THE SEARCH FOR $0\nu\beta\beta$ IN ^{136}XE WITH EXO-200

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

Kevin O'Sullivan
August 2012

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

Giorgio Gratta, 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.

Roger Blandford

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

Blas Cabrera

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost Graduate Education

Abstract

This analysis reports a search for the neutrinoless double decay ($0\nu\beta\beta$) of ^{136}Xe with the EXO-200 detector. $0\nu\beta\beta$ is a hypothetical nuclear process which only occurs if neutrinos are Majorana particles, which are there own anti-particles, as opposed to Dirac particles which have separate anti-particles. The search for $0\nu\beta\beta$ is currently the most experimentally viable technique to address the Dirac or Majorana nature of neutrinos. The EXO-200 experiment is a liquid xenon time projection chamber operating 655 meters underground at the Waste Isolation Pilot Plant (WIPP) with 200 kg of xenon enriched to 80.6% in the isotope ^{136}Xe. A total of one event in the one sigma region of interest is observed with background of $(1.5 \pm .1) \times 10^{-3}$ counts/keV/kg/yr. This places a limit on the $0\nu\beta\beta$ half-life of ^{136}Xe is of 1.59×10^{25}.

Acknowledgements

To my friends, without whom graduate would not have been half as enjoyable. I would particularly like to thank Chris, Cindy, Andrei, Felix, Anna, Drew, Lindsay, Muni, Lu, Evan, Dan, other Dan, Will, Sean, Erin, Raghvir, and Aurelien. To Karlheinz, Mehmet, John, Scott, and Matt in the Varian machine shop, without whom I wouldn't appreciate the joy, or the art, of metal work. To the entire EXO collaboration, without whom there wouldn't be an experiment. I would especially like to thank some of my collaborators who were so helpful on the analysis, particularly Manuel, Guillaume, Brian, Simon, Mike, Wolfi, and Clayton. To my local collaborators, both at Stanford and SLAC, without whom my work would have felt too difficult to handle. I'd particularly like to thank Marcia (best admin in history), Karl, Michelle, Phil, Maria, Steve, Derek, Jesse, Matt, Axel, Thomas, Igor, Amanda, Delia, Ralph, Liang, Bob, Al, Naoko, Russell, Bjorn, Rob, both Peters, Francisco, and Andrea. To my advisor Giorgio Gratta, without whose guidance I would still be an amateur. To my family, and to my parents, without whom I wouldn't exist.

Contents

Abstract **iv**

Acknowledgements **v**

1 Introduction **1**

2 Massive Neutrinos and $0\nu\beta\beta$ Decay **3**

2.1 Neutrino Oscillations 3

2.2 Majorana Neutrinos 4

2.3 Double Beta Decay 5

2.4 Experimental Challenges of Studying $0\nu\beta\beta$ decay 11

3 The EXO-200 Detector **14**

3.1 Liquid Xenon Time Projection Chambers 14

3.1.1 Measuring Ionization 14

3.1.2 Measuring Scintillation 17

3.1.3 Determining Energy From Both Scintillation and Ionization . 18

3.1.4 Liquid Xenon as a Detector for $0\nu\beta\beta$ 18

3.2 Backgrounds and Construction Considerations 19

3.3 The Time Projection Chamber 20

3.3.1 Detector Regions 21

3.4 Detector Infrastructure 25

3.4.1 The Cryostat 25

3.4.2 The Calibration System 26

3.4.3 Xenon Purification 26
3.5 Detectors Run 1 and 2 29

4 Electronics and Data Acquisition 30
4.1 FECs 30
4.1.1 Wire Gains 30
4.1.2 Signal Shaping 32
4.2 DAQ 38

5 Event Reconstruction 40
5.1 Overview 40
5.2 Signal Finding 40
5.2.1 Matched Filter 40
5.2.2 Charge Reconstruction Threshold 43
5.2.3 Waveform Unshaping 44
5.3 Pulse Fitting and Parameter Extraction 47
5.4 Signal Clustering 48
5.5 Position Determination 49
5.5.1 Drift Velocity Measurements 50
5.6 Pulse Rise Times 51

6 PDF Generation and Background Model 55
6.1 Overview of PDF Generation 55
6.1.1 Particle Tracking 55
6.1.2 Digitization and Reconstruction 56
6.1.3 Energy Spectrum Generation 56
6.2 Backgrounds Model 57
6.2.1 Backgrounds from Copper Vessel and TPC Components . . . 57
6.2.2 Backgrounds from Outside the Copper Vessel 58
6.2.3 Backgrounds Dissolved in the Xenon 58

7 Fits to Data **62**
7.1 Analysis Cuts 62
7.1.1 Fiducial Volume 62
7.1.2 Three Dimensional Position Reconstruction 63
7.1.3 Coincidence Cuts 64
7.1.4 Charge to Light Ratio Cut 64
7.1.5 Other Cuts 66
7.2 Masking Procedure 67
7.3 Position Dependent Corrections 68
7.3.1 Ionization Signal Corrections 68
7.3.2 Scintillation Signal Corrections 69
7.4 Calibrations 71
7.4.1 Pair Production Events 72
7.5 Source Agreement 73
7.5.1 Shape Agreement 75
7.5.2 Rate Agreement 78
7.5.3 Single Site Fraction Agreement 79
7.5.4 Agreement for Different Fiducial Volumes 79
7.5.5 Effect of Three Dimensional Reconstruction Cut on Source Agreement 82
7.5.6 Effect of Triangular Reshaping on Source Agreement 83
7.6 Data Taken Between 4/3/12 and 4/15/12 84
7.7 Fits to Low Background Data 84
7.7.1 ^{135}Xe 86
7.7.2 Systematic Uncertainties 88
7.8 Rates as a Function of Fiducial Cut 90
7.9 Effect of Triangular Reshaping on Low Background Data 91
7.10 $0\nu\beta\beta$ Limit 92
7.11 Comparison of Results to Other Experiments 93

A The Compton Telescope Technique 94
A.1 Principle of a Compton Telescope 94
A.2 Application to Source Calibration Runs 96

Bibliography 101

List of Tables

2.1 Half-life lower limits for $0\nu\beta\beta$ in various isotopes. 8

2.2 Natural abundances and $Q_{\beta\beta}$ of $\beta\beta$ decay candidate isotopes. 12

2.3 Measured radioactivities for various materials. 13

3.1 Calibration source activities at time of purchase and at the beginning of Run 2 (10/5/11). 27

4.1 Shaping times relevant to the transfer functions of different channel types. In Run 2 the final differentiation stage for the u-wires is measured from fits to charge calibration pulses. 33

7.1 χ^2/NDF goodness of fits for the single site and multi site spectra of ^{60}Co and ^{228}Th source agreement. 77

7.2 Percentage of toy Monte Carlos for each PDF which have a *Kolmogorov–Smirnov* distance greater than or equal that of the data from the given PDF. 77

7.3 Rate disagreement as a function of source distance from the detector for ^{60}Co and ^{228}Th sources. Percentage is $\frac{Measured-Expected}{Measured}$ 79

7.4 Source rate and single site fraction agreement versus hexagonal fiducial cut. The Z cuts here are on the absolute value of Z. 80

7.5 Expected number of events from backgrounds near $Q_{\beta\beta}$ of ^{136}Xe from the fit to the low background physics data. 92

List of Figures

2.1 Nuclear binding energies for nuclei with mass number 136. 6

2.2 Simplified Feynman diagram representing the standard model process $2\nu\beta\beta$. 6

2.3 Simplified Feynman diagram representing the hypothetical process $0\nu\beta\beta$ where the nucleus emits two electrons but the neutrinos are not produced because of a virtual annihilation. This particular diagram shows only the phenomenon where $0\nu\beta\beta$ occurs as a result of an exchange of light Majorana neutrinos, but the process can instead proceed by the exchange of putative heavy particles. 7

2.4 Energy spectrum from [1] with a fit to a linear background plus six Gaussian, including four known lines from ^{214}Bi at 2011, 2017, 2022, and 2053 keV, $0\nu\beta\beta$ at 2039 keV, and an unidentified line near 2030 keV. 9

2.5 $m_{\beta\beta}$ versus m_{min} for the inverted hierarchy (labeled "IS") and the normal hierarchy (labeled "NS"). Included are the current limits on $m_{\beta\beta}$ from [2][3][4] and $m_{min} = \frac{1}{3}\Sigma m_i$ from [5]. Figure is taken from [6]. 10

2.6 A hypothetical mechanism for $0\nu\beta\beta$ decay mediated by a "blackbox." 11

2.7 The same "blackbox" mechanism from figure 2.6 can lead to a process which converts a $\bar{\nu_e}$ into a ν_e. 11

2.8 A prototypical energy spectrum for $\beta\beta$ decay showing both the $2\nu\beta\beta$ spectrum and a peak due to $0\nu\beta\beta$. The energy is normalized to be 1 at the $Q_{\beta\beta}$. Figure taken from [7]. 13

3.1 A schematic showing the configuration of the u and v wires in EXO-200. The blue lines show a simulated drift path to the u-wires for an electron cloud. Although the v and u wire planes are shown as parallel, in reality they are at a 60° angle. 15
3.2 Electron drift velocity as a function of electric field for gas, liquid, and solid xenon. EXO-200 operates at a field of 376 V/cm. Figure taken from [8]. 16
3.3 (a) Energy deposition in liquid xenon results in electron ion pairs and excited xenon atoms. Some of the ions recombine, and this results in the creation of dimers. Dimers de-excite to give off scintillation light at 175 nm. (b) The emission spectrum of Xenon in gas, liquid, and solid phase taken from [9]. 17
3.4 A 2D histogram of scintillation versus ionization for a ^{228}Th source calibration of EXO-200. 19
3.5 A schematic of the EXO-200 detector. 21
3.6 The u and v wires are mounted as shown. There is a copper ring (1), to which the six acrylic parts (2) are attached in a hexagonal configuration. Each of the 38 v-wire triplets and 38 u-wires triplets (3) attach to the acrylic parts. The red lines show the configuration of the v-wires and the blue lines show the configuration of the u-wires. . . . 22
3.7 Number of charge deposits versus radius for ^{228}Th source calibrate data (black points) and simulation (red curve). The red curve is normalized to have the same integral as the black points. Outside a radius of 163 mm the number of charge deposits in data drops off, while the number of events in simulation continues to increase. This indicates not all the charge deposits outside 163 mm radius are detected. 23
3.8 Close up image of a region of the detector near the teflon wall. The area circled in red shows the v-wire tabs inside the teflon radius, meaning there is a region where there are no u-wires. 23

3.9 (a) X-Y and (b) X-Z views of the detector highlighting different regions. Shown are the inactive xenon (green), dead layer (or dead xenon) (blue), fiducial xenon (red), and non-fiducial active xenon (grey). . . 24
3.10 A schematic of the copper vessel and the surrounding infrastructure. . 25
3.11 The calibration tubing surrounds the detector providing a way to deploy radioactive sources. The green dots represent typical source locations for calibration runs. 27
3.12 X-Y distribution of events for a ^{228}Th source deployed to positive Y. . 28
3.13 Electron lifetime and xenon recirculation rate over time. 29

