{Bq}$.
- The remaining $(78.2 \pm 46.0) \mu\text{Bq}$ is used as the constraint for ^{222}Rn and its daughters in the active xenon.
- There is also 30.2 kg of inactive xenon. It is assumed that it mixes well with the active xenon. The activity in this volume is constrained to $(110.3 \pm 11.0) \text{ Bq}$.

These constraints are imposed as Gaussian constraints (eq. (8.8)) on the N_i (as in eq. (8.6)) for the PDFs, taking into account the efficiencies of all cuts, including the α -particle-eliminating scintillation-to-ionization ratio cut.

8.5.4 Radon in the Air Gap

Radon can seep through the cracks in the lead wall and decay in the small gap between the lead and the cryostat. γ rays from decays of ^{214}Bi in this volume can reach the TPC. An external instrument [70] monitors the level of ^{222}Rn in the clean room air. The average activity during periods when the detector was live was $(7.2 \pm 0.5) \text{ Bq m}^{-3}$. Measurements sampling air from the gap directly have confirmed that radon levels in the gap are similar to those in the clean room air. The volume of air in the gap is known to be $(0.280 \pm 0.015) \text{ m}^3$, and so the total activity of $(2.01 \pm 0.11) \text{ Bq}$ is used, along with the efficiencies from simulation, to constrain the number of decays N_i for ^{214}Bi in the gap. This is applied as a one-sided Gaussian constraint on the upper number of decays.

8.5.5 Normalization

The systematic rate uncertainty described in section 8.3.1 is included in the fit as an overall factor multiplying the normalization. This is not so much a constraint as a way to account for the uncertainty. The constraint has the factor centered around 1, so it does not significantly affect the best fit value. The width is given by the rate uncertainty given in table 8.3, and so it contributes to the uncertainty on half-life measurements or limits.

8.6 Measurement of $2\nu\beta\beta$

The best fit of the background and signal PDFs to the data is shown in fig. 8.13. The fiducial volume contains 81.9 kg of xenon, of which 80.6 % is ^{136}Xe . The total live time is 116.5 days and the efficiency for detecting $2\nu\beta\beta$ after all cuts is 60 %. The best fit using the maximum likelihood method says the data set contains 1.91×10^4

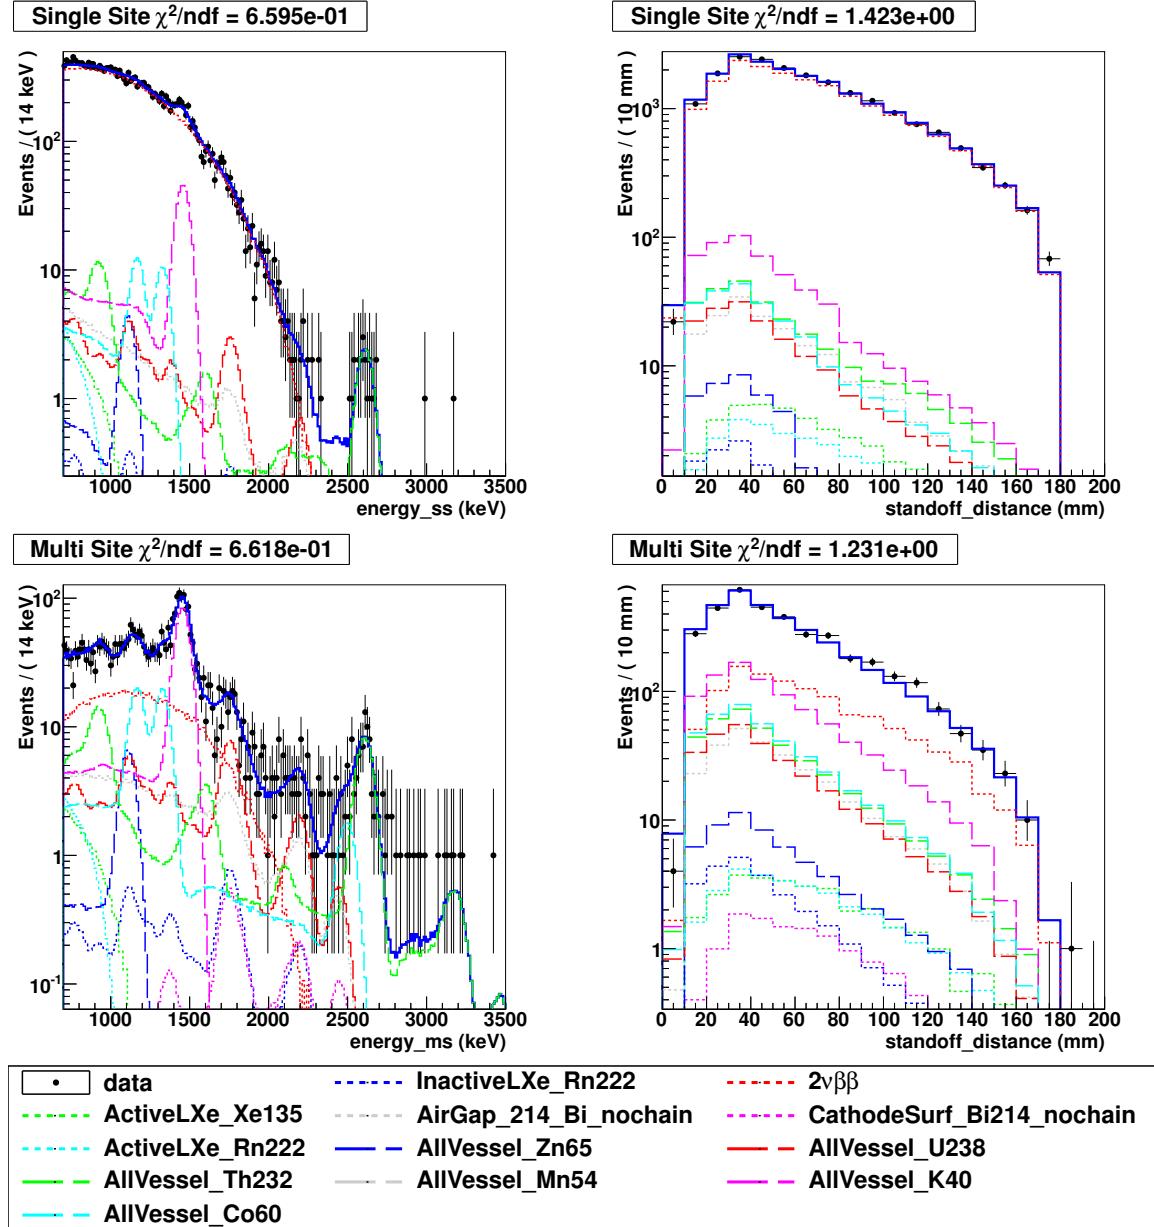


Figure 8.13: The best fit of $2\nu\beta\beta$ and all background components to the data. The legend is shown at bottom. A projection of the energy spectrum is shown on the left, and a projection of the standoff distance spectrum is shown right. Single site events are shown at the top, while multiple site events are shown below.

$2\nu\beta\beta$ decays. This implies a half-life of 2.04×10^{21} yr. Figure 8.13 shows projections of the best fit spectra.

8.6.1 Effects of Systematic Errors

Ultimately, the profile likelihood method gives an overall error, not broken down into statistical and systematic components. However, the relative contributions to the overall error can be obtained by fixing various parameters in the fit and observing the change in the width of the likelihood profile. The statistical uncertainty is obtained by fixing all parameters except the number of $2\nu\beta\beta$ events to their best fit values. Systematic uncertainty can likewise be obtained by looking at just the effect of the single site fraction uncertainty, the contributions of the background PDFs, and the normalization uncertainty. Table 8.5 breaks down these different sources of uncertainty.

Table 8.5: A breakdown of how various uncertainties contribute to the overall uncertainty on the measurement of the $2\nu\beta\beta$ half-life. Each contribution is given as a percentage of the best-fit value obtained by letting only the specified parameter float when performing a profile scan of the likelihood function. Correlations mean they do not sum in quadrature.

Uncertainty	Contribution (%)
Statistical	0.76
SS fraction	0.78
Backgrounds	0.83
Normalization	2.35
Total	3.76

Thus, the measurement of the $2\nu\beta\beta$ half-life for ^{136}Xe is

$$T_{1/2} = (2.04 \pm 0.015(\text{stat}) \pm 0.075(\text{sys})) \times 10^{21}\text{yr} \quad (8.9)$$

8.6.2 Consistency Checks

Time and Position Independence

As fig. 8.14 shows, the half-life obtained by fitting over subsets of the data is consistent with the half-life obtained from the entire data set. Thus, the half-life seems steady over time. Likewise, as fig. 8.15 shows, the $2\nu\beta\beta$ decays seem uniformly distributed throughout the volume.

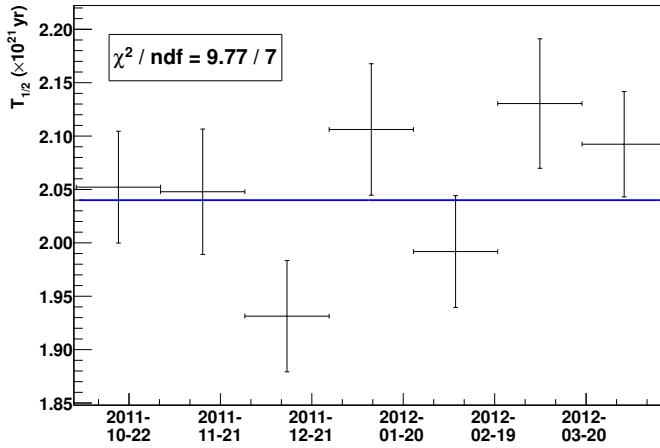


Figure 8.14: The half-life of $2\nu\beta\beta$ throughout the data set. The data set has been broken up into 4 week slices and each individual period has been fit in order to determine the half-life. The error bars do not include the normalization uncertainty, which is a systematic error shared by all fits. The blue line represents the best fit half-life for the entire data set, as reported in eq. (8.9). The χ^2 statistic comparing this value to the individual time periods shows good agreement.

Energy-Only Fit

It is worthwhile to verify that the inclusion of standoff distance in the fit is not somehow biasing the results. Repeating the fit with the standoff distance ignored, the best fit for the $2\nu\beta\beta$ half-life is $T_{1/2} = (2.06 \pm 0.08) \times 10^{21} \text{ yr}$, in agreement with the half-life reported above.

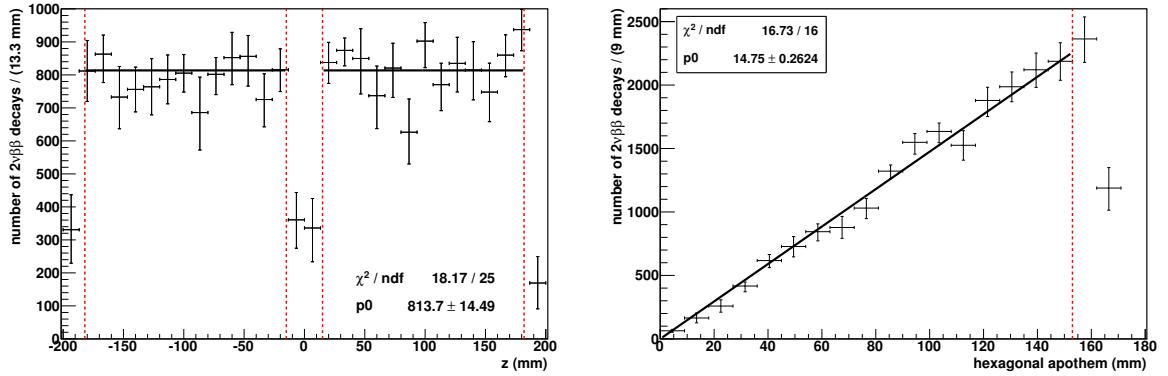


Figure 8.15: The number of $2\nu\beta\beta$ decays returned by the fit when applied to small volumes of the detector. As the plot on the left shows, the rate is consistent with uniformity in z . For the right plot, the volume is proportional to the square of the apothem. Therefore, the volume in a shell bounded by apothem r and $r + dr$ should increase linearly as a function of r . The plot and fit, then, show that the $2\nu\beta\beta$ decays are consistent with being uniformly distributed throughout the fiducial volume.

8.6.3 Comparison to Previous Results

Figure 8.16 shows previous measurements of the $2\nu\beta\beta$ half-life made with the EXO-200 and KamLAND-Zen experiments. The EXO-200 Run 2a data was previously used to report $T_{1/2} = (2.23 \pm 0.017(\text{stat}) \pm 0.22(\text{sys})) \times 10^{21}\text{yr}$. This measurement suffered from a 3% overestimate in the detector live time, and correcting for this brings the measurement closer to the one reported here. That measurement also used a larger fiducial volume, including regions near the detector edges where the volume is less well-known (this is demonstrated by the apparent drop-off in events near the walls and anode/cathode planes in fig. 8.15). So while this measurement is consistent with previous measurements made with the EXO-200 experiment, it is now in tension with the measurements from KamLAND-Zen.