4.1 A schematic showing the structure of the EXO-200 electronics. 31
4.2 A histogram of pulse amplitudes for a precision pulser calibration of a u-wire channel. 32
4.3 V-wire gains measured for each channel with errors. 33
4.4 A schematic of the first four shaping stages for the u-wires. 34
4.5 Pulse shaping of a u-wire signal. A raw pulse (1) is a step function. This pulse is shaped with a 1.5μs integrator (2), a 40μs differentiator (3), a second 1.5μs integrator (4), a second 40μs differentiator (5), and finally a 60μs differentiator to arrive at the final pulse shape (6). . . . 35
4.6 37
4.7 U-wire differentiation times for each u-wire channel. 38

5.1 (a) A u-wire pulse from a detector event. (b) the matched filter output for this waveform. The $\frac{peak}{RMS}$ is 3.6 for the raw waveform and 8.2 for the matched filter output. 42
5.2 Fit of the low energy tail of reconstructed charge clusters to an error function to determine charge reconstruction threshold. 43
5.3 (a) a raw waveform from a u-wire channel with two pulses (b) the matched filter output for this waveform (c) an unshaped sub-trace of this waveform (d) the sub-trace after application of a 2μs triangular filter. 45

5.4 (a) Single site spectrum for thorium source calibration simulation (black) without the reshaping algorithm and the energy spectrum of events moved from single site to multiple site by the algorithm scaled by a factor of ten (red). (b) Fraction of events moved from single site to multi site versus energy for thorium source calibration simulation. (c) Fraction of events moved from single site to multi site in $2\nu\beta\beta$ decay simulation versus energy. In all cases the energy is the energy is taken to be the energy measured without the reshaping algorithm. 46
5.5 The fraction of events moved from single site to multi site by the triangular reshaping algorithm in thorium calibration data versus the energy of the event measured without the reshaping algorithm. 47
5.6 The ratio of the reconstructed energy when using the reshaping algorithm, E_1, to the reconstructed energy when not using the reshaping algorithm, E_2 from ^{228}Th source calibration data for event moved from single site to multi site. 48
5.7 Electron drift time in TPC2. The full spectrum is shown in (a) whereas (b) is a zoom in on the falling edge. There are some events with drift times longer than the measured maximum drift time, these are events where a charge cluster has been associated with the wrong scintillation signal. There is also a peak close to a drift time of zero. This peak exists because some charge deposits occur between the APD plane and the u-wires. This is a 6 mm gap between and drift times of these clusters can be a few μs. This peak contains both charge clusters from between the u-wires and the v-wires and from between the u-wires and the APD plane. 50
5.8 Electron drift velocities based on this drift time end point based method (black) and using muons which pass through both halves of the TPC (red). 51

5.9 The stages of unshaping of a raw u-wire pulse. (a) is the raw pulse, stages (b)-(d) are the unshaping of the 1.5 μs integration, 40 μs differentiation, 1.5 μs integration again, and 40 μs differentiation again. The resulting pulse at (e) is a fast rise exponential decay and is used for rise time measurement. 52
5.10 Shows rise time versus pulse amplitude for u-wire pulses from thorium source calibration data (red), low background data (green), and u-wire charge injection runs (black). The amplitude is in ADC counts. To translate this to an energy, in the charge one ADC counts corresponds to 5.23 keV after calibration on average. For example, a 2458 keV charge cluster would have an amplitude of roughly 470 ADC counts. This number is channel dependent due to gain variations 53
5.11 Example of a raw waveform with large rise time. In this example a two charge deposits are apparent, one roughly 20μs earlier than the other. 54

6.1 V-wire detection threshold is determined by taking the fraction of charge clusters with given energy that have a detected v-wire signal for both simulation (dashed blue line) and data (black points). The red line is a fit a sigmoid function to the data. V-wire detection is fully efficient above 350 keV. The blue line shows the effective detection threshold in simulation after a stochastic application of the threshold from data. 57
6.2 The ^{222}Rn decay chain through ^{210}Pb. Vertical arrows indicate α decays and diagonal lines indicate β decays. Half-lives and decay Q-values are shown. 59
6.3 Ratio of ionization to scintillation for β's (black) and α's (red) taken from ^{214}Bi to ^{214}Po coincidences. 60

7.1 The fraction of events passing all other cuts that have full three dimensional reconstruction for single site events in data (black) and simulation (red) for (a) ^{60}Co, (b) ^{228}Th, and (c) low background physics data, which is compared to $2\nu\beta\beta$ simulation. 90% of single site low background events are $2\nu\beta\beta$, so this last comparison is reasonable. While the error bars on the data are purely statistical, the errors bars on the simulation include a 9.4% uncertainty determined from calibration source agreement studies (see section 7.5.2). 65

7.2 Scintillation signal versus ionization signal for ^{228}Th events. The dashed red lines shows the threshold for the light to charge ratio cut. The red axis represents the the energy scale for the calibrated energy after combining the scintillation and ionization signals (see section 7.4). 66

7.3 Scintillation signal versus ionization signal for low background physics data. The dashed red line represents the light to charge ratio cut. . . 67

7.4 Fits to the ^{208}Tl double escape peak for two different u-wires channels. 69

7.5 Fitted peak position versus Z for ^{232}Th source calibration, which is fit to an exponential to extract the electron lifetime. 70

7.6 ^{232}Th source calibration versus Z after the purity correction with a fit to the $\frac{1}{d}$ dependence for shielding grid inefficiency. 70

7.7 Top: Residuals for energy calibrations for single site events (circles, solid lines) and multi site events (boxes, dashed lines). The lines show the one sigma error band for each for each calibration. Bottom: Energy resolution as a function of energy for single site (circles, solid lines) and multi site (boxes, dashed lines). 72

7.8 Measured peak position using the detector calibration divided by true peak energy over time. Deviation is less than 0.4% over all of Run 2. 73

7.9 Fits to the (a) 511 keV and (b) 1593 keV charge only peaks from pair production events. 74

7.10 Residuals for ionization signal only energy calibration including the 1593 keV and 511 keV peaks from pair production events. 75

7.11 Maximum likelihood fit of a ^{228}Th source calibration PDF (blue line) to calibration data (black points). 76

7.12 Maximum likelihood fit of a ^{60}Co source calibration PDF (blue line) to calibration data (black points). 76

7.13 Single Site fraction discrepancy between simulation and data for ^{60}Co and ^{228}Th as a function of energy bin. The largest discrepancy observed is 8.5%. 80

7.14 Fits of ^{228}Th PDFs to data for different fiducial cuts. The fiducial cuts are hexagonal, and defined by the apothem of the hexagonal (the radius of the circle inscribed in the hexagon). The different cuts shown are (a) 133mm, (b) 103mm, (c) 73mm, and (d) 43mm. 81

7.15 Source agreement plots where the three dimensional reconstruction cut is not applied in simulation (blue curve) or data (black points). . . . 82

7.16 Source agreements plots for (a) ^{60}Co and (b) ^{228}Th without the triangular reshaping algorithm applied in reconstruction. The reduced χ^2's are included. 83

7.17 Comparison of ^{228}Th source data before April 3rd to a weighted average of source data over the course of Run 2. 85

7.18 (a) Fit to low background physics data to a background model plus $2\nu\beta\beta$ and $0\nu\beta\beta$. The grey area is the fit to $2\nu\beta\beta$. Other components include: ^{232}Th in the copper vessel (dotted magenta), ^{222}Rn in the air around the cryostat (dashed green), ^{238}U in the vessel (dotted black), ^{214}Bi on the cathode (dashed cyan), ^{222}Rn dissolved in the inactive xenon (dotted turquoise), ^{222}Rn in the active liquid xenon (dashed brown), ^{60}Co in the vessel (dotted blue), ^{54}Mn in the vessel (dotted brown), ^{40}K in the vessel (dotted orange), and ^{135}Xe in the active xenon (dashed blue). (b) A zoom in on the plot in (a) on the area around $Q_{\beta\beta}$ of ^{136}Xe. The solid vertical red line represents the one sigma window around $Q_{\beta\beta}$ whereas the dashed red lines show the two sigma window. 87

7.19 Profile likelihood of ^{135}Xe in the fit to the low background physics data with a fit to a parabola. Plot shows the difference between the negative log likelihood (NLL) from the minimum NLL versus number of ^{135}Xe events in the fit. The profile measures a total of 635 ± 135 ^{135}Xe events in the dataset, 4.7σ different from zero. 88

7.20 The single site fraction pulls for the various PDFs. The pull here is defined as the difference between the fit value for the single site fraction and the predicted single site fraction from Monte Carlo divided by the measured single site fraction. 89

7.21 Rates versus fiducial cut coming from maximum likelihood fits to the low background physics data for various signals. Shown: $2\nu\beta\beta$ (black triangles), ^{40}K in the vessel X10 (red triangles), ^{232}Th in the vessel X10 (green circles), ^{238}U in the vessel X10 (blue squares). 90

7.22 Black: faction of single site events that were moved to the multi site spectrum by the triangular reshaping algorithm in signal finding as a function of energy. Red: fraction of masked single site events which are gamma events according to the maximum likelihood fits to the low background data as a function of energy. 91

7.23 Comparison of $0\nu\beta\beta$ limits for pairs of isotopes for various RQRPA (labeled RQRPA-1) [10], IBM-2 [11], GCM [12], NSM [13], and QRPA-2 [14]. Both plots show the the 90% confidence level limit from KamLAND-ZEN [15]. (a) Comparison of ^{136}Xe half-life limit on the x-axis to ^{76}Ge half-life on the y-axis, including the limit from the Heidelberg-Moscow experiment [2]. Also included is the 68% confidence interval for the claimed discovery by a subset of the Heidelberg-Moscow experiment [1]. (b) Comparison of ^{136}Xe half-life limit on the x-axis to ^{130}Te half-life on the y-axis. Included is the best limit from the COURICINO experiment [3]. 93

A.1 A three site Compton scatter with the energy and spatial coordinates measured at each site in the detector. In the case of a three site event this defines two vectors in space, between site one and site two, the other between site two and site three. Using the spatial coordinator the angle between these two vectors can be calculated. 96
A.2 Compton Telescope image of a calibration source at positive y. These images are used to identify the source locations, which is then used to generate simulated calibration runs for comparison to data. 98
A.3 A horizontal slice of the 2D histogram shown in figure A.2, which is fit to a Gaussian to determine the location of the point-like source. . . . 99
A.4 Z distribution of events for a source calibration run at positive x. . . 100

Chapter 1

Introduction

This dissertation describes a search for neutrinoless double beta decay in ^{136}Xe with the EXO-200 experiment. Double beta decay is a process where a nucleus emits two electrons and (typically) two anti-electron neutrinos, changing its charge by two units. If neutrinos are their own anti-particles, the putative process of nuetrinoless double beta decay becomes possible. From studies of neutrino oscillations in atmospheric neutrinos, solar neutrinos, reactor neutrinos, and neutrino beams, neutrinos are known to be massive [16] [17] [18] [19] [20]. The mass scale for neutrinos is heavily constrained, both from terrestrial experiments and cosmological measurements [21] [22] [23]. According to these limits, neutrinos are at least five orders of magnitude lighter than the next lightest particle, the electron, which may point to new physics in the neutrino sector. Understanding whether or not neutrinos are their own anti-particles can help shed light on the physics behind their mass scale.

EXO-200 is a liquid phase time projection chamber using 200 kg of xenon enriched to 80.6% in the isotope ^{136}Xe, housed underground at the WIPP site near Carlsbad, NM. The mine's overburden provides shielding from cosmic rays. The detector is shielded from external gamma rays by 25 cm of lead and 50 cm of HFE-7000, a fluorocarbon heat transfer fluid, on all sides and all construction materials are qualified for low radioactivity. In Run 1, EXO-200 was the first detector to observe the two neutrino double beta decay of ^{136}Xe.

This thesis focuses on Run 2, with better energy resolution and lower background than Run 1. The data places a constraint of $T_{1/2}^{0\nu} > 1.59 \times 10^{25}$ years for ^{136}Xe. ^{76}Ge is the only other isotope for which the limit on its half-life is longer than 10^{25} years.

Chapter 2

Massive Neutrinos and $0\nu\beta\beta$ Decay

2.1 Neutrino Oscillations

From neutrino oscillation experiments, neutrinos are known to have three mass eigenstates ν_1, ν_2, and ν_3 which are different from their flavor eigenstates ν_e, ν_μ, and ν_τ, with squared mass splittings $\Delta m^2_{21} = (7.58^{+.14}_{-.13}(\text{stat}) \pm .15(\text{sys})) \times 10^{-5} eV$ and $\Delta m^2_{32} = (2.43 \pm .13) \times 10^{-3} eV$ [16] [17]. That $m_1 < m_2$ is known from solar neutrino studies, but whether $m_3 > m_1$ (call the normal hierarchy) or $m_3 < m_1$ (called the inverted hierarchy) is unknown. The mass of a given neutrino flavor state can be written

$$\nu_i = \sum_{j=0}^{3} U_{ij} m_j \tag{2.1}$$

where U is the neutrino mixing matrix

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} e^{i\alpha/2} & 0 & 0 \\ 0 & e^{i\beta/2} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{2.2}$$

$c_{ij} = cos(\theta_{ij})$, $s_{ij} = sin(\theta_{ij})$, and δ is a CP-violating phase. The matrix is broken down into four components: one for each pair of mass eigenstates and a fourth component of phases α and β which are allowed if neutrinos are Majorana particles, which are their own anti-particles. Particles which have separate anti-particles are referred to as Dirac particles. All three mixing angles have been measured and are non-zero [18] [19] [20]. Although most of the parameters of this matrix have been measured, the nature of the Majorana phases α and β remain unknown (in addition to δ), as there is no definitive evidence as to whether or not neutrinos are Majorana or Dirac in nature.

The current upper limit from the Mainz and Troitsk tritium endpoint experiments is m_{ν_e} <2 eV at the 90% confidence level [23]. This is less than one part in 10^5 of the mass of the next lightest particle, the electron. Such a large difference in mass scales can be explained if neutrinos are Majorana particles.

2.2 Majorana Neutrinos

Neutrinos are intrinsically parity violating particles, with all ν produced with left-handed helicity and all $\bar{\nu}$ produced with right-handed helicity. A proper pair of particle and anti-particle will have opposite charge and opposite helicity. Because neutrinos are neutral, it is possible that ν and $\bar{\nu}$ are the same particle. In this case, the left-handed state would is what is referred to as ν and the right-handed state is referred to as $\bar{\nu}$.

The Lagrangian for neutrino masses containing both Dirac and Majorana mass terms is

$$\mathcal{L}_{D+M} = \mathcal{L}_D + \mathcal{L}_M = -m_D(\overline{\Psi_L}\Psi_R + H.c.) - \frac{m_L}{2}[\overline{(\Psi_L^c)}\Psi_L + H.c.] - \frac{m_R}{2}[\overline{(\Psi_R^c)}\Psi_R + H.c.] \tag{2.3}$$

Defining

$$n_L = \begin{bmatrix} \Psi_L \\ (\Psi_R)^c \end{bmatrix} \tag{2.4}$$

2.3 becomes

$$\mathcal{L}_{D+M} = -\frac{1}{2}\overline{(n_L)^c}\mathcal{M}n_L + H.c. \tag{2.5}$$

where

$$\mathcal{M} = \begin{pmatrix} m_L & m_D \\ m_D & m_R \end{pmatrix} \tag{2.6}$$

is the neutrino mass matrix, m_L and m_R are the left-handed and right-handed Majorana mass terms respectively, and m_D is the Dirac mass term. If $m_R \gg m_D \gg m_L$, diagonalizing the matrix yields two eigenvalues, $m_1 \approx -\frac{m_D^2}{m_R}$ and $m_2 \approx m_R$. This model can be extended to N neutrino flavors with the same result [24] [25]. This model connects the lightness of the known neutrinos to the heaviness of some yet unknown, heavy right-handed neutrino. The known neutrino has mass m_1 which is proportional to $\frac{1}{m_2}$, so the heaviness of one particle suppresses the mass of the other. This model is referred to as the Seesaw Mechanism and requires that neutrinos are Majorana particles.

Differentiating between Dirac and Majorana neutrinos is challenging because in the massless limit the two cases are indistinguishable. This is referred to as the Confusion Theorem. For example, the amplitude for ν to $\bar{\nu}$ oscillations from a beam is suppressed by a factor of $(\frac{m_\nu}{E})^2$, where E is the energy of the beam [26] [27]. As neutrino masses will be negligible compared to any energy scale, the difference between Majorana and Dirac neutrinos in the majority of experiments will be heavily suppressed. The best known technique for probing this question is to study the nuclear process of double beta decay.

2.3 Double Beta Decay

Figure 2.1 shows the nuclear binding energy versus atomic number, Z, for nuclei with mass number, A, equal to 136 amu. ^{136}Xe is less massive than ^{136}Cs and therefore its β-decay to ^{136}Cs is energetically forbidden. However, ^{136}Xe is heavier than ^{136}Ba, and there is a second order process known as double beta decay ($\beta\beta$-decay) whereby

$$^{A}_{Z}N \rightarrow ^{A}_{Z+2} N + 2e^- + 2\bar{\nu}_e \tag{2.7}$$

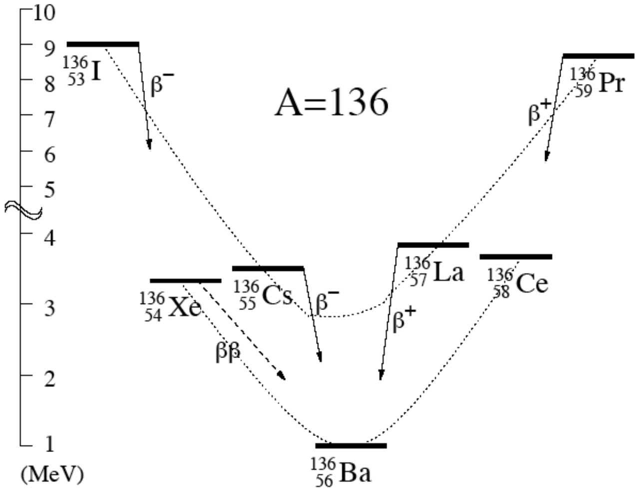

Figure 2.1: Nuclear binding energies for nuclei with mass number 136.

where $\bar{\nu_e}$ is the electron anti-neutrino. This is the slowest observed nuclear process in nature, with decay half-lives ranging from 7.1×10^{18} years for ^{100}Mo to 7.7×10^{24} years for ^{128}Te [28][29]. Even the half-life of ^{100}Mo is more than 10^8 times the age of the universe. Figure 2.2 shows a simplified Feynman diagram for this process, typically referred to as two neutrino double beta decay ($2\nu\beta\beta$). This is a standard model process, observed in several isotopes.

If neutrinos were Majorana particles, this would imply finite amplitude for the

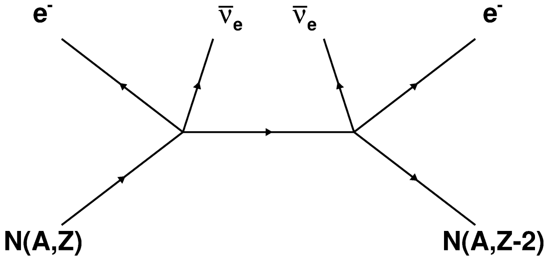

Figure 2.2: Simplified Feynman diagram representing the standard model process $2\nu\beta\beta$.

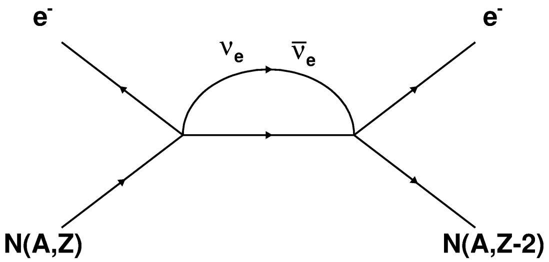

Figure 2.3: Simplified Feynman diagram representing the hypothetical process $0\nu\beta\beta$ where the nucleus emits two electrons but the neutrinos are not produced because of a virtual annihilation. This particular diagram shows only the phenomenon where $0\nu\beta\beta$ occurs as a result of an exchange of light Majorana neutrinos, but the process can instead proceed by the exchange of putative heavy particles.

process described in figure 2.3 where there is a virtual annihilation of electron neutrinos. The resulting decay is referred to as neutrinoless double beta decay ($0\nu\beta\beta$). The observation of the this process would imply:

- Neutrinos are Majorana particles
- Lepton number is not conserved (as this is a ΔL=2 process)
- B-L is not conserved (Δ(B-L)=-2)

Limits on the $0\nu\beta\beta$ half-lives of several isotopes obtained before the present work are listed in table 2.1. ^{76}Ge is the only isotope for which the limit is longer than 10^{25} years.

There is a claimed evidence for $0\nu\beta\beta$ in ^{76}Ge at a half-life of $(2.23^{+.44}_{-.31})\times 10^{25}$ years with 4.2 σ significance by part of the Heidelberg-Moscow collaboration [1]. The energy spectrum near the Q value is shown in figure 2.4. This result is controversial [33][34][35][36] for several reasons

- The authors of [1] have performed multiple analyses claiming evidence for $0\nu\beta\beta$ [37][38][39] and determine the probability of a peak in the data at any given energy. Although this technique does indicate a statistical significance for a peak

Table 2.1: Half-life lower limits for $0\nu\beta\beta$ in various isotopes.