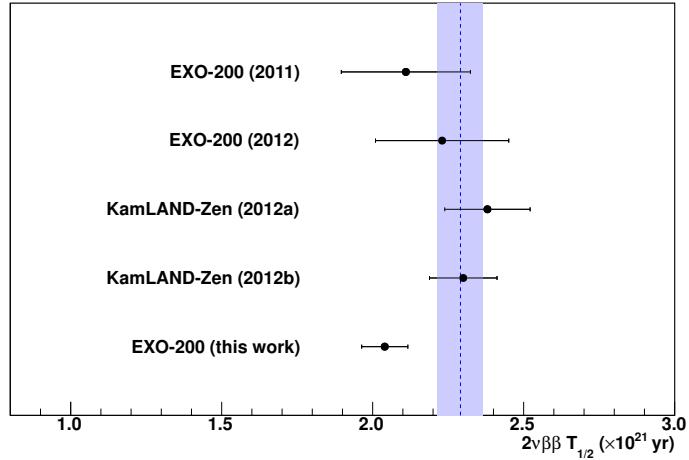


Figure 8.16: A comparison of measurements of the $2\nu\beta\beta$ half-life for ^{136}Xe . The EXO-200 experiment reported results in 2011 [5] and 2012 [6] using independent data sets. The KamLAND-Zen experiment reported two results in 2012, denoted “a” [71] and “b” [30], with some data in common between the two analyses. The blue line and blue shaded area are the maximum likelihood estimate and the 1σ error for the half-life ($T_{1/2} = (2.29 \pm 0.076) \times 10^{21} \text{ yr}$) based on these previous measurements, assuming all are independent. The bottom point is the measurement reported in this work.

8.7 Limits on $0\nu\beta\beta\chi^0(\chi^0)$

The Majoron-emitting modes are each separately included in the fit, and confidence intervals are constructed for each using the bounded likelihood method described above in section 8.4.1. All modes fit to a value consistent with zero events at the 68% confidence level. The likelihood profiles used to extract the 90% confidence limits are shown in fig. 8.17. Limits on the number of events are translated into limits on half-lives using the known live time and number of ^{136}Xe nuclei in the fiducial volume. The efficiency for detection is 90.0% for the mode with spectral index 1; 85.7% for the mode with spectral index 2; 77.4% for the mode with spectral index 3; and 45.0% for the mode with spectral index 7. Figure 8.18 shows projections of the single site energy spectra with the 90% confidence limits for the Majoron-emitting modes. Table 8.6 summarizes the limits obtained in the context of a number of models.

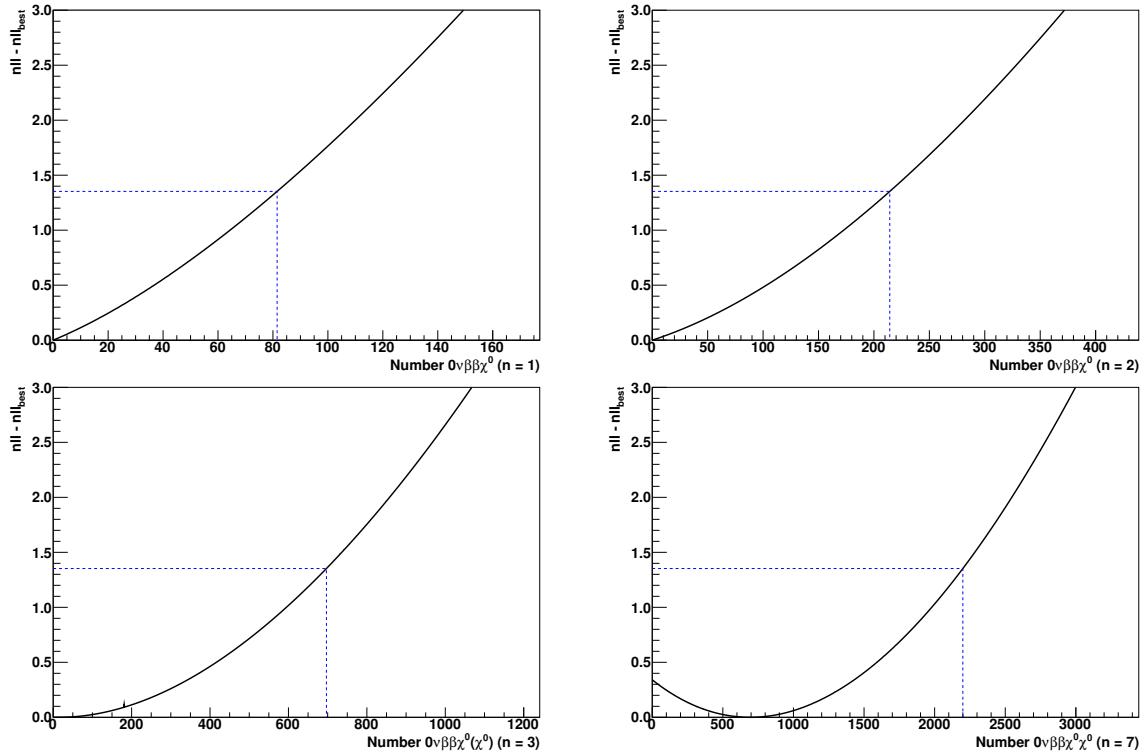


Figure 8.17: Profile likelihood scans for the Majoron-emitting modes. Though some modes fit to a non-zero value best fit, they are not significant at the 90% confidence level. The change in negative log likelihood corresponding to the limit, and the limit, are shown in blue. Upper left is spectral index 1. Upper right is spectral index 2. Bottom left is spectral index 3. Bottom right is spectral index 7.

Table 8.6: The half-life limits obtained for the various Majoron emitting modes in ^{136}Xe by KamLAND-Zen [30] and in this work. All limits are at the 90% confidence level.

Decay mode	n	$T_{1/2}(\text{yr})$	
		KamLAND-Zen	EXO-200 (this work)
$0\nu\beta\beta\chi^0$	1	$> 2.6 \times 10^{24}$	$> 7.15 \times 10^{23}$
$0\nu\beta\beta\chi^0$	2	$> 1.0 \times 10^{24}$	$> 2.59 \times 10^{23}$
$0\nu\beta\beta\chi^0(\chi^0)$	3	$> 4.5 \times 10^{23}$	$> 7.35 \times 10^{22}$
$0\nu\beta\beta\chi^0\chi^0$	7	$> 1.1 \times 10^{22}$	$> 1.33 \times 10^{22}$

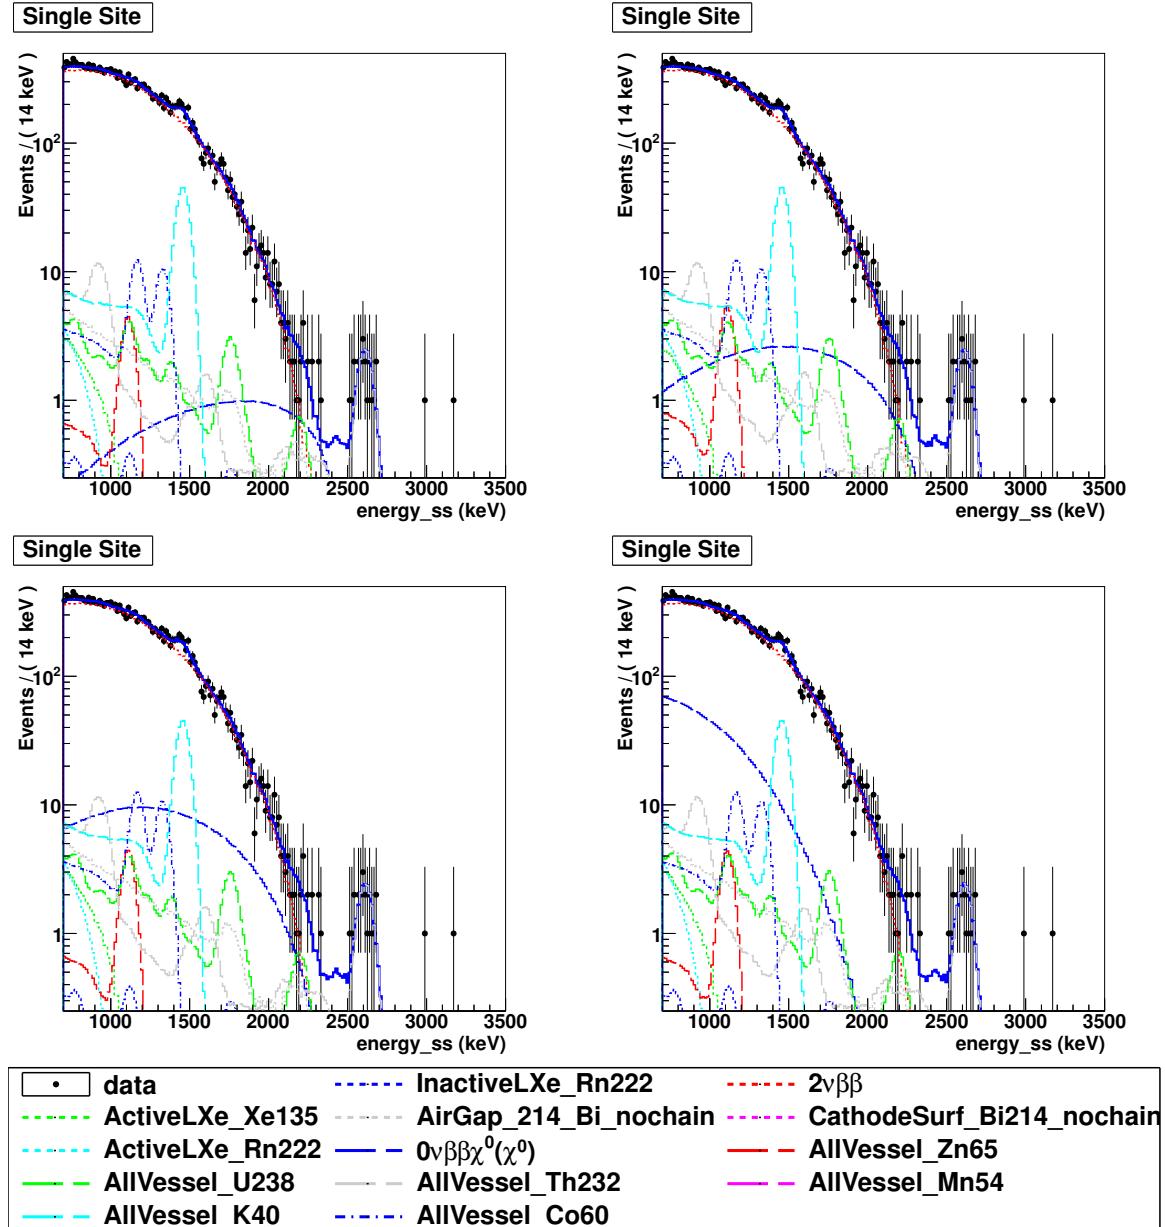


Figure 8.18: The best fits of $2\nu\beta\beta$, $0\nu\beta\beta\chi^0(\chi^0)$, and all background components to the data. Only projections of the single site energy spectra are shown. For each of the Majoron-emitting modes, the 90% confidence level upper limit is shown. Upper left is spectral index 1. Upper right is spectral index 2. Bottom left is spectral index 3. Bottom right is spectral index 7.

8.7.1 Effects of Systematic Errors

Table 8.7 tabulates the effect of the various uncertainties on the limits obtained for the Majoron-emitting modes. All of the limits are only slightly degraded below their statistical limits. For all modes, uncertainty in the backgrounds has a modest effect on the limit. Being able to better constrain the backgrounds, both with more data and through external constraints, could improve the limits. For all modes, the statistical limits are weakened more by the uncertainty in the single site fraction. Since these events are similar to $2\nu\beta\beta$ events, this suggests the limits could be improved modestly by using the $2\nu\beta\beta$ single site fraction as a constraint on the single site fraction for these modes. The normalization uncertainty does not weaken the limit significantly, and so enlarging the fiducial volume to increase the number of xenon atoms would likely improve the limit, despite worsening the normalization uncertainty.

Table 8.7: A breakdown of how various uncertainties affect the limits obtained on the Majoron-emitting modes. These are obtained by fixing parameters to their best-fit values and then letting only the specified parameters float when doing the profile scan. $2\nu\beta\beta$ is considered a background for these results.