Isotope	90% CL Limit	Technique	Reference
^{48}Ca	5.8×10^{22} yr	CaF$_2$ scintillating crystal	ELEGANT VI Collaboration [30]
^{76}Ge	1.9×10^{25} yr	High purity Ge detector	Heidelberg-Moscow Experiment [2]
^{82}Se	3.6×10^{23} yr	Thin-foil tracking calorimeter	NEMO-3 Collaboration [4]
^{96}Zr	9.2×10^{21} yr	Thin-foil tracking calorimeter	NEMO-3 Collaboration [31]
^{100}Mo	1.1×10^{24} yr	Thin-foil tracking calorimeter	NEMO-3 Collaboration [4]
^{116}Cd	1.7×10^{23} yr	^{116}CdWO$_4$ crystal scintillators	Solotvina Experiment [32]
^{130}Te	3.0×10^{24} yr	TeO$_2$ bolometers	COURICINO Collaboration [3]
^{136}Xe	5.7×10^{24} yr	Xenon doped scintillator	KamLAND-ZEN Collaboration [15]
^{150}Nd	1.8×10^{22} yr	Thin-foil tracking calorimeter	NEMO-3 Collaboration [4]

at 2039 keV, it also indicates statistical significance for peaks with unidentified origins.

- [35], [34], and [36] reanalyze the data and find less statistical significance for the peak than do the authors of [38][39].
- The authors of [38][39][1] do not study the effects of changing their analysis window on the result [33].

If $0\nu\beta\beta$ occurs and is due to a virtual annihilation of light Majorana neutrinos as shown in figure 2.3, one over its half-life, $(T^{0\nu}_{1/2})^{-1}$, can be calculated:

$$(T^{0\nu}_{1/2})^{-1} = G_{0\nu}(Q_{\beta\beta}, Z) \mid M_{0\nu} \mid^2 \langle m_{\beta\beta} \rangle^2 \tag{2.8}$$

where $G_{0\nu}(Q_{\beta\beta},Z)$ is a phase space factor, which is a function of the Q-value of the decay, $Q_{\beta\beta}$, and Z. $M_{0\nu}$ is a nuclear matrix element and $\langle m_{\beta\beta} \rangle$ is the Majorana mass of the electron neutrino, defined as

$$\langle m_{\beta\beta} \rangle = \sum_{i=1}^{3} \mid U_{e,i} \mid^2 m_i \epsilon_i \tag{2.9}$$

where m_i are the values of the three mass eigenstates, and ϵ_i are phase factors. A limit on the half-life of $0\nu\beta\beta$ decay translates into a limit on the mass that the electron neutrino could have were it a Majorana particle, assuming no cancellations. Figure 2.5 shows the possible combinations of $m_{\beta\beta}$ and m_{min}, the mass of the lightest neutrino

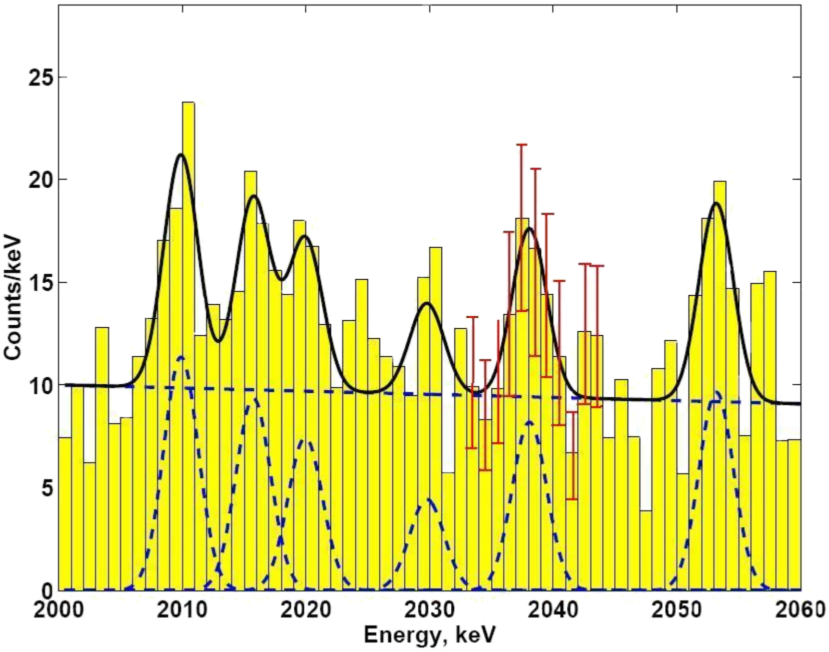

Figure 2.4: Energy spectrum from [1] with a fit to a linear background plus six Gaussian, including four known lines from ^{214}Bi at 2011, 2017, 2022, and 2053 keV, $0\nu\beta\beta$ at 2039 keV, and an unidentified line near 2030 keV.

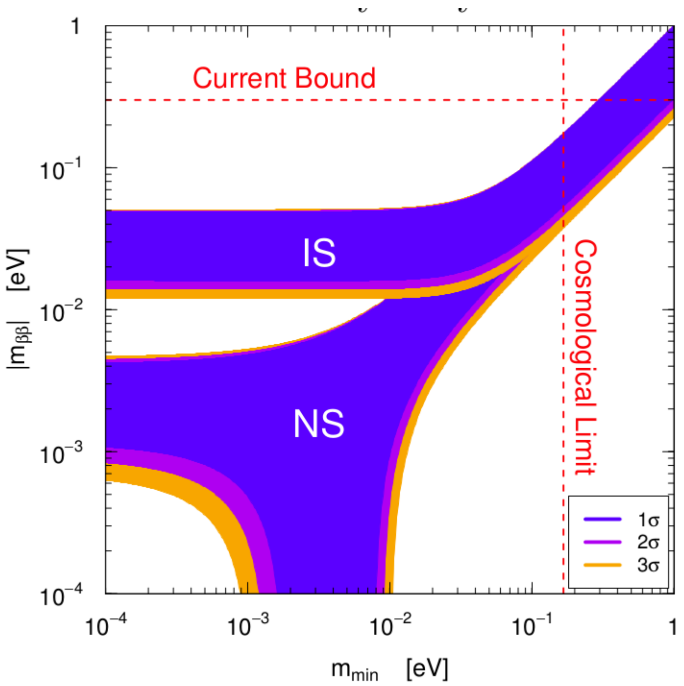

Figure 2.5: $m_{\beta\beta}$ versus m_{min} for the inverted hierarchy (labeled "IS") and the normal hierarchy (labeled "NS"). Included are the current limits on $m_{\beta\beta}$ from [2][3][4] and $m_{min} = \frac{1}{3}\Sigma m_i$ from [5]. Figure is taken from [6].

(m_1 in the normal hierarchy, and m_3 in the inverted hierarchy). Included are current limits on $m_{\beta\beta}$ from [2][3][4] and m_{min}, from cosmology, which places a limit $\Sigma m_i <$ 0.5 eV[5]. The figure assumes the limit on m_{min} is $\frac{1}{3}\Sigma m_i$. Although this also places a more stringent limit on the mass of ν_e than tritium end point experiments, it is also dependent on cosmological models.

The calculation of nuclear matrix elements is non-trivial and there are several different competing models [10] [11] [12] [13] [14]. There can be a factor of 2-3 difference between these models. Even if the matrix elements could be calculated exactly, their application is dependent on the assumption that a light Majorana neutrino exchange is responsible for neutrinoless double beta decay. Studies of lepton flavor violating processes beyond the standard model will also affect the interpretation of these results [40]. However, even if $0\nu\beta\beta$ decay were to arise from a process that did not involve

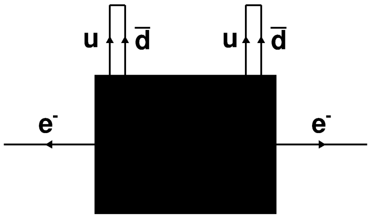

Figure 2.6: A hypothetical mechanism for $0\nu\beta\beta$ decay mediated by a "blackbox."

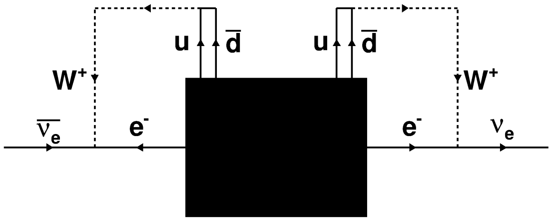

Figure 2.7: The same "blackbox" mechanism from figure 2.6 can lead to a process which converts a $\bar{\nu_e}$ into a ν_e.

light neutrino annihilation, its discovery would still imply that neutrinos are massive Majorana particles.

Consider figure 2.6 where $0\nu\beta\beta$ decay proceeds by some totally unknown ("black box") mechanism. Regardless of what this mechanism is, it implies finite amplitude for the process shown in figure 2.7 which converts a neutrino into an anti-neutrino [41]. This means the discovery of $0\nu\beta\beta$ decay implies that neutrinos are Majorana particles, regardless of the microscopic process mediating it.

2.4 Experimental Challenges of Studying $0\nu\beta\beta$ decay

The Enriched Xenon Observatory (EXO) searches for the $0\nu\beta\beta$ decay of ^{136}Xe, which has a $Q_{\beta\beta}$ of 2458 keV [42]. Properties of select $0\nu\beta\beta$ candidate isotopes are listed

Table 2.2: Natural abundances and $Q_{\beta\beta}$ of $\beta\beta$ decay candidate isotopes.

Isotope	Natural Abundance	$Q_{\beta\beta}$ (keV)
^{48}Ca	0.2%	4271
^{76}Ge	7.7%	2039
^{82}Se	8.7%	2995
^{96}Zr	2.8%	3350
^{100}Mo	9.8%	3034
^{116}Cd	7.5%	2805
^{130}Te	34.1%	2530
^{136}Xe	8.9%	2458
^{150}Nd	5.6%	3367

in table 2.2. EXO-200 is the first stage of the EXO project, using 200 kg of xenon enriched in the isotope ^{136}Xe. Because $\beta\beta$-decay is so rare, even ambient radioactivity is a background for the experiment. Table 2.3 compares measured activities for various materials, including ^{136}Xe. Not only must a search for $0\nu\beta\beta$ be shielded from external radiation, even components of the experiment itself can easily be too radioactive. Although stainless steel has several advantages as a construction material, it is not used for parts close to the detector because it is too radioactive, copper and copper alloys are used instead. Teflon is a good choice of insulator, whereas most ceramics are not. Even 1L of air in the EXO-200 clean rooms (.122 kg) is several orders of magnitude more radioactive than the xenon because of its radon content, and this is a background for EXO-200 (see section 7.10).