Uncertainty	90% CL limit on $T_{1/2}$			
	$n = 1 (10^{23} \text{ yr})$	$n = 2 (10^{23} \text{ yr})$	$n = 3 (10^{22} \text{ yr})$	$n = 7 (10^{22} \text{ yr})$
Statistical only	7.42	2.71	7.99	1.38
SS fraction	7.22	2.62	7.50	1.37
Backgrounds	7.39	2.69	7.92	1.36
Normalization	7.42	2.71	7.99	1.38
Reported	7.15	2.59	7.35	1.33

8.7.2 Coverage Tests and Sensitivity

As discussed in section 8.4.1, when using the bounded likelihood method of Rolke et al., it is important to verify that the intervals constructed have coverage. To study this, the best fit spectrum with no $0\nu\beta\beta\chi^0(\chi^0)$ decays included is used to randomly generate toy Monte Carlo data sets. Sampling from the $0\nu\beta\beta\chi^0(\chi^0)$ PDF, a known number of Majoron-emitting decay events are injected into the toy data set, and the

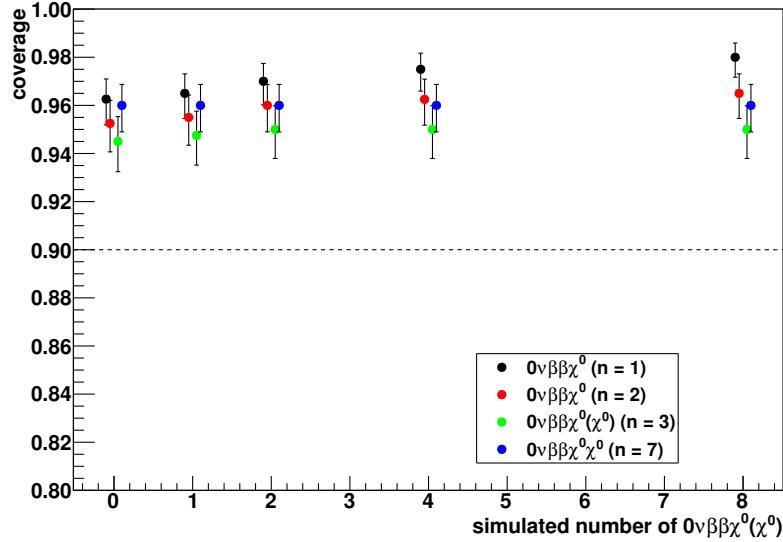


Figure 8.19: The results of coverage tests for the confidence intervals produced by the fits for Majoron-emitting modes when setting a limit. For a nominal 90% confidence level, good coverage is shown for all modes. Each point represents the coverage fraction for 400 tests with the indicated number (0, 1, 2, 4, or 8) of $0\nu\beta\beta\chi^0(\chi^0)$ events injected. In all cases, an integral number of events were injected; the points are spaced out for ease of visualization.

profile likelihood method is used to construct confidence intervals. Doing this many times provides an estimate of how likely the intervals are to cover the true value. As fig. 8.19 shows, the method provides adequate coverage.

Using the best fit spectrum with no $0\nu\beta\beta\chi^0(\chi^0)$ to again create toy simulations, it is possible to study the sensitivity to Majoron-emitting modes. Each simulation is used to construct a 90% confidence interval. Figure 8.20 shows the distributions of the limits obtained from many simulations. For all modes, the expected sensitivity is higher than the observed limit. For the modes with spectral indices 1 and 2, higher limits are expected 73% and 78% of the time, respectively. For the mode with spectral index 3, 95% of simulations had a better limit. Looking at the distributions for these modes, the limits obtained are close to the most likely limit. For the mode with spectral index 7, only two simulations in 500 had a comparable limit. Higher limits would be expected nearly all of the time. This suggests that the model used for the toy simulations is inadequate. As discussed in section 8.2.3 the simulations

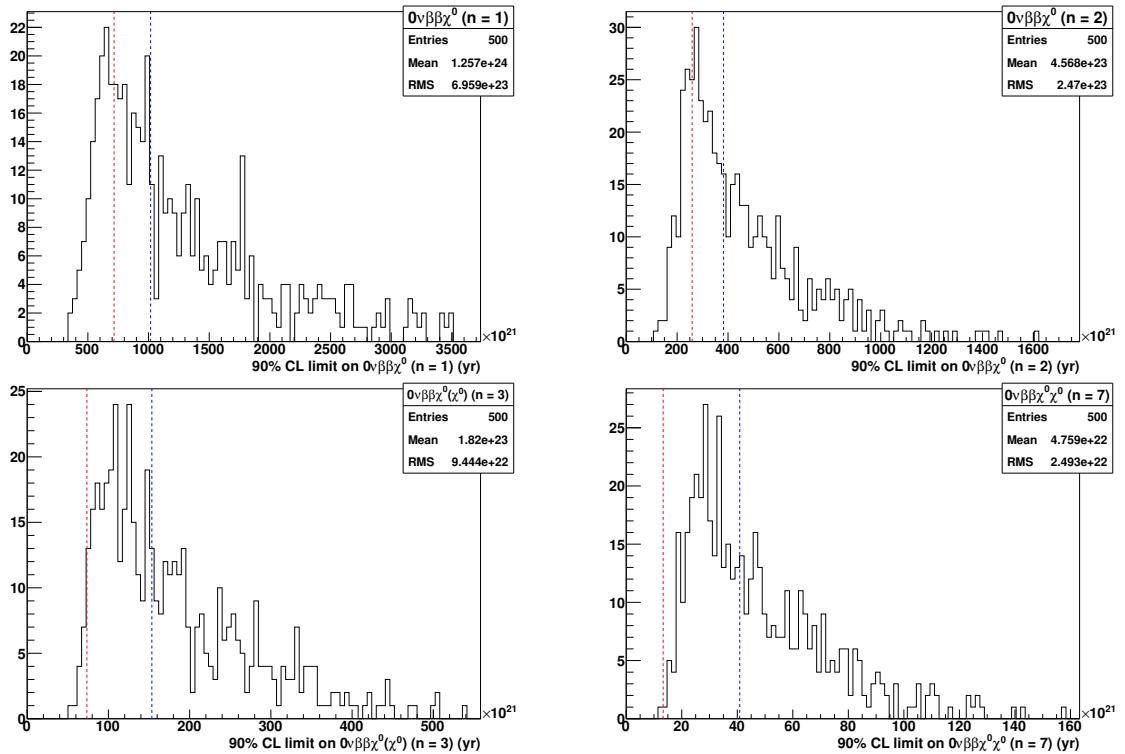


Figure 8.20: The limits obtained for toy simulations containing zero Majoron-emitting events, made from the best-fit to the data. The median limit is shown in blue, while the obtained limit is shown in red.

do not have quite as many single site events at low energies as in data. Since the $0\nu\beta\beta\chi^0\chi^0$ spectrum is primarily single site and peaks at lower energies than $2\nu\beta\beta$, the upper limit on the number of decays is larger when considering data than in the toy simulations.

8.7.3 Comparison to Previous Results

As table 8.6 shows, the limits on the Majoron-emitting modes with spectral indices 1, 2, and 3 are weaker than those reported by the KamLAND-Zen collaboration. The weaker limits are primarily due to statistics, as discussed in section 8.7.1. This is not surprising, as the KamLAND-Zen results come from 38.4 kg yr exposure of ^{136}Xe , whereas this analysis only contains 21.0 kg yr. For the spectral index 7 mode, data from EXO-200 produces a modestly improved limit. Looking at fig. 8.21, backgrounds in KamLAND-Zen begin to dominate over the $2\nu\beta\beta$ mode below 1 MeV. EXO-200, on the other hand, has much smaller backgrounds at this low energy, which allow a stronger limit, even with reduced exposure.

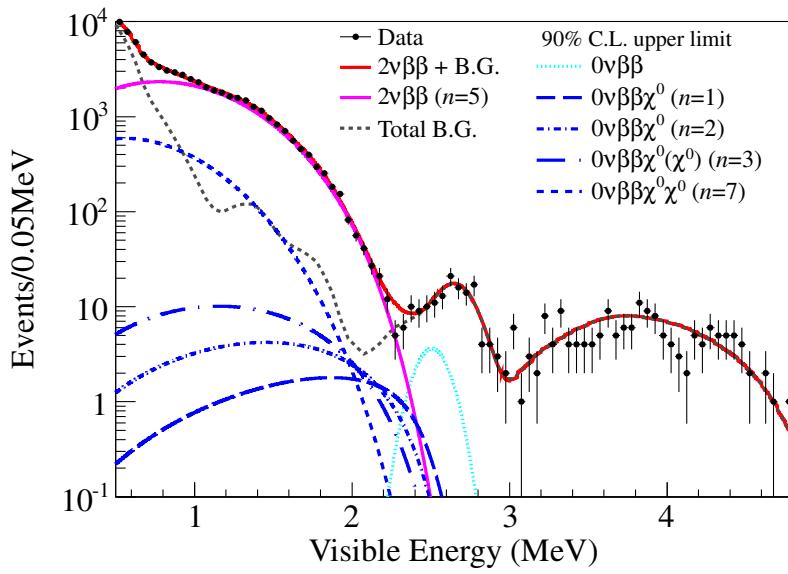


Figure 8.21: The energy spectrum reported by KamLAND-Zen. Backgrounds begin to become comparable to the $2\nu\beta\beta$ signal below 1 MeV, reducing sensitivity to the Majoron-emitting mode with spectral index 7. (Figure reproduced from [30].)

Chapter 9

Conclusions

Much work has gone into better understanding the EXO-200 detector. An improved measurement of the vertical muon flux at the WIPP underground site ($\Phi_v = (4.01 \pm 0.04(\text{stat})^{+0.04}_{-0.05}(\text{sys})) \times 10^{-7} \text{ Hz/cm}^2\text{sr}$) made with EXO-200 data allows better estimates of cosmogenic backgrounds. More importantly, the analysis techniques have been improved and the uncertainties on rate measurements and fiducial volume have been reduced compared to previous analyses. The measurement of the $2\nu\beta\beta$ half-life of ^{136}Xe : $T_{1/2} = (2.04 \pm 0.015(\text{stat}) \pm 0.075(\text{sys})) \times 10^{21}\text{yr}$ is the most precise measurement for this isotope to date. Low-background construction and good background rejection allows EXO-200 to place a competitive limit on the Majoron-emitting decay mode $0\nu\beta\beta\chi^0\chi^0$ with spectral index 7: $T_{1/2} > 1.33 \times 10^{22}\text{yr}$ at 90% CL, even with roughly half the exposure of the KamLAND-Zen experiment.

Since the end of the Run 2a data set, EXO-200 has been running steadily for over a year and collecting more data. The analysis techniques and improved detector understanding demonstrated in this analysis will be used in conjunction with this increased exposure to yield improved sensitivity to $0\nu\beta\beta$ and the Majoron-emitting $0\nu\beta\beta\chi^0(\chi^0)$ modes in future EXO-200 analyses.

Appendix A

Tagging Electronic Noise

A.1 Motivation

Ideally, EXO-200 would only trigger on events that are due to particles interacting in the liquid xenon. Unlike the detectors at hadron colliders that have a large event rate and must selectively trigger, EXO-200 only sees about one event every 10 seconds. Therefore, the triggers do not need to be sophisticated, and indeed err on the side of including events that may not be interesting.

These simple triggers, however, allow events that are not actually due to interactions in the detector to make it into the data. One common source of spurious events is electronic noise. One source of this noise is microphonic vibration of the signal readout cables due to loud acoustic noises. When these noise signals go through processing, they can slow down the processing due to the odd shapes and large multiplicity of the signals in the events. Furthermore, noise signals can masquerade as real signals. As described below, some of the most common types of noise can be tagged as a TPC muon. These noise events are frequent enough that the dead time enforced after a TPC muon would cause a significant hit to live time. Therefore algorithms were developed to identify and tag noise events.

A.2 Types of Noise Events

A.2.1 Unphysically Negative Signals on the Collection Wires

The collection wires normally collect drifting ionization, which produces a positive-going signal. Some collection wire channels may show a small negative-going induction signal when a signal is collected on a neighboring channel. However, there should not be events in which a large number of collection wire channels have a negative signal without being preceded by or simultaneous to a positive signal.

These events are identified by summing all collection wire waveforms. The baseline of this sum is taken by averaging the first 256 samples of the waveform. If the sum waveform drops 270 ADC counts below baseline before any individual wire goes above 30 ADC counts, the event is tagged as noise.

Figure A.1 shows an example of a typical event caught by this check. These events would otherwise be problematic for the muon tagging algorithm (chapter 7) because they exhibit lines of positive collection signals and negative inductions signals that resemble a muon passing parallel to the wire plane.

A.2.2 “Glitch” Events

For a period of time, the high-voltage supply for the EXO-200 cathode would occasionally cause many channels in the TPC to completely saturate. Events with 100 or more saturated channels are tagged as “glitch” events. The exact origin of these events is still unknown. They seem to be electronic, since they were still observed saturating all the APD channels even when their gain was reduced to unity.

Figure A.2 shows an example of a “glitch” event. The only other type of event that saturates a large number of channels are high-energy TPC muons, or muons stopping and decaying in the TPC. However, these are rare, and so this noise check only slightly reduces the efficiency of muon tagging.

A.2.3 APD “Bouncing” Events

In some events, the APD signals seem to spontaneously saturate, go far below their baseline, rebound, and then bounce below or high above the baseline again. This seems to be accompanied by noise in the muon veto system, though it has proven difficult to diagnose.

These events are identified by looking for any APD channels that saturate, then fall below 20 ADC counts absolute (without considering baseline), recovering to above 1024 ADC counts (again absolute, which is 1/4 full scale), and then either dipping back below 1024 ADC counts or bouncing above 2048 ADC counts (typical baselines are \sim 1600 ADC counts).

Figure A.3 shows an example of this type of noise. This noise events could otherwise be tagged as muons, since the signal often bleeds over into the wire signals.

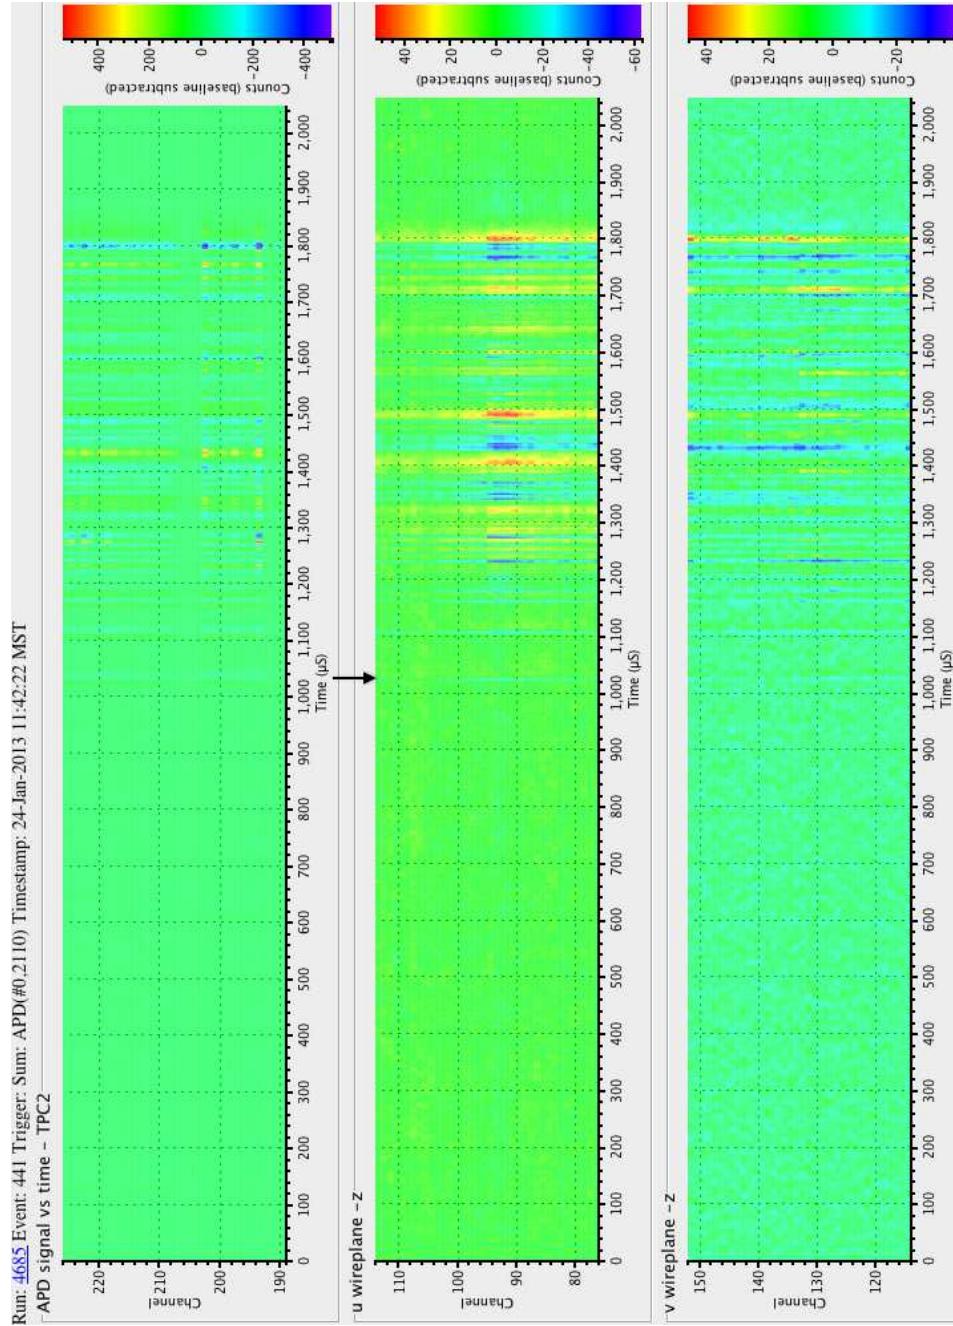


Figure A.1: A noise event which shows the collection (u) wires going unphysically negative (blue color) before any individual wire shows a positive signal (red color). The arrow indicates the first time at which the sum signal goes negative, but more extreme examples are visible later in the event.

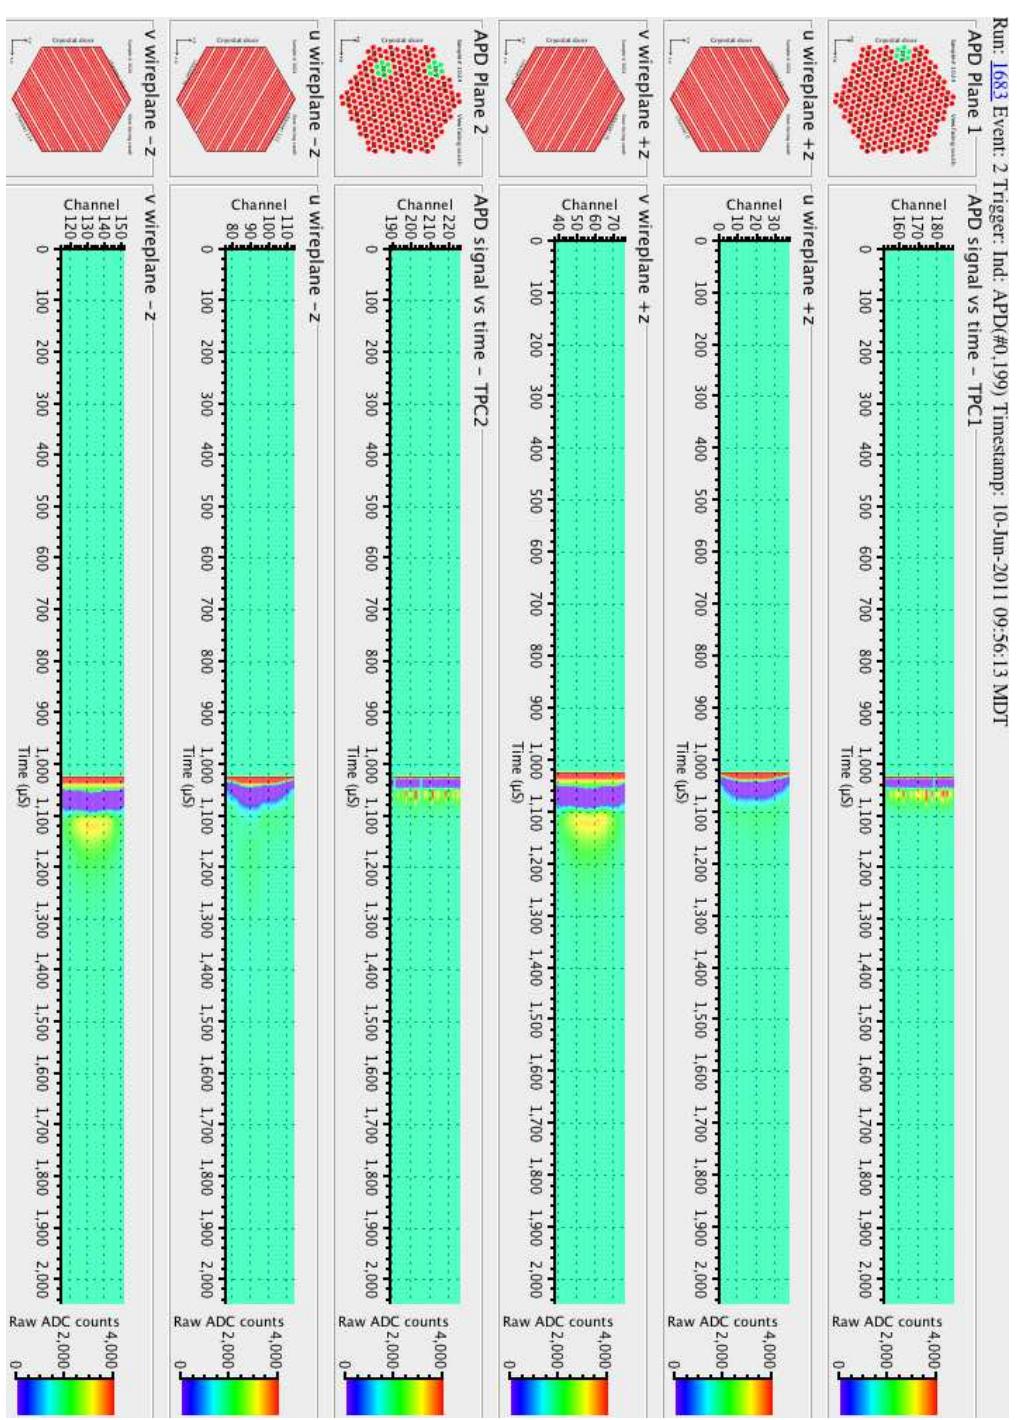


Figure A.2: An example of a “glitch” event showing a large number of saturated (indicated by red) channels.

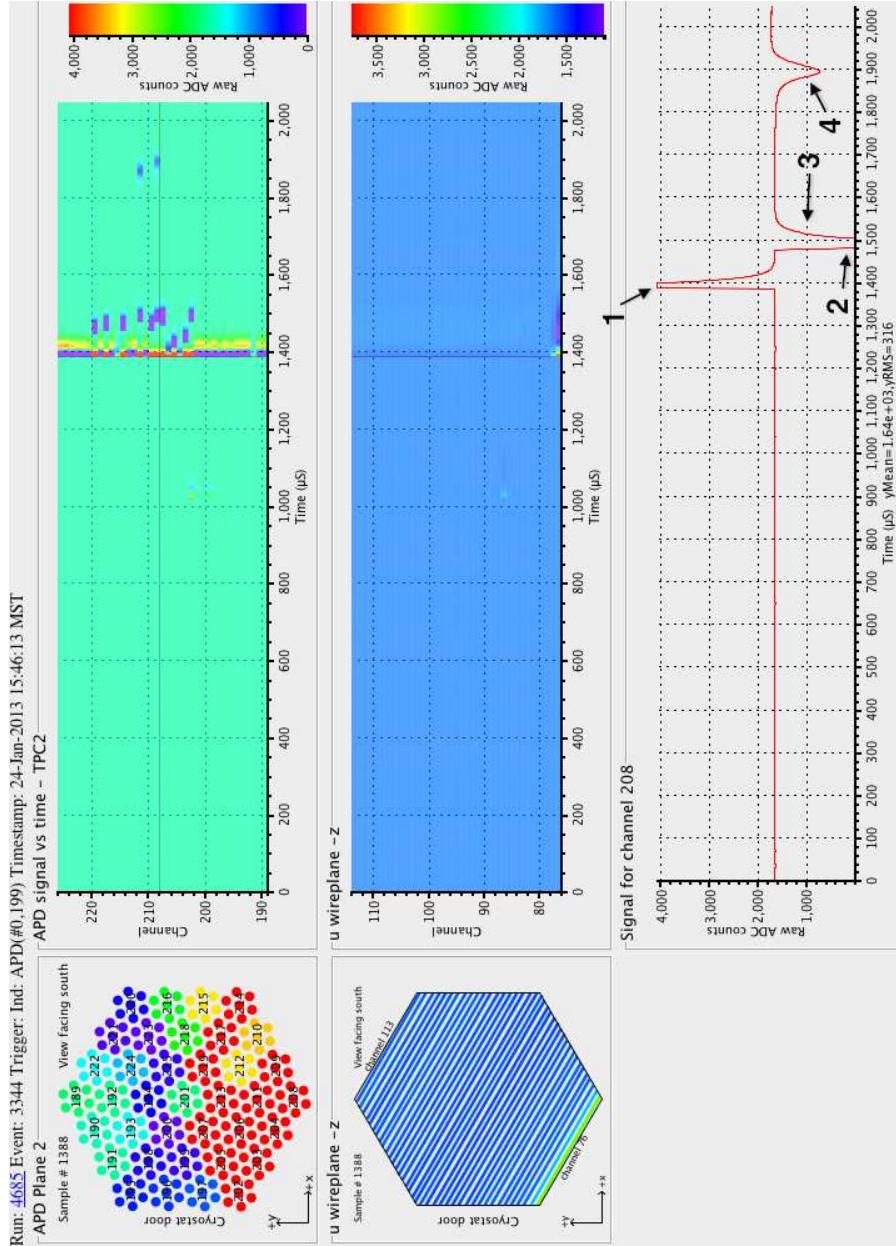


Figure A.3: A noise event with an APD channel signal ‘bouncing’. The top image shows signals for the APD plane and u wires in TPC 2, while the bottom image shows the waveform for one APD channel. No baseline subtraction has been applied in order to show the characteristic noise signal: 1) The APD channel saturates; 2) It dips below 20 ADC counts absolute; 3) It rises back above one quarter full scale; 4) Then dips back below one quarter full scale again.