$2\nu\beta\beta$ is also a background for $0\nu\beta\beta$. Because in $0\nu\beta\beta$ no neutrinos are produced, all the energy goes into the two electrons, resulting in a delta function peak at $Q_{\beta\beta}$, smeared by the energy resolution of the detector. Figure 2.8 shows an example energy spectrum for $2\nu\beta\beta$ and $0\nu\beta\beta$. The end point of $2\nu\beta\beta$ is at the energy of $0\nu\beta\beta$ and the narrower the $0\nu\beta\beta$ peak, the better the two can be separated.

Table 2.3: Measured radioactivities for various materials.

Material	mBq/kg	Reference
Copper from EXO-200 Vessel	<1.9	[43] and EXO-200 Run 2 Data
Stainless Steel from GERDA	6.6-66.9	[44]
EXO-200 Teflon reflector	<.06	[45]
Macor Insulator	$.8\text{-}18.7\times10^5$	[43]
Air in EXO-200 clean rooms	4900-6500	Monitored in situ
^{136}Xe ($2\nu\beta\beta$)	.046±0.5	[46]
^{136}Xe ($0\nu\beta\beta$)	$<1.7\times10^{-5}$	[15]

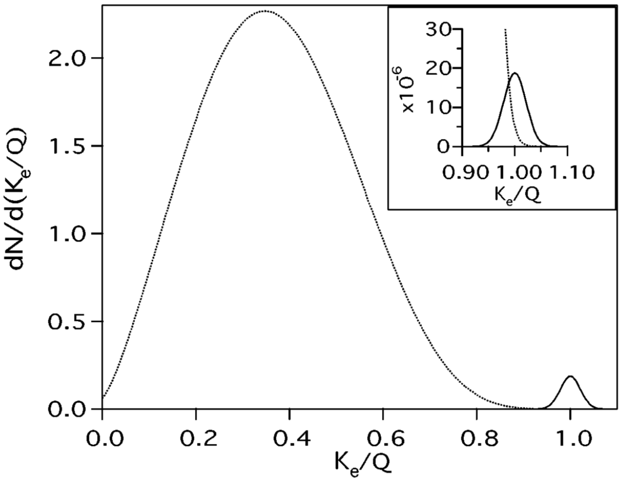

Figure 2.8: A prototypical energy spectrum for $\beta\beta$ decay showing both the $2\nu\beta\beta$ spectrum and a peak due to $0\nu\beta\beta$. The energy is normalized to be 1 at the $Q_{\beta\beta}$. Figure taken from [7].

Chapter 3

The EXO-200 Detector

3.1 Liquid Xenon Time Projection Chambers

When ionizing radiation deposits energy in liquid xenon, it creates electron ion pairs and excited xenon atoms. This leads to two separate channels from which the energy of the event can be determined. First, because xenon is a noble gas, it is available at a purity high enough that charge can be drifted to an anode and measured directly. Second, excited xenon atoms, and xenon ions which recombine, produce scintillation light. By measuring the amount of charge liberated or the amount of light produced, one can measure the energy deposited by a given event. The measurement can be made even more precise by simultaneously measuring the ionization and scintillation signals.

3.1.1 Measuring Ionization

Figure 3.1 shows a schematic of the charge collection in EXO-200. Charge is drifted past the v-wires, which see a signal from induction but otherwise let the electrons through, and is collected on the u-wires. Based on which u and v channels see signals, the position of the charge deposit the in the U-V plane can be determined. In EXO-200 the u and v wires are strung as triplets, with three wires to a channel (see figure 3.6). The u and v wires are not orthogonal, they are at a 60° angle with respect to

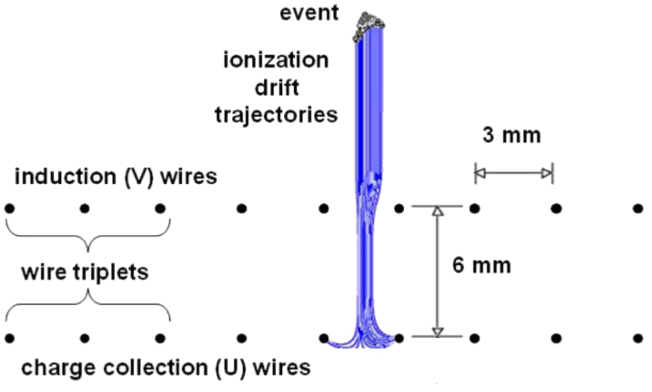

Figure 3.1: A schematic showing the configuration of the u and v wires in EXO-200. The blue lines show a simulated drift path to the u-wires for an electron cloud. Although the v and u wire planes are shown as parallel, in reality they are at a 60° angle.

each other, and the U and V coordinates are translated to orthogonal coordinates X and Y.

The third spatial coordinate, Z, is reconstructed by measuring the drift time of the charge. Electrons drifting through liquid xenon see an effective viscosity and have a terminal velocity referred to as the drift velocity. The drift velocity is a function of the electric field, and has been measured to be 1.71 mm/μs for in EXO-200, which is consistent with previous work [8], also describing drift velocities for a variety of electric fields in xenon as shown in figure 3.2. The increase in drift velocity with electric field is faster at low electric fields but the increase is slower above 1 kV/cm. The event drift time is taken to be the time between when the scintillation is observed and when the charge is collected, as the scintillation is measured within nanoseconds of the event.

Spatial reconstruction of charge deposits in X, Y, and Z makes discrimination between single site events, which have only one charge deposit, and multi site events, which have more than one charge deposit, possible. This is useful for background rejection as $0\nu\beta\beta$ events are predominately single site. Events which contain multiple scatters close together may appear to be single site to the analysis, these will also be referred to as single site. As will be discussed in more detail in section 3.3.1, position

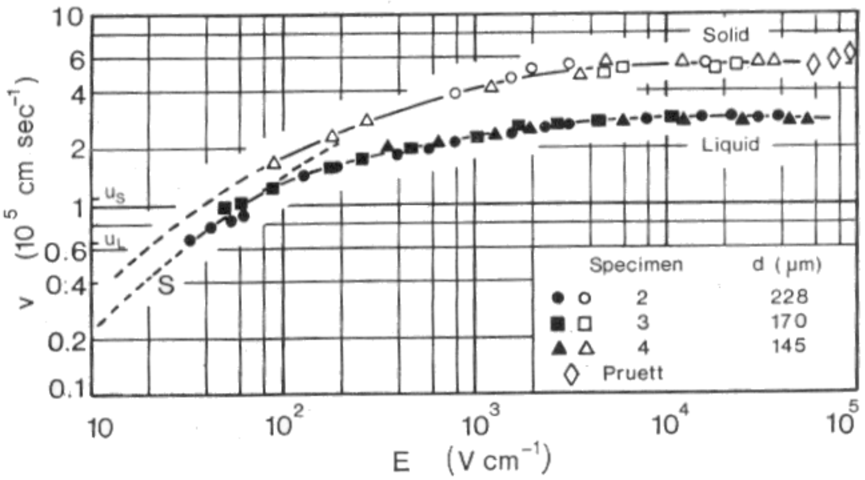

Figure 3.2: Electron drift velocity as a function of electric field for gas, liquid, and solid xenon. EXO-200 operates at a field of 376 V/cm. Figure taken from [8].

reconstruction also allows for rejecting events in volumes within the detector that are less well understood.

In order to have efficient charge collection, high xenon purity is necessary. The purity of the xenon in such a detector is often parameterized by the electron lifetime, defined as the mean time for which an electron can drift in the xenon before being absorbed by an electronegative impurity. If a total charge Q drifts for a time t in liquid xenon with an electron lifetime τ, then the charge collected at the anode is

$$Qe^{-\frac{t}{\tau}} \tag{3.1}$$

The longer the electron lifetime, the less charge that is lost to impurities and the more charged that is detected. A correction for charge loss to impurity as a function of charge deposit drift time is also typically necessary.

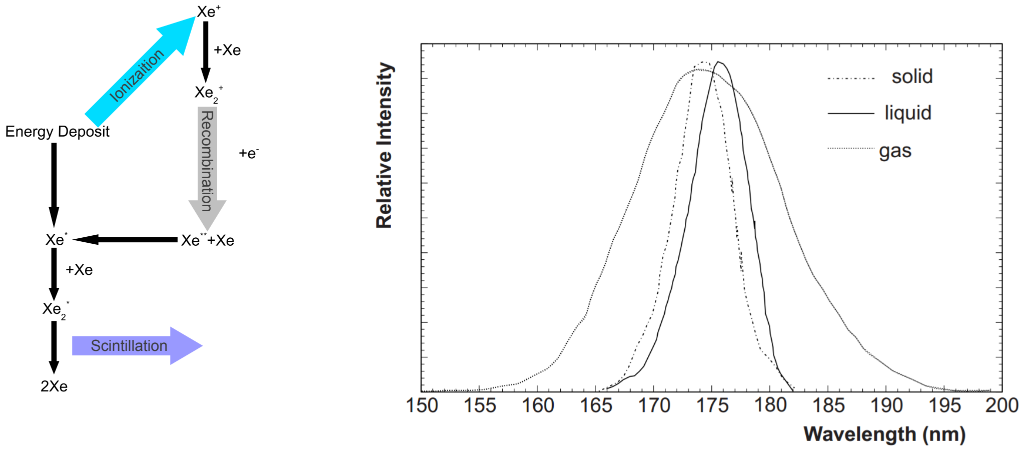

(a)

(b)

Figure 3.3: (a) Energy deposition in liquid xenon results in electron ion pairs and excited xenon atoms. Some of the ions recombine, and this results in the creation of dimers. Dimers de-excite to give off scintillation light at 175 nm. (b) The emission spectrum of Xenon in gas, liquid, and solid phase taken from [9].

3.1.2 Measuring Scintillation

Figure 3.3 shows the mechanism for the production of scintillation light in liquid xenon and the resulting emission spectrum. Incident radiation creates electron ion pairs and excited xenon atoms. The excited xenon atoms can combine with ground state xenon atoms to form excited xenon dimers which de-excite to give off 175nm light. Xenon ions can also lead to xenon dimers as some of these ions will recombine to form excited state xenon atoms. Because this is a molecular transition, the wavelength of photon emitted does not correspond to an absorption line of atomic xenon and the xenon will not reabsorb its own scintillation.

3.1.3 Determining Energy From Both Scintillation and Ionization

Because electrons recombining with xenon ions creates a scintillation signal, there is a trade-off between scintillation and ionization in such a detector. For every recombined electron there is one less unit of charge which can be detected and one more unit of light. From event-to-event, even at a fixed energy, the fraction of electrons which recombine will fluctuate. This implies that if one uses only ionization or only scintillation to measure the energy of an event, the apparent energy will be affected by the amount of recombination. This worsens the energy resolution, or the energy uncertainty, when measuring charge or light alone.

As shown by the EXO collaboration [47], and later by [48], one can improve the energy resolution in liquid xenon detectors by measuring both the ionization and scintillation signals. Figure 3.4 shows a ^{228}Th source calibration of EXO-200. The 2615 keV gamma line from ^{208}Tl (a progeny of ^{228}Th) is visible as an island and there is an anti-correlation between the detected charge and the detected light. Using both charge and scintillation to measure the energy, the energy resolution, $\frac{\sigma}{E}$, is 1.6% at 2615 keV, as opposed to 3.6% when using only the ionization.