Appendix B

Correcting for Light Collection Efficiency

B.1 Motivation

Two events with the same energy occurring in different locations in the liquid xenon will not necessarily produce the same scintillation signal. This will lead to a degraded energy resolution. For example, an event occurring near a corner of the detector will have more light escape through the gap between the PTFE reflectors and the APD platter. Additionally, while the trim voltages applied to each APD gang attempt to normalize the performance, there is still a 2.5% spread in the response. While the latter can potentially be accounted for with measurements using the laser pulser system, the system was not operating for most of Run 2a. Such a correction would still not correct for geometric efficiency, in any case.

Although attempts have been made to simulate light collection in the detector (see, for example [72]), this takes a very thorough understanding of the detector construction and the response of all materials to 178 nm light. Individually tracking photons is also extremely computationally intensive. Therefore, data from source calibrations is used to map out the light response of the detector.

B.2 Building the Light Map

B.2.1 Calibration Runs

In order to illuminate the entire detector, long calibration runs at both anode positions and all three cathode locations are needed. These runs are taken as part of larger calibration campaigns, or whenever electronic components are changed or adjusted in a way that might affect the APD response. The thorium source is used because its 2615 keV gamma rays penetrate deep into the detector and produce a large scintillation signal.

B.2.2 Event Selection

Only single site full-absorption events are used to make the light map, since both their location and energies are known. Full absorption events are identified by looking at the ionization spectrum. The mean and sigma of the full-absorption peak are fit with a Gaussian + complementary error function model as described in section 5.4.1. Only events with ionization energies between $+0.33\sigma$ and $+3.0\sigma$ are used. Figure B.1 provides an illustration. This cut retains 37% of the full-absorption events, and should only retain 3% of the Compton scatter events with energies within 2σ of the peak. Since scintillation is correlated with ionization, and the ionization response is uniform throughout the detector (after electron lifetime and shielding grid corrections), this should give a sample of events with the same scintillation throughout the detector.

B.2.3 Binning the Detector Volume

The full detector volume is divided into 1352 spatial bins (see figs. B.3 and B.4 for an illustration of this binning). This binning consists of 8 ϕ bins, 13 radial bins, and 13 z bins. The ϕ binning is evenly-spaced. For the z binning, the full detector is divided evenly into 11 slices. The central slice is then divided into 3 even slices, since the response is observed to change rapidly near the cathode. The radial binning consists of 1 bin from 0 mm to 3 mm, bins every 20 mm from 30 mm to 90 mm, bins every 10 mm from 90 mm to 120 mm, and bins every 8 mm from 120 mm to 168 mm. This

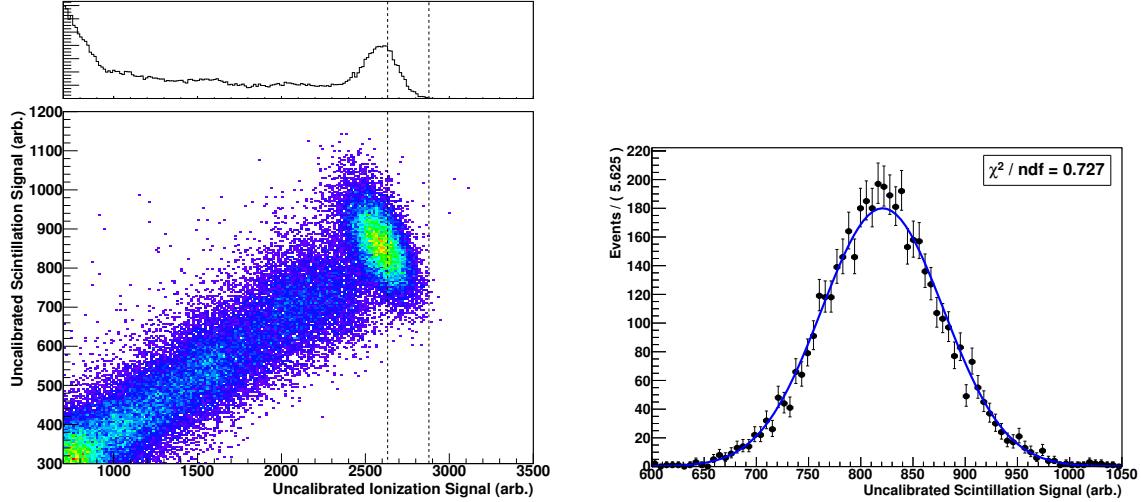


Figure B.1: Selecting events for the light map from a ^{228}Th calibration run. (Left) Only events between $+0.33\sigma$ and $+3.0\sigma$ (indicated by the dashed lines) of the full-absorption peak in ionization are used to form the light map. (Right) The selected events have a scintillation spectrum well-described by a Gaussian distribution, indicating that the cut is successful in sampling mostly full-absorption events.

binning is chosen to ensure adequate statistics within all bins, and to optimally map the response in regions with a high light collection gradient.

Full absorption, single site events from the calibrations are placed into the bins according to the event location. Then the mean of each bin is stored in a 3-dimensional histogram. This histogram stores the scintillation response in each bin of the detector and is referred to as the “light map”. Figures B.3 and B.4 show an example light map. Histograms of the number of events in each bin, and the error on the mean, are also retained to ensure an adequate amount of data was used.

B.3 Correction Function

B.3.1 Interpolating in Cylindrical Polar Coordinates

Suppose a function $F(r_i, \phi_j, z_k)$ is defined only at discrete points, and $f(r, \phi, z)$ is a smooth transition between the defined points of F . The simplest way to do so is with

trilinear interpolation. Define

$$r_d = \frac{r_{i+1} - r_i}{r_{i+1} + r_i}; \phi_d = \frac{\phi_{j+1} - \phi_j}{\phi_{j+1} + \phi_j}; z_d = \frac{z_{k+1} - z_k}{z_{k+1} + z_k} \quad (\text{B.1})$$

where $r_i \leq r < r_{i+1}$ and $z_k \leq z < z_{k+1}$. In cylindrical coordinates, ϕ is cyclical, and so ϕ lies between ϕ_j and $\phi_{j+1} \bmod 2\pi$. Multiples of 2π may need to be added or subtracted from the angles at which F is defined to ensure $0 \leq \phi_d < 1$.

Now f can be constructed:

$$\begin{aligned} f(r, \phi, z) = & F(r_i, \phi_j, z_k)(1 - r_d)(1 - \phi_d)(1 - z_d) \\ & + F(r_i, \phi_j, z_{k+1})(1 - r_d)(1 - \phi_d)z_d \\ & + F(r_i, \phi_{j+1}, z_k)(1 - r_d)\phi_d(1 - z_d) \\ & + F(r_i, \phi_{j+1}, z_{k+1})(1 - r_d)\phi_dz_d \\ & + F(r_{i+1}, \phi_j, z_k)r_d(1 - \phi_d)(1 - z_d) \\ & + F(r_{i+1}, \phi_j, z_{k+1})r_d(1 - \phi_d)z_d \\ & + F(r_{i+1}, \phi_{j+1}, z_k)r_d\phi_d(1 - z_d) \\ & + F(r_{i+1}, \phi_{j+1}, z_{k+1})r_d\phi_dz_d \end{aligned} \quad (\text{B.2})$$

There is a complication, however, when the lower neighbor bin for r is the z -axis. In order to ensure that f is single-valued at $r = 0$, define the value of F at the axis as the mean value for the surrounding points:

$$F(r_0 = 0, \phi, z_k) = \frac{1}{N_\phi} \sum_{j=1}^{N_\phi} F(r_1, \phi_j, z_k) \quad (\text{B.3})$$

Then eq. (B.2) can still be used. In this case, ϕ_d will not matter in the terms involving r_i .

B.3.2 Correcting for Light Response

To correct for the light response, first the histogram containing the light map is normalized such that the average response is 1. Let $F(r_i, \phi_j, z_k)$ be the value of the normalized light map for the bin that has center (r_i, ϕ_j, z_k) . Then let $f(r, \phi, z)$ as in eq. (B.2) be the trilinear interpolation of F . Figures B.5 and B.6 show an example of this interpolation. The normalized scintillation response for an event is

$$g = \frac{\sum_{m=1}^N E_m f(r_m, \phi_m, z_m)}{\sum_{m=1}^N E_m} \quad (\text{B.4})$$

where the sums are over the N ionization clusters in the event, each with a position (r_m, ϕ_m, z_m) and an ionization energy E_m . This assumes that the fraction of total scintillation light produced by a charge cluster at a given position is the same as that cluster's fraction of total ionization energy.

If E_s is the uncorrected scintillation signal, then

$$E_{s,\text{corrected}} = \frac{E_s}{g} \quad (\text{B.5})$$

is the corrected signal.

Table B.1 shows the improvement in the scintillation-only energy resolution for calibration data taken with a ^{228}Th source at several positions around the detector. The mean improvement in σ/E was 2.0 % for single site events, and 1.8 % for multiple site data. An example of the improvement is shown in fig. B.2.

Table B.1: The improvement in the scintillation-only energy resolution (σ/E) after applying the correction for both single site (SS) and multiple site (MS) events. This is based on fitting a Gaussian + complimentary error function model (see 5.4.1) to ^{228}Th calibration data taken at a number of positions.

Source Position	Multiplicity	σ/E uncorrected (%)	σ/E corrected (%)
cathode	+x	SS	8.01 ± 0.18
	+x	MS	8.03 ± 0.08
	+y	SS	8.73 ± 0.28
	+y	MS	8.99 ± 0.12
	-y	SS	8.17 ± 0.24
	-y	MS	8.48 ± 0.15
	+z	SS	7.65 ± 0.26
	+z	MS	7.86 ± 0.09
	-z	SS	7.21 ± 0.21
	-z	MS	7.56 ± 0.10

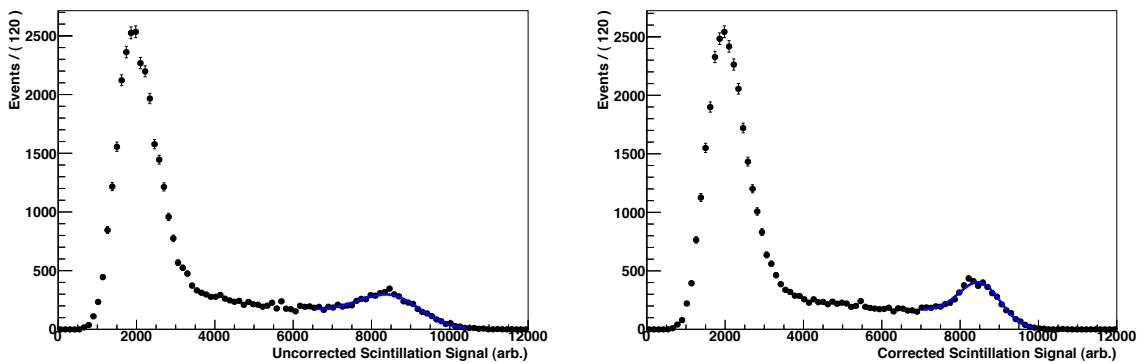


Figure B.2: The scintillation energy spectrum for a ^{228}Th source at the cathode in the $+y$ position. The uncorrected spectrum (left) has $\sigma/E = 8.7\%$. The corrected spectrum (right) has $\sigma/E = 6.1\%$.