3.1.4 Liquid Xenon as a Detector for $0\nu\beta\beta$

As a detector for a low radioactivity experiment, such as a search for $0\nu\beta\beta$, a liquid xenon TPC is a promising technology. The TPC is a monolithic detector which can distinguish between single site and multi site events. The xenon can be repurified over the course of the experiment to ensure detector performance. The energy resolution, while not as a good a solid state detectors, is much better than liquid scintillator-based detectors. Lastly, scaling up the detector mass is much easier for a liquid xenon TPC than for solid state detectors, which are much more expensive per unit mass.

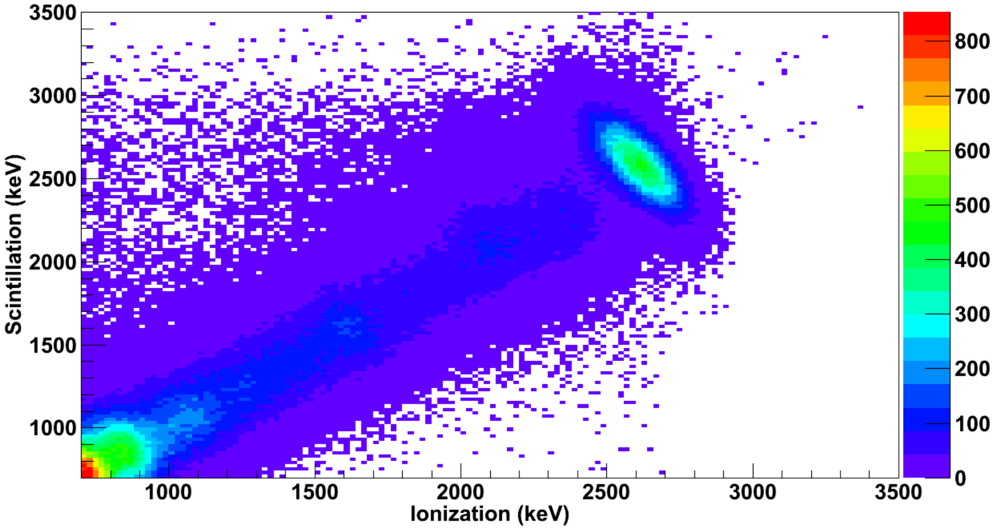

Figure 3.4: A 2D histogram of scintillation versus ionization for a ^{228}Th source calibration of EXO-200.

3.2 Backgrounds and Construction Considerations

As mentioned in section 2.4, any search for $0\nu\beta\beta$ is sensitive to radioactivity from construction materials. The EXO collaboration has had an extensive program of materials counting and qualification [43]. There are also backgrounds from cosmic rays and gases dissolved in the xenon. Comic ray backgrounds correspond not only to muons traversing the detector, but also muon induced isotopes and neutrons. For example, neutron capture on ^{136}Xe produces ^{137}Xe, which β-decays with a 3.8 minute half-life and 4.2 MeV end point. There are two ways to address such backgrounds: going deep undergrounds to gain shielding and vetoing muon induced events by coincidence with scintillator panels outside the detector. EXO-200 is located 2150 feet (655m) underground at the Waste Isolation Pilot Plant (WIPP) near Carlsbad, NM. The experiment is housed inside of a clean room which is surrounded by scintillation panels which identify muons that travel near the detector. This system vetoes $(95.5\pm.6)\%$ of muons traversing the TPC.

There are also backgrounds from ^{222}Rn dissolved in the xenon and in the air outside the cryostat. ^{222}Rn has a decay chain which includes ^{214}Bi, which has a gamma line at 2448 keV, within 10 keV of $Q_{\beta\beta}$ of ^{136}Xe. Internal system components and/or plumbing can emanate ^{222}Rn, and as a result these emanation rates are measured to qualify such parts for use. As will be discussed in section 6.2.3, the radon content of the xenon can be measured using the data.

^{222}Rn is also common in the air between the cryostat and the lead shielding, posing a potential background, as some of these gammas will penetrate the lead shielding and the HFE. The cryostat is surrounded by an enclosure which will be used to flush this area with low-radon air.

3.3 The Time Projection Chamber

Figure 3.5 shows a schematic of the detector, which is cylindrical and split into two TPCs (referred to as TPC1 and TPC2) by a high voltage cathode in the center. This minimizes the drift length of charge deposits in the detector, minimizing the effects of electronegative impurities. The Z coordinate is defined to be positive in TPC1, negative in TPC2, and zero at the cathode. The active xenon is surrounded by a low activity (<.06 mBq/kg) teflon reflector to improve light collection efficiency for VUV scintillation. Outside the teflon is a set of field shaping rings which create a uniform electric field of 376 V/cm throughout the fiducial volume of the TPC. The field shaping rings are held in place by combs made from low activity (<0.16 mBq/kg) acrylic [45].

The TPC is housed inside a copper vessel which is mounted onto the cryostat door using six legs. These legs also act as conduits for cables and xenon plumbing. Although copper is more radio-clean than most other metals, the mass of the copper vessel needs to be minimized. The vessel was made as thin as possible, at a thickness of only 1.37mm. As a result, the pressure differential between the inside and outside of the vessel must be maintained at lower than 250 Torr. Minimizing the mass of copper also means the vessel has no flanges, it is welded shut. All welds used in detector construction are of the electron beam type except for the field weld closing the end caps, which are TIG.

There are separate u and v wire planes on each end of the detector. One of two

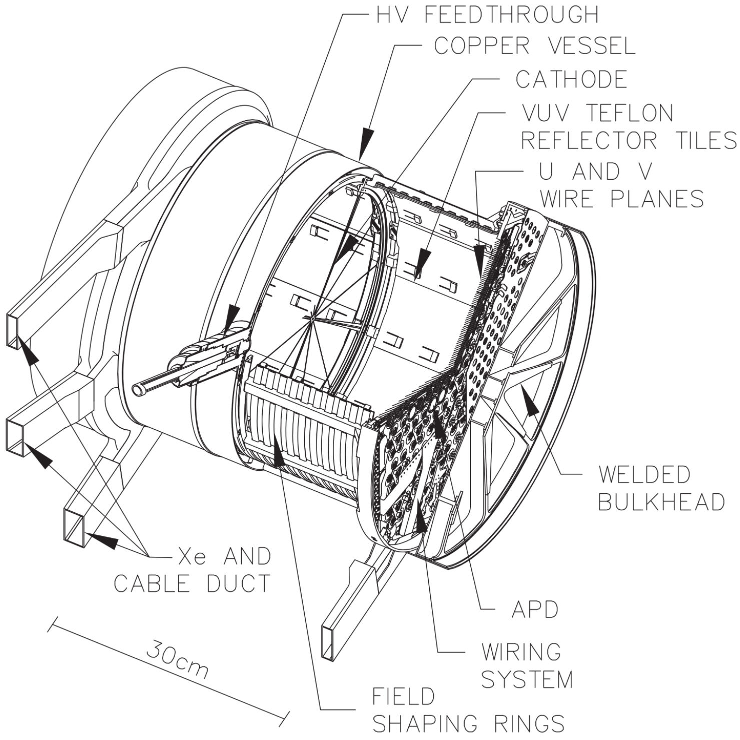

Figure 3.5: A schematic of the EXO-200 detector.

sets of u and v planes is shown in figure 3.6 and a single wire triplet is labeled (3). The wire triplets are made from photo-etched phosphor bronze with tabs at the end to provide tension. Behind the u-wire plane are 234 windowless avalanche photodiodes (APDs). The APDs are read out in groups of seven.

3.3.1 Detector Regions

Some areas close to the teflon reflector have lower charge collection efficiency. Figure 3.7 compares the radial distribution of events in ^{228}Th source calibration data and simulation. At the outer edge of the detector there are fewer events in data than the simulation predicts. There are two issues leading to this problem. The first is that when one gets close to the field shaping rings, the field lines become no longer parallel

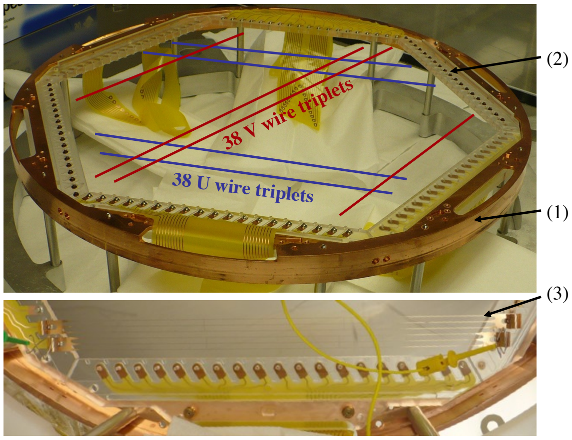

Figure 3.6: The u and v wires are mounted as shown. There is a copper ring (1), to which the six acrylic parts (2) are attached in a hexagonal configuration. Each of the 38 v-wire triplets and 38 u-wires triplets (3) attach to the acrylic parts. The red lines show the configuration of the v-wires and the blue lines show the configuration of the u-wires.

and some of the field lines will actually end on the field shaping rings rather than the u-wires. The full extent of this particular issue is not yet fully understood and work is ongoing to improve the electrostatic model of the detector to be accurate in three dimensions. The second is related to how the u and v wires are mounted. Figure 3.6 shows the acrylic parts to which the u and v wires are attached, which form a hexagon. The wire triplets stretch across the detector as shown by the red and blue lines. The teflon reflector is a cylinder, and parts of the acrylic mounts overlaps with the region inside the teflon, as shown in figure 3.8 in the area circled in red. In this image, one is viewing the charge collection plane from the cathode, so the v-wires are in front and the u-wires are mounted in back.

In the area indicated in the figure, there are tabs from several v-wires triplets inside the teflon radius. As the tabs are connected to the acrylic, there are no u-wires behind these tabs, this creates a region from which charge collection is more difficult. This is referred to as the “dead layer” (although this nomenclature can be misleading as the charge collection efficiency is lower, but not zero).