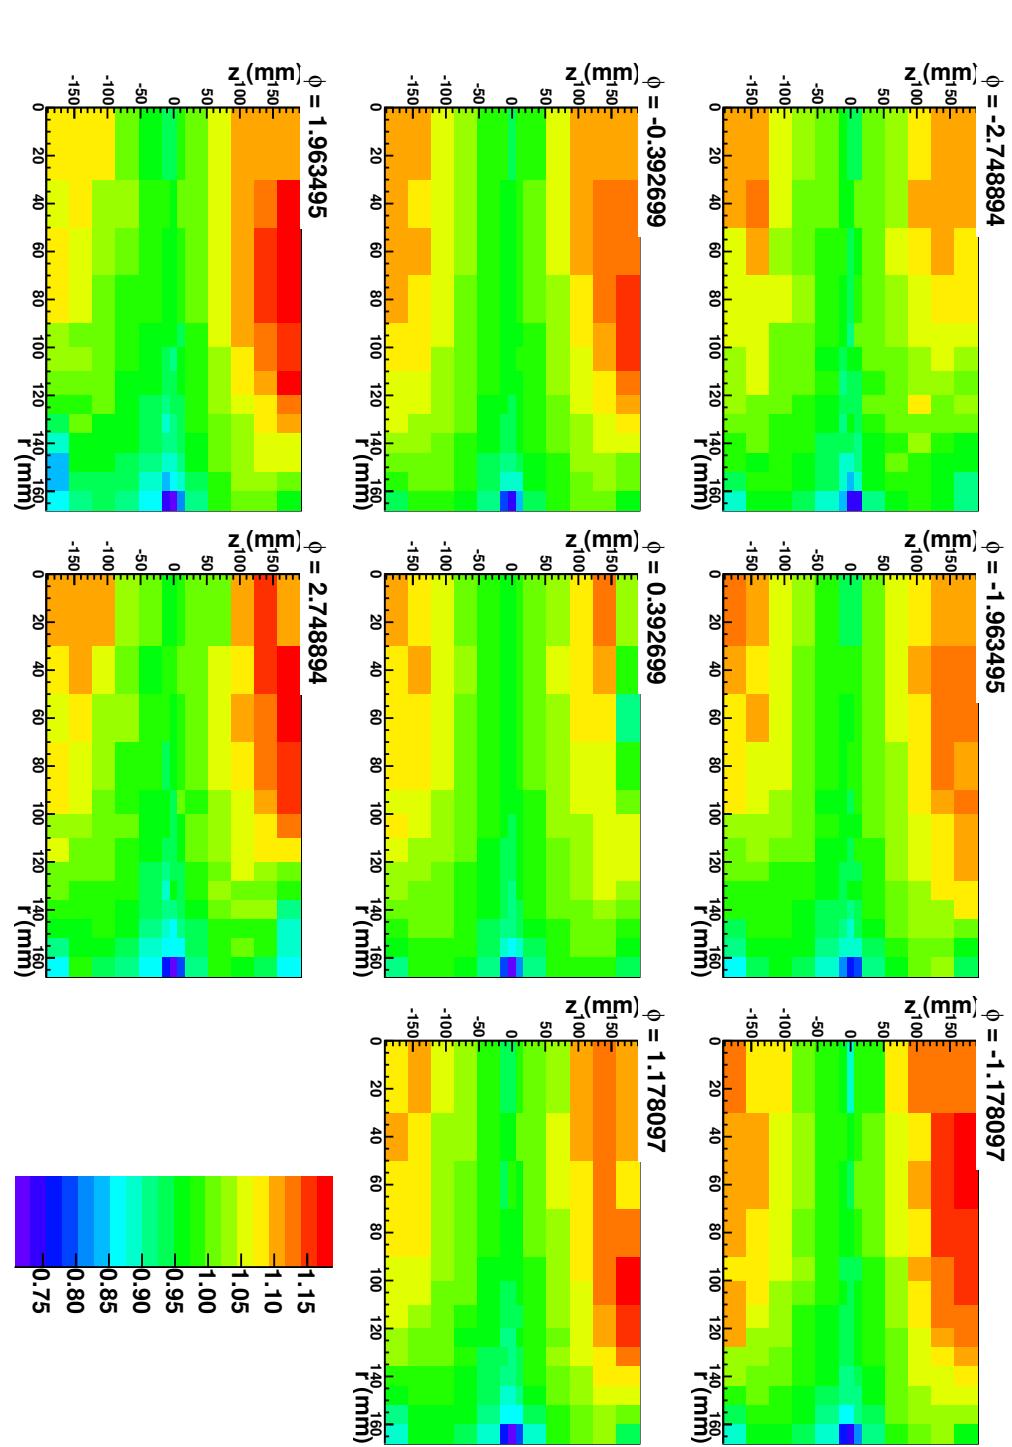


Figure B.3: The light map, showing the bins in r (horizontal axis) and z (vertical axis) for each of the 8 ϕ bins. Note the finer binning around the cathode at $z = 0$.

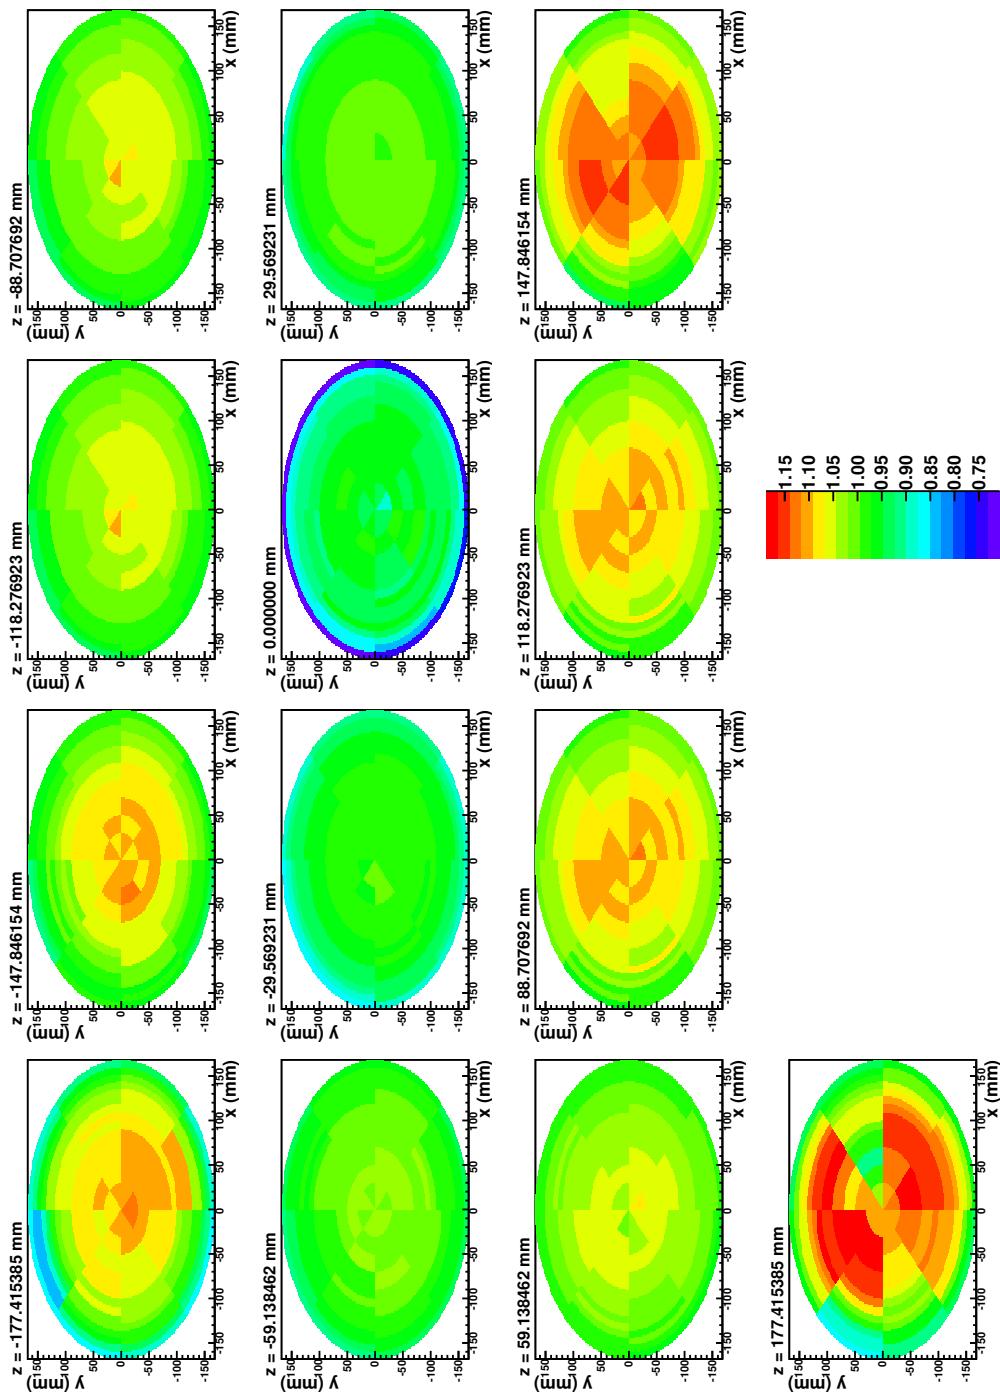


Figure B.4: The light map, showing the bins in r and ϕ for each of the 13 z bins.

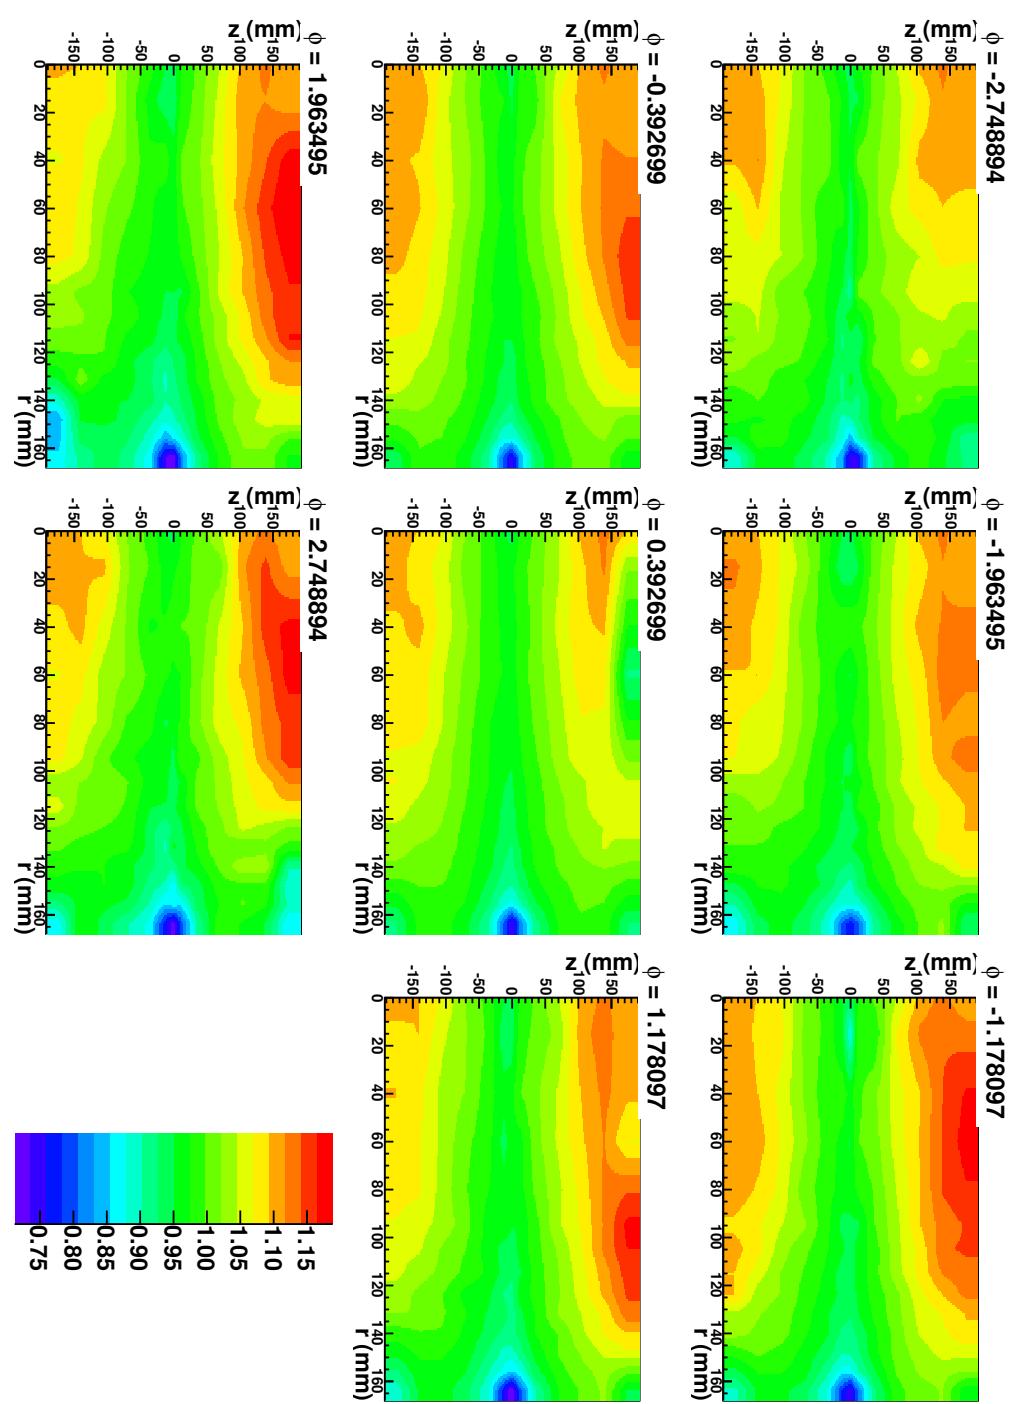


Figure B.5: The light response function $f(r, \phi, z)$, evaluated at 8 discrete values of ϕ over the full range of r (horizontal axis) and z (vertical axis).