Figure 3.9 shows the regions of the detector. Inside the copper vessel but outside

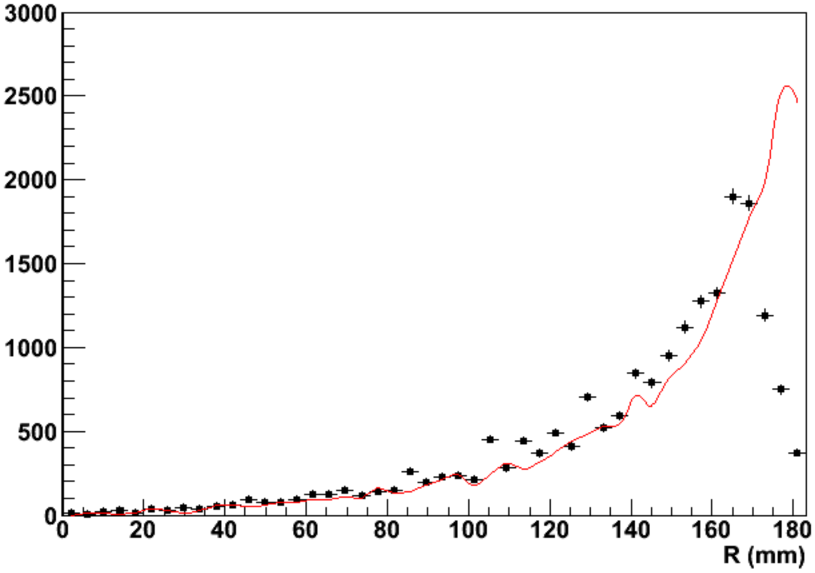

Figure 3.7: Number of charge deposits versus radius for ^{228}Th source calibrate data (black points) and simulation (red curve). The red curve is normalized to have the same integral as the black points. Outside a radius of 163 mm the number of charge deposits in data drops off, while the number of events in simulation continues to increase. This indicates not all the charge deposits outside 163 mm radius are detected.

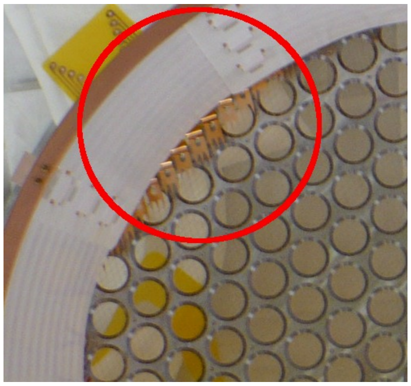

Figure 3.8: Close up image of a region of the detector near the teflon wall. The area circled in red shows the v-wire tabs inside the teflon radius, meaning there is a region where there are no u-wires.

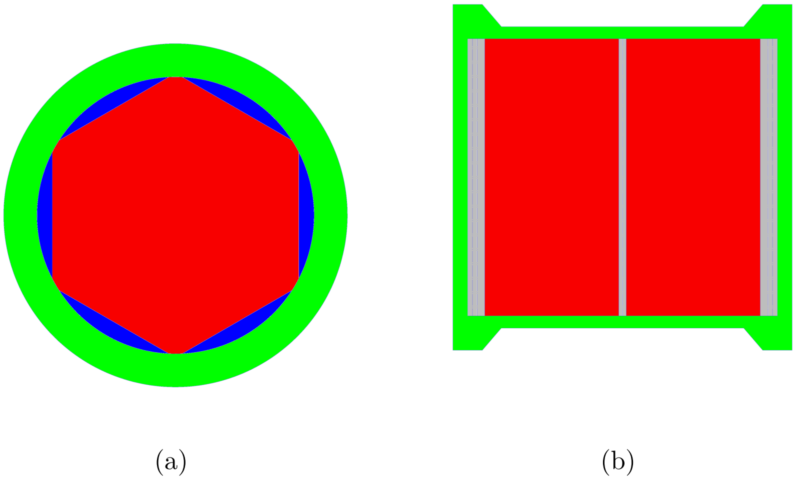

Figure 3.9: (a) X-Y and (b) X-Z views of the detector highlighting different regions. Shown are the inactive xenon (green), dead layer (or dead xenon) (blue), fiducial xenon (red), and non-fiducial active xenon (grey).

the teflon reflector and/or behind the APD planes is the inactive xenon (in green). Neither charge nor light from charge deposits in this region are detected. Within the radius of the teflon and between the APD planes is the active xenon, which is further broken up into three subsections. The fiducial xenon is the volume of xenon where the detector response is understood and the charge collection if fully efficient (in red). Previous work has demonstrated that the appropriate shape for this region is a hexagon as this mimics the geometry of the u and v wires. Close to the walls of the teflon is the dead xenon (in blue), where the scintillation light is detected but the charge collection efficiency is unknown. The regions close to the anode planes or to the cathode are not used for the analysis. Close to the anodes the charge signal is heavily affected by the shielding grid inefficiency of the v-wire plane (see section 7.3.1) and near the cathode there is evidence of light collection inefficiency. This is referred to as the non-fiducial active xenon (in grey).

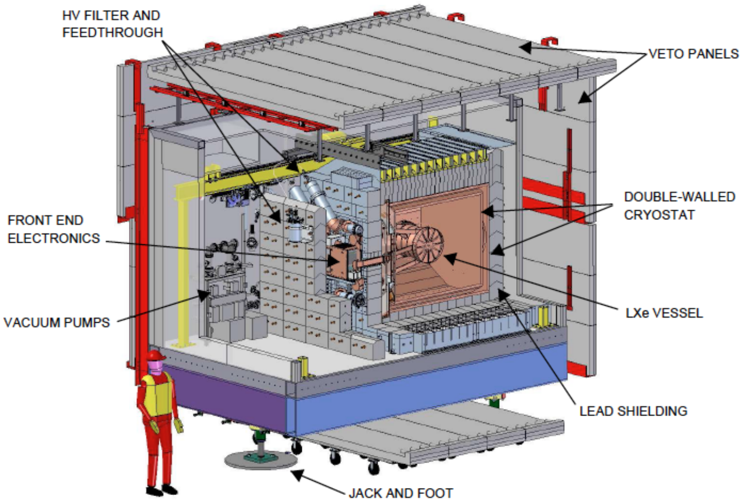

Figure 3.10: A schematic of the copper vessel and the surrounding infrastructure.

3.4 Detector Infrastructure

Figure 3.10 shows copper vessel and the surrounding infrastructure. The vessel is mounted inside a two walled cryostat, both made from low activity (>1.9 mBq/kg) copper, which is surrounded by at least 25cm of low activity (<0.8 mBq/kg) lead on all sides [45]. The readout electronics for the experiment are located in between these two sets of lead walls, just outside the cryostat. All this resides inside a clean room surrounded by the muon veto panels. Because the experiment takes place in a salt mine, the clean rooms need to be able to adapt to shifts in the ground underneath, so they are mounted on adjustable legs.

3.4.1 The Cryostat

The inner vessel of the cryostat contains four tons of HFE-7000, a fluorocarbon heat transfer fluid manufactured by the DuPont corporation. The HFE-7000 also provides shielding, covering at least 50cm on every side of the vessel. HFE-700 is a good material for shielding because it is particularly radio-clean. The entire four tons of HFE-7000 has a total activity from ^{40}K, ^{238}U, and ^{232}Th of less than 20 mBq [45].

The outer vessel of the cryostat provides vacuum insulation.

3.4.2 The Calibration System

Figure 3.11 shows the copper tubing through which calibration sources are deployed. The calibration sources are .7mm in diameter and attached to the end of a flexible steel cable which can be fed through the tubing to reach the desired location near the detector. There are a total of five calibration locations in use, three near the cathode and one at each anode. Of the three cathode positions, one is at positive X (the far side from the cryostat door), one is at positive Y (closer to the ceiling), and one is at negative Y (closer to the floor). Most calibration runs are either at the positive X side of the cathode or at the negative Z anode as the other locations require deploying the cable deeper into the tubing, which is higher risk for the system. Figure 3.12 shows an X-Y distribution for events in the detector from a calibration run taken at positive Y. As expected, the events tend to be closer to the location of the source.

In total three different isotopes are in use for detector calibrations: ^{137}Cs with a line at 662 keV, ^{60}Co with lines at 1173 keV and 1333 keV, and ^{228}Th which has a decay chain including ^{208}Tl with a line at 2615 keV. In total there are five sources, with two sources each of ^{60}Co and ^{228}Th with different activities, referred to as the strong sources (higher activity) and weak sources (lower activity). Table 3.1 summarizes the activities of these sources at relevant times. The weak sources cause little to no pile up in the data, and are useful for studies of very specific aspects of detector performance. The strong sources are for high statistics measurements of calibration gamma lines. There is a calibration run almost every day for 1-2 hours to monitor detector performance and measure the electron lifetime. Several long calibration runs have been taken to study other position dependent corrections to the detector response.

3.4.3 Xenon Purification

The xenon is constantly recirculated through hot zirconium getters to maintain high purity. To ensure that recirculation does not introduce impurities or radioactive

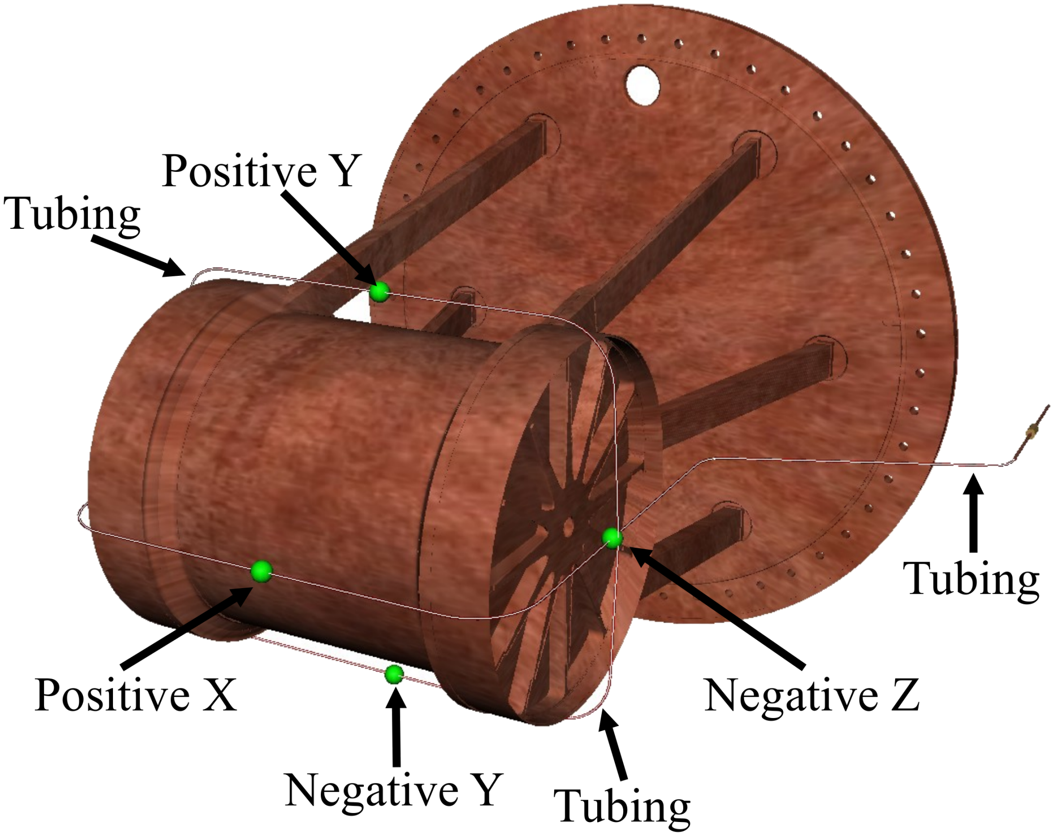

Figure 3.11: The calibration tubing surrounds the detector providing a way to deploy radioactive sources. The green dots represent typical source locations for calibration runs.