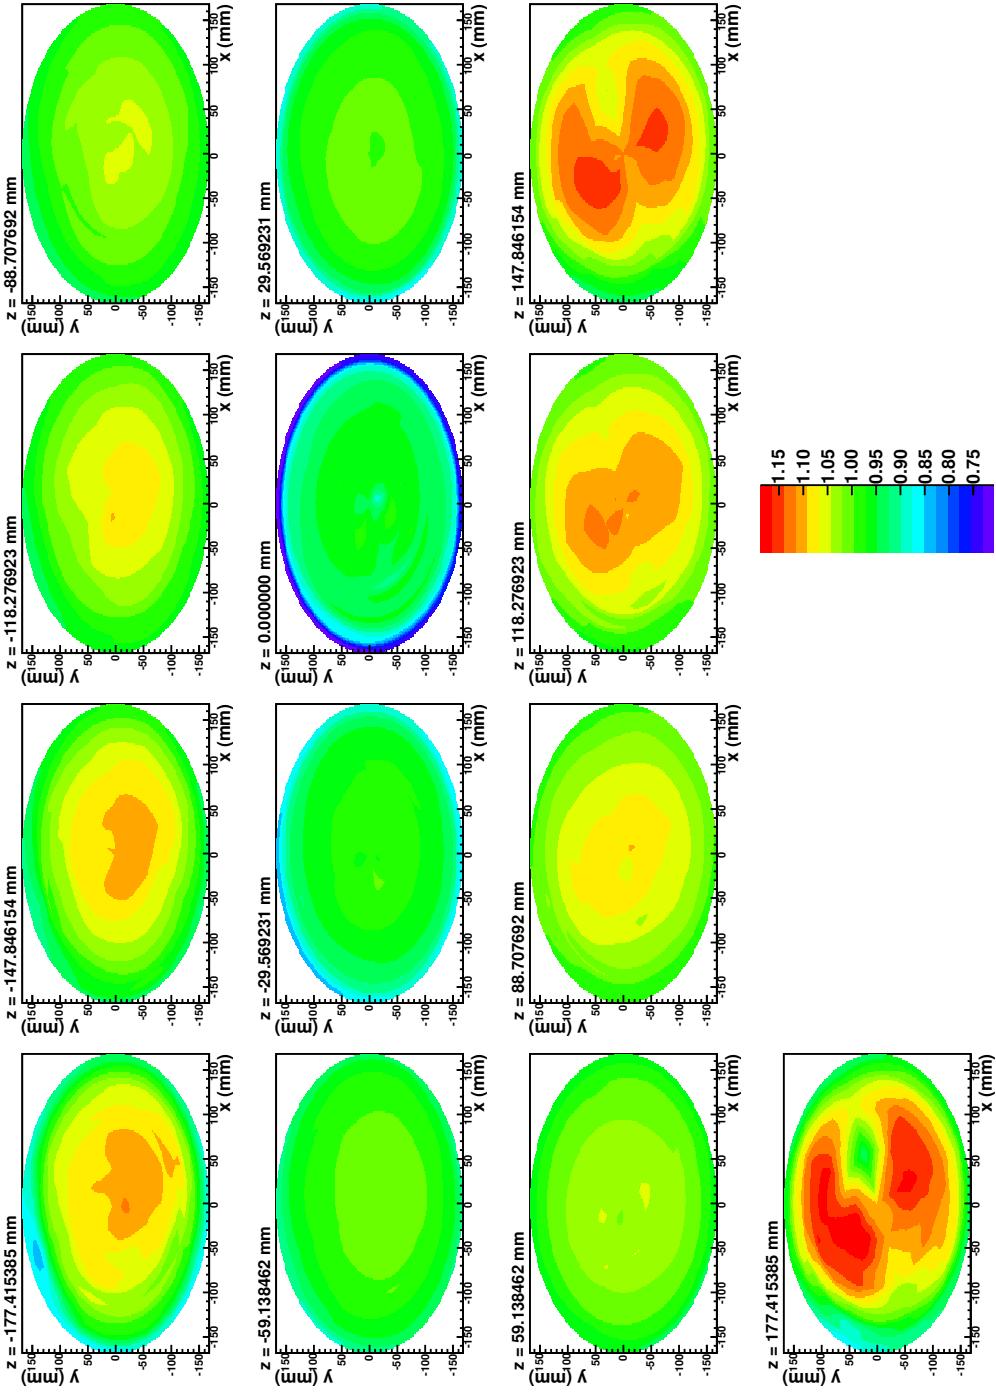


Figure B.6: The light response function $f(r, \phi, z)$, evaluated at 13 discrete values of z over the full range of x (horizontal axis) and y (vertical axis). Notice that the most light is collected for events near the anodes. Near the $+z$ anode, just above the $+x$ axis there is a region of low collection due to a disabled APD gang.

Bibliography

- [1] Y. Fukuda et al. “Evidence for Oscillation of Atmospheric Neutrinos”. In: *Phys. Rev. Lett.* 81.8 (Aug. 1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. arXiv:hep-ex/9807003. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.81.1562>.
- [2] Q. R. Ahmad et al. “Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by ^8B Solar Neutrinos at the Sudbury Neutrino Observatory”. In: *Phys. Rev. Lett.* 87.7 (July 2001), p. 071301. DOI: 10.1103/PhysRevLett.87.071301. arXiv:nucl-ex/0106015. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.87.071301>.
- [3] K. Eguchi et al. “First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance”. In: *Phys. Rev. Lett.* 90.2 (Jan. 2003), p. 021802. DOI: 10.1103/PhysRevLett.90.021802. arXiv:hep-ex/0212021. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.90.021802>.
- [4] F. P. An et al. “Observation of Electron-Antineutrino Disappearance at Daya Bay”. In: *Phys. Rev. Lett.* 108.17 (Apr. 2012), p. 171803. DOI: 10.1103/PhysRevLett.108.171803. arXiv:1203.1669 [hep-ex]. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.108.171803>.
- [5] N. Ackerman et al. “Observation of Two-Neutrino Double-Beta Decay in ^{136}Xe with EXO-200”. In: *Phys. Rev. Lett.* 107 (Nov. 2011), p. 212501. DOI: 10.1103/PhysRevLett.107.212501. arXiv:1108.4193 [nucl-ex]. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.107.212501>.

- [6] M. Auger et al. “Search for Neutrinoless Double-Beta Decay in ^{136}Xe with EXO-200”. In: *Phys. Rev. Lett.* 109 (July 2012), p. 032505. DOI: 10.1103/PhysRevLett.109.032505. arXiv:1205.5608 [hep-ex]. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.109.032505>.
- [7] H. Klapdor-Kleingrothaus and I. Krivosheina. “The Evidence for the Observation of $0\nu\beta\beta$ Decay: The Identification of $0\nu\beta\beta$ Events from the Full Spectra”. In: *Mod. Phys. Lett. A* 21 (2006), pp. 1547–1566. ISSN: 1793-6632. DOI: 10.1142/S0217732306020937. URL: <http://www.worldscientific.com/doi/abs/10.1142/S0217732306020937>.
- [8] E. Fermi. “Versuch einer Theorie der β -Strahlen. I”. German. In: *Zeitschrift für Physik* 88.3-4 (1934), pp. 161–177. ISSN: 0044-3328. DOI: 10.1007/BF01351864. URL: <http://dx.doi.org/10.1007/BF01351864>.
- [9] C. L. Cowan et al. “Detection of the Free Neutrino: a Confirmation”. In: *Science* 124.3212 (1956), pp. 103–104. DOI: 10.1126/science.124.3212.103. URL: <http://www.sciencemag.org/content/124/3212/103.short>.
- [10] G. Danby et al. “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos”. In: *Phys. Rev. Lett.* 9.1 (July 1962), pp. 36–44. DOI: 10.1103/PhysRevLett.9.36. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.9.36>.
- [11] B. T. Cleveland et al. “Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector”. In: *Astrophys. J.* 496.1 (1998), p. 505. URL: <http://stacks.iop.org/0004-637X/496/i=1/a=505>.
- [12] S. Weinberg. “Baryon- and Lepton-Nonconserving Processes”. In: *Phys. Rev. Lett.* 43.21 (Nov. 1979), pp. 1566–1570. DOI: 10.1103/PhysRevLett.43.1566. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.43.1566>.
- [13] G. Gelmini and M. Roncadelli. “Left-Handed Neutrino Mass Scale and Spontaneously Broken Lepton Number”. In: *Phys. Lett. B* 99.5 (1981), pp. 411–415. ISSN: 0370-2693. DOI: 10.1016/0370-2693(81)90559-1. URL: <http://www.sciencedirect.com/science/article/pii/0370269381905591>.

- [14] H. M. Georgi, S. L. Glashow, and S. Nussinov. “Unconventional Model of Neutrino Masses”. In: *Nucl. Phys. B* 193.2 (1981), pp. 297–316. ISSN: 0550-3213. DOI: 10.1016/0550-3213(81)90336-9. URL: <http://www.sciencedirect.com/science/article/pii/0550321381903369>.
- [15] Y. Chikashige, R. Mohapatra, and R. Peccei. “Are There Real Goldstone Bosons Associated with Broken Lepton Number?” In: *Phys. Lett. B* 98.4 (1981), p. 265. ISSN: 0370-2693. DOI: 10.1016/0370-2693(81)90011-3. URL: <http://www.sciencedirect.com/science/article/pii/0370269381900113>.
- [16] Z. Berezhiani, A. Smirnov, and J. Valle. “Observable Majoron Emission in Neutrinoless Double Beta Decay”. In: *Phys. Lett. B* 291.1-2 (1992), pp. 99–105. ISSN: 0370-2693. DOI: 10.1016/0370-2693(92)90126-0. URL: <http://www.sciencedirect.com/science/article/pii/0370269392901260>.
- [17] P. Bamert, C. Burgess, and R. Mohapatra. “Multi-Majoron Modes for Neutrinoless Double-Beta Decay”. In: *Nucl. Phys. B* 449.1-2 (1995), pp. 25–48. ISSN: 0550-3213. DOI: 10.1016/0550-3213(95)00273-U. URL: <http://www.sciencedirect.com/science/article/pii/055032139500273U>.
- [18] C. Kraus et al. “Final Results From Phase II of the Mainz Neutrino Mass Search in Tritium β Decay”. In: *Eur. Phys. J. C* 40.4 (2005), pp. 447–468. ISSN: 1434-6044. DOI: 10.1140/epjc/s2005-02139-7. arXiv:0412056. URL: <http://dx.doi.org/10.1140/epjc/s2005-02139-7>.
- [19] V. N. Aseev et al. “An Upper Limit on Electron Antineutrino Mass from Troitsk Experiment”. In: *Phys. Rev. D* 84 (2011), p. 112003. DOI: 10.1103/PhysRevD.84.112003. arXiv:1108.5034 [hep-ex].
- [20] A. Osipowicz et al. “KATRIN: A Next Generation Tritium Beta Decay Experiment with Sub-eV Sensitivity for the Electron Neutrino Mass. Letter of Intent”. In: *Unpublished* (2001). arXiv:hep-ex/0109033.
- [21] P. Ade et al. “Planck 2013 Results. XVI. Cosmological Parameters”. In: *Preprint* (2013). arXiv:1303.5076 [astro-ph.CO].

- [22] K. Abazajian et al. “Cosmological and Astrophysical Neutrino Mass Measurements”. In: *Astropart. Phys.* 35 (July 2011), pp. 177–184. DOI: 10.1016/j.astropartphys.2011.07.002. arXiv:1103.5083 [astro-ph.CO].
- [23] R. Arnold et al. “First Results of the Search for Neutrinoless Double-Beta Decay with the NEMO 3 Detector”. In: *Phys. Rev. Lett.* 95.18 (Oct. 2005), p. 182302. DOI: 10.1103/PhysRevLett.95.182302. arXiv:nucl-ex/0507083. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.95.182302>.
- [24] S. Umebara et al. “Neutrino-less Double- β Decay of ^{48}Ca Studied by $\text{CaF}_2(\text{Eu})$ Scintillators”. In: *Phys. Rev. C* 78.5 (Nov. 2008), p. 058501. DOI: 10.1103/PhysRevC.78.058501. arXiv:0810.4746 [nucl-ex]. URL: <http://link.aps.org/doi/10.1103/PhysRevC.78.058501>.
- [25] H. Klapdor-Kleingrothaus et al. “Latest Results from the Heidelberg - Moscow Double Beta Decay Experiment”. In: *Eur. Phys. J. A* 12.2 (2001), pp. 147–154. ISSN: 1434-6001. DOI: 10.1007/s100500170022. arXiv:hep-ph/0103062. URL: <http://dx.doi.org/10.1007/s100500170022>.
- [26] A. Barabash and V. Brudanin. “Investigation of Double-Beta Decay with the NEMO-3 Detector”. In: *Phys. Atom. Nucl.* 74.2 (2011), pp. 312–317. ISSN: 1063-7788. DOI: 10.1134/S1063778811020062. arXiv:1002.2862 [nucl-ex]. URL: <http://dx.doi.org/10.1134/S1063778811020062>.
- [27] F. A. Danevich et al. “Search for 2β Decay of Cadmium and Tungsten Isotopes: Final Results of the Solotvina Experiment”. In: *Phys. Rev. C* 68 (3 Sept. 2003), p. 035501. DOI: 10.1103/PhysRevC.68.035501. URL: <http://link.aps.org/doi/10.1103/PhysRevC.68.035501>.
- [28] E. Andreotti et al. “ ^{130}Te Neutrinoless Double-Beta Decay with CUORICINO”. In: *Astropart. Phys.* 34.11 (2011), pp. 822–831. ISSN: 0927-6505. DOI: 10.1016/j.astropartphys.2011.02.002. arXiv:1012.3266 [nucl-ex]. URL: <http://www.sciencedirect.com/science/article/pii/S0927650511000351>.

- [29] A. Gando et al. “Limit on Neutrinoless $\beta\beta$ Decay of ^{136}Xe from the First Phase of KamLAND-Zen and Comparison with the Positive Claim in ^{76}Ge ”. In: *Phys. Rev. Lett.* 110.6 (Feb. 2013), p. 062502. DOI: 10.1103/PhysRevLett.110.062502. arXiv:1211.3863 [hep-ex]. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.110.062502>.
- [30] A. Gando et al. “Limits on Majoron-Emitting Double- β Decays of ^{136}Xe in the KamLAND-Zen Experiment”. In: *Phys. Rev. C* 86.2 (Aug. 2012), p. 021601. DOI: 10.1103/PhysRevC.86.021601. arXiv:1205.6372 [hep-ex]. URL: <http://link.aps.org/doi/10.1103/PhysRevC.86.021601>.
- [31] D. V. Forero, M. Tórtola, and J. W. F. Valle. “Global Status of Neutrino Oscillation Parameters after Neutrino-2012”. In: *Phys. Rev. D* 86.7 (Oct. 2012), p. 073012. DOI: 10.1103/PhysRevD.86.073012. arXiv:1205.4018 [hep-ph]. URL: <http://link.aps.org/doi/10.1103/PhysRevD.86.073012>.
- [32] M. Redshaw et al. “Mass and Double-Beta-Decay Q Value of ^{136}Xe ”. In: *Phys. Rev. Lett.* 98.5 (Feb. 2007), p. 053003. DOI: 10.1103/PhysRevLett.98.053003. URL: <http://link.aps.org/doi/10.1103/PhysRevLett.98.053003>.
- [33] T. Takahashi et al. “Average Energy Expended per Ion Pair in Liquid Xenon”. In: *Phys. Rev. A* 12.5 (Nov. 1975), pp. 1771–1775. DOI: 10.1103/PhysRevA.12.1771. URL: <http://link.aps.org/doi/10.1103/PhysRevA.12.1771>.
- [34] U. Fano. “Ionization Yield of Radiations. II. The Fluctuations of the Number of Ions”. In: *Phys. Rev.* 72.1 (July 1947), pp. 26–29. DOI: 10.1103/PhysRev.72.26. URL: <http://link.aps.org/doi/10.1103/PhysRev.72.26>.
- [35] T. Doke et al. “Estimation of Fano Factors in Liquid Argon, Krypton, Xenon and Xenon-Doped Liquid Argon”. In: *Nucl. Instrum. Meth.* 134.2 (Apr. 1976), pp. 353–357. ISSN: 0029-554X. DOI: 10.1016/0029-554X(76)90292-5. URL: <http://www.sciencedirect.com/science/article/pii/0029554X76902925>.
- [36] E. Aprile and T. Doke. “Liquid Xenon Detectors for Particle Physics and Astrophysics”. In: *Rev. Mod. Phys.* 82.3 (July 2010), pp. 2053–2097. DOI: 10.

1103/RevModPhys . 82 . 2053. URL: <http://link.aps.org/doi/10.1103/RevModPhys.82.2053>.

[37] T. Doke et al. “LET Dependence of Scintillation Yields in Liquid Argon”. In: *Nucl. Instrum. Meth. A* 269.1 (1988), pp. 291–296. ISSN: 0168-9002. DOI: 10.1016/0168-9002(88)90892-3. URL: <http://www.sciencedirect.com/science/article/pii/0168900288908923>.

[38] J. Thomas and D. A. Imel. “Recombination of Electron-Ion Pairs in Liquid Argon and Liquid Xenon”. In: *Phys. Rev. A* 36.2 (July 1987), pp. 614–616. DOI: 10.1103/PhysRevA.36.614. URL: <http://link.aps.org/doi/10.1103/PhysRevA.36.614>.

[39] E. Conti et al. “Correlated Fluctuations Between Luminescence and Ionization in Liquid Xenon”. In: *Phys. Rev. B* 68 (2003), p. 054201. DOI: 10.1103/PhysRevB.68.054201. arXiv:hep-ex/0303008 [hep-ex].

[40] E. Aprile et al. “Observation of Anticorrelation Between Scintillation and Ionization for MeV Gamma Rays in Liquid Xenon”. In: *Phys. Rev. B* 76.1 (July 2007), p. 014115. DOI: 10.1103/PhysRevB.76.014115. arXiv:0704.1118 [astro-ph]. URL: <http://link.aps.org/doi/10.1103/PhysRevB.76.014115>.

[41] M. Berger et al. *XCOM: Photon Cross Section Database (version 1.5)*. 2010. URL: <http://physics.nist.gov/xcom>.

[42] R. Neilson et al. “Characterization of Large Area APDs for the EXO-200 Detector”. In: *Nucl. Instrum. Meth. A* 608.1 (Sept. 2009), pp. 68–75. DOI: 10.1016/j.nima.2009.06.029. arXiv:0906.2499 [physics.ins-det]. URL: <http://www.sciencedirect.com/science/article/pii/S0168900209012376>.

[43] D. Leonard et al. “Systematic Study of Trace Radioactive Impurities in Candidate Construction Materials for EXO-200”. In: *Nucl. Instrum. Meth. A* 591 (2008), pp. 490–509. DOI: 10.1016/j.nima.2008.03.001. arXiv:0709.4524 [physics.ins-det].

[44] 3M HFE-7000. URL: <http://www.3m.com>.

- [45] M Auger et al. “The EXO-200 Detector, Part I: Detector Design and Construction”. In: *J. Instrum.* 7.05 (2012), P05010. doi: 10.1088/1748-0221/7/05/P05010. arXiv:1202.2192 [physics.ins-det]. URL: <http://stacks.iop.org/1748-0221/7/i=05/a=P05010>.
- [46] F. LePort et al. “A Magnetically-Driven Piston Pump for Ultra-Clean Applications”. In: *Rev. Sci. Instrum.* 82 (2011), p. 105114. doi: 10.1063/1.3653391. arXiv:1104.5041 [physics.ins-det].
- [47] *SAES MonoTorr*. URL: <http://www.saespuregas.com>.
- [48] A. Dobi et al. “A Xenon Gas Purity Monitor for EXO”. In: *Nucl. Instrum. Meth. A* 659.1 (2011), pp. 215–228. ISSN: 0168-9002. doi: 10.1016/j.nima.2011.09.017. arXiv:1106.1812 [physics.ins-det]. URL: <http://www.sciencedirect.com/science/article/pii/S0168900211017694>.
- [49] D. North. “An Analysis of the Factors which Determine Signal/Noise Discrimination in Pulsed-Carrier Systems”. In: *Proceedings of the IEEE* 51.7 (July 1963), pp. 1016–1027. ISSN: 0018-9219. doi: 10.1109/PROC.1963.2383.
- [50] O. Bunemann, T. E. Cranshaw, and J. A. Harvey. “Design of Grid Ionization Chambers”. In: *Canadian Journal of Research* 27a.5 (1949), pp. 191–206. doi: 10.1139/cjr49a-019. URL: <http://www.nrcresearchpress.com/doi/abs/10.1139/cjr49a-019>.
- [51] *ANSYS Maxwell*. URL: <http://www.ansoft.com/>.
- [52] E. Aprile et al. *Elementary Processes Affecting Generation of Signals*. Wiley-VCH Verlag GmbH & Co. KGaA, 2006. ISBN: 9783527610020. doi: 10.1002/9783527610020. URL: <http://dx.doi.org/10.1002/9783527610020>.
- [53] G. Bakale, U. Sowada, and W. F. Schmidt. “Effect of an Electric Field on Electron Attachment to Sulfur Hexafluoride, Nitrous Oxide, and Molecular Oxygen in Liquid Argon and Xenon”. In: *J. Phys. Chem.* 80.23 (1976), pp. 2556–2559. doi: 10.1021/j100564a006. URL: <http://pubs.acs.org/doi/abs/10.1021/j100564a006>.

- [54] P. Hough. “Machine Analysis of Bubble Chamber Pictures”. In: *Proceedings of the International Conference on High Energy Accelerators and Instrumentation* (1959), pp. 554–558.
- [55] R. O. Duda and P. E. Hart. “Use of the Hough Transformation to Detect Lines and Curves in Pictures”. In: *Commun. ACM* 15.1 (Jan. 1972), pp. 11–15. ISSN: 0001-0782. DOI: 10.1145/361237.361242. URL: <http://doi.acm.org/10.1145/361237.361242>.
- [56] S. Agostinelli et al. “GEANT4: A Simulation Toolkit”. In: *Nucl. Instrum. Meth. A* 506 (2003), pp. 250–303. DOI: 10.1016/S0168-9002(03)01368-8.
- [57] S. Miyake. “Rapporteur Paper on Muons and Neutrinos”. In: *Proceedings of the 13th International Cosmic Ray Conference* 5 (1973), pp. 3638–3655.
- [58] E.-I. Esch et al. “The Cosmic Ray Muon Flux at WIPP”. In: *Nucl. Instrum. Meth. A* 538 (2005), pp. 516–525. DOI: 10.1016/j.nima.2004.09.005. arXiv:astro-ph/0408486.
- [59] T. Gaisser. *Cosmic Rays and Particle Physics*. Cambridge University Press, 1990. ISBN: 9780521339315. URL: <http://books.google.com/books?id=qJ7Z6oIMqeUC>.
- [60] D. E. Groom, N. V. Mokhov, and S. I. Striganov. “Muon Stopping Power and Range Tables 10 MeV–100 TeV”. In: *Atomic Data and Nuclear Data Tables* 78.2 (2001), pp. 183 –356. ISSN: 0092-640X. DOI: 10.1006/adnd.2001.0861. URL: <http://www.sciencedirect.com/science/article/pii/S0092640X01908617>.
- [61] D.-M. Mei and A. Hime. “Muon-Induced Background Study for Underground Laboratories”. In: *Phys. Rev. D* 73.5 (Mar. 2006), p. 053004. DOI: 10.1103/PhysRevD.73.053004. arXiv:astro-ph/0512125. URL: <http://link.aps.org/doi/10.1103/PhysRevD.73.053004>.

- [62] M. Crouch. “An Improved World Survey Expression for Cosmic Ray Vertical Intensity vs. Depth in Standard Rock”. In: *Proceedings of the 20th International Cosmic Ray Conference*. International Cosmic Ray Conference 6 (1987), pp. 165–168.
- [63] G. Schenter and P. Vogel. “A Simple Approximation of the Fermi Function in Nuclear Beta Decay”. In: *Nucl. Sci. Eng.* 83.3 (Mar. 1983), pp. 393–395.
- [64] E. Aprile, R. Mukherjee, and M. Suzuki. “Performance of a Liquid Xenon Ionization Chamber Irradiated with Electrons and Gamma-Rays”. In: *Nucl. Instrum. Meth. A* 302.1 (1991), pp. 177–85. ISSN: 0168-9002. DOI: 10.1016/0168-9002(91)90507-M. URL: <http://www.sciencedirect.com/science/article/pii/016890029190507M>.
- [65] E. Aprile, R. Mukherjee, and M. Suzuki. “Ionization of Liquid Xenon by ^{241}Am and ^{210}Po Alpha Particles”. In: *Nucl. Instrum. Meth. A* 307.1 (1991), pp. 119–125. ISSN: 0168-9002. DOI: 10.1016/0168-9002(91)90138-G. URL: <http://www.sciencedirect.com/science/article/pii/016890029190138G>.
- [66] A. Ferrari et al. *FLUKA: A Multi-Particle Transport Code (Program Version 2005)*. CERN-2005-010, SLAC-R-773, INFN-TC-05-11. 2005.
- [67] F. James and M. Roos. “Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations”. In: *Comput. Phys. Commun.* 10 (1975), pp. 343–367. DOI: 10.1016/0010-4655(75)90039-9.
- [68] S. S. Wilks. “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses”. In: *The Annals of Mathematical Statistics* 9.1 (1938), pp. 60–62. ISSN: 00034851. URL: <http://www.jstor.org/stable/2957648>.
- [69] W. A. Rolke, A. M. López, and J. Conrad. “Limits and Confidence Intervals in the Presence of Nuisance Parameters”. In: *Nucl. Instrum. Meth. A* 551.2-3 (2005), pp. 493–503. ISSN: 0168-9002. DOI: 10.1016/j.nima.2005.05.068. arXiv:physics/0403059. URL: <http://www.sciencedirect.com/science/article/pii/S016890020501291X>.
- [70] *DURRIDGE RAD7*. URL: <http://www.durridge.com>.

- [71] A. Gando et al. “Measurement of the Double- β Decay Half-Life of ^{136}Xe with the KamLAND-Zen Experiment”. In: *Phys. Rev. C* 85.4 (Apr. 2012), p. 045504. DOI: 10.1103/PhysRevC.85.045504. arXiv:1201.4664 [hep-ex]. URL: <http://link.aps.org/doi/10.1103/PhysRevC.85.045504>.
- [72] D. Mackay. “The Enriched Xenon Observatory, 200kg”. PhD thesis. Stanford University, 2011.