Table 3.1: Calibration source activities at time of purchase and at the beginning of Run 2 (10/5/11).

Source	Half-life	Activity on 9/1/09	Activity on 10/5/11
^{137}Cs	30.1 years	2820 Bq	2687 Bq
Weak ^{60}Co	5.27 years	529.5 Bq	402 Bq
Strong ^{60}Co	5.27 years	7060 Bq	5362 Bq
Weak ^{228}Th	1.91 years	1417 Bq	664 Bq
Strong ^{228}Th	1.91 years	34040 Bq	15944 Bq

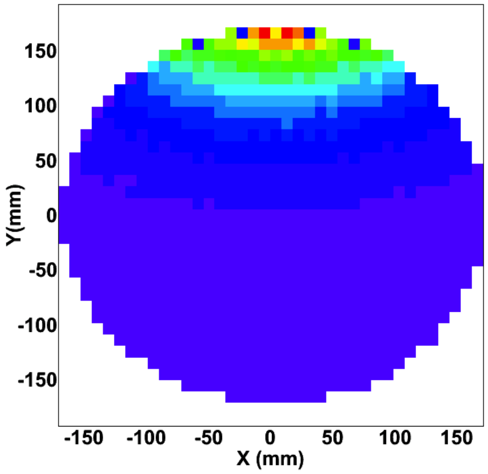

Figure 3.12: X-Y distribution of events for a ^{228}Th source deployed to positive Y.

species, an ultra clean pump using a magnetically driven teflon piston was developed [49]. Figure 3.13 shows the purity and pump flow rate versus time over the course of the data taking. The dashed red lines and arrows show the bounds of the time periods where the purity is stable enough for use in the analysis.

There is a period at the beginning where the electron lifetimes are relatively short, between 200 μs and 300 μs. This is referred to as Run 1 of the detector. After Run 1, the recirculation rate was raised, increasing the electron lifetime $\sim$ 3 ms. The maximum drift time for any event in the detector is 116 μs, so at an electron lifetime of 300 μs this corresponds to a loss of 32% of the charge signal for charge deposits at the cathode. At 3 ms electron lifetime, only 4% of the charge signal is lost.

After the xenon recirculation rate is raised above 10 slpm, the only periods of poor purity come as a result of the recirculation having to stop temporarily, usually due to power outages. In these cases the purity drops very fast but recovers over the course of a few days when recirculation is restored. The recirculation is critical to obtaining high purity and in the current setup and makes long electron lifetimes possible.

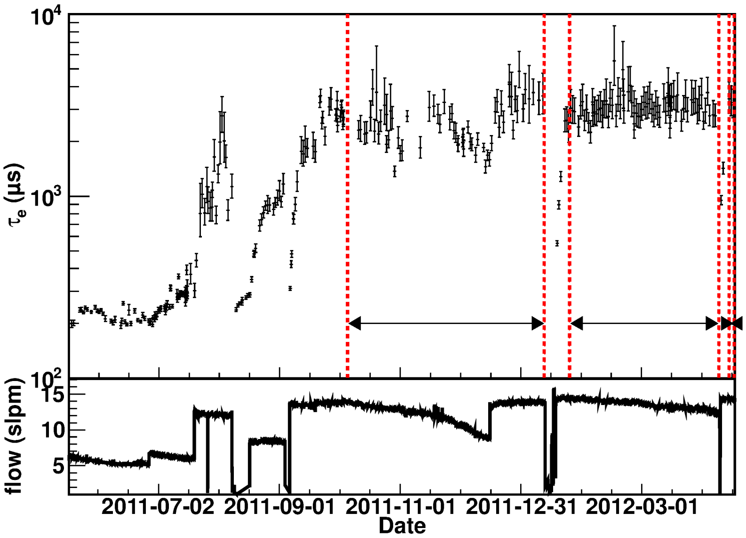

Figure 3.13: Electron lifetime and xenon recirculation rate over time.

3.5 Detectors Run 1 and 2

The first run of EXO-200 collected data from May 21, 2011 to July 9, 2011 and made the first measurement of $2\nu\beta\beta$ in ^{136}Xe [46]. This run collected 752.66 of low background physics data. The scintillation signal was only used for determining the Z coordinate of events and measuring the radon content of the detector, the energy of events was determined using only the measured ionization. Most of the analysis discussed here will pertain to a separate data set, referred to as Run 2. As already mentioned (see section 3.4.3), the xenon purity is improved in Run 2. An additional lead wall is added on one side of the cryostat to shield from external backgrounds. Also, the differentiation times of the u-wire signals were lengthened to improve fits to u-wire signals (see section 4.1.2).

Chapter 4

Electronics and Data Acquisition

Figure 4.1 shows the structure of the electronics. Charge and light are collected at the wires and APDs respectively and fed through kapton insulated flex cables to the front end electronics cards (FECs). Signals are amplified by charge preamps before being shaped and going to analog-to-digital converters (ADCs). From the ADCs the signals are passed into the data acquisition system (DAQ) which consists of a trigger event module (TEM) and a computer with a raid array.

4.1 FECs

4.1.1 Wire Gains

U-Wire Gains

Individual u-wire channels are calibrated using a series of three different data sets: precision pulser calibration, internal charge injection, and external charge injection. Both internal and external charge injection calibrations use an on-board calibrator which injects charge directly to the preamp. Several pulses of fixed amounts of charge are injected into each channel. An example histogram of pulse heights for a precision pulser run is shown in figure 4.2. The gain of the channel, G, is

$$G = \frac{injected\,charge}{\#\,ADC\,counts} \tag{4.1}$$

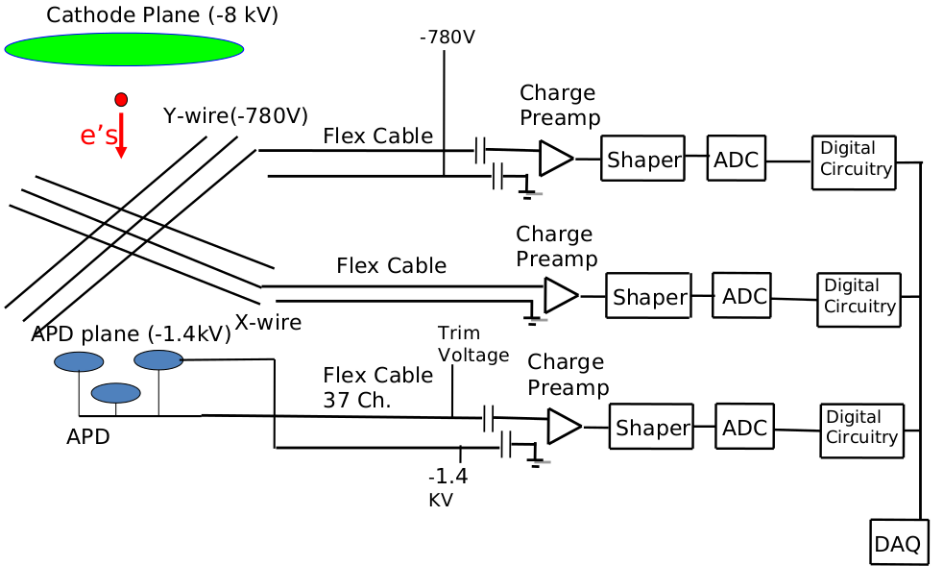

Figure 4.1: A schematic showing the structure of the EXO-200 electronics.

Precision pulser calibration injects charge through a capacitor whose value is measured to better than 1%, which is more precise than the charge injections calibrations. An internal charge injection run, where the channels are connected to their respective loads from the TPC, is compared to an external charge injection run, where the channels are not connected to their respective loads from the TPC, to calculate this correction, C

$$C = \frac{G_{internal}}{G_{external}} \tag{4.2}$$

where $G_{internal}$ is the gain measured from internal charge injection and $G_{external}$ is the gain measured from external charge injection. This correction is 0.5% or less for any given channel. U-wire gains measured using electronics calibrations are used for the Run 1 analysis. In the Run 2 analysis, a more accurate method for wire gain determination which utilizes source calibration data is applied (see section 7.3.1).

V-Wire Gains

The V position of a charge deposit is determined using a weighted mean of its v-wire signals (see section 5.5). The gains of the v channels range from 300 to 360

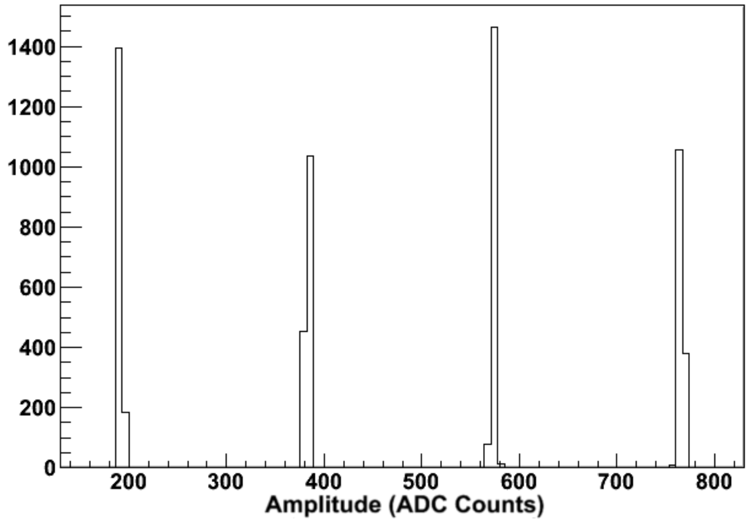

Figure 4.2: A histogram of pulse amplitudes for a precision pulser calibration of a u-wire channel.

electrons/ADC counts, as shown in figure 4.3, so a gain correction is made to the v-wire amplitudes to improve V position determination. The gains of the v-wires are measured using a series of pulses from an external calibrator, similar to the u-wire gain calibration run, but no corrections are made for the load of the TPC.

4.1.2 Signal Shaping

The signals are shaped using one integration time, τ_1, and two differentiation times, τ_2 and τ_3. The shaping of each channel is two integrators with time constant τ_1, two differentiators with time constant τ_2, and one differentiator with time constant τ_3. The integrators act as a filter for high frequency noise while the differentiators filter out low frequency noise, primarily from microphonics. The first four shaping stages for the u-wires are shown in figure 4.4. The u-wire shaping times were changed between Run 1 and Run 2 to reduce the information loss in the shaping.

The final shaping stage comes from the input to the FET. For the u-wires, this time constant is determined by a 1 pF capacitor and a 60 MΩ resistor. A 1 pF capacitor is only accurate to within ±0.5 pF, meaning this 60 μs differentiation time has a large uncertainty and will vary from channel to channel. The same problem may exist for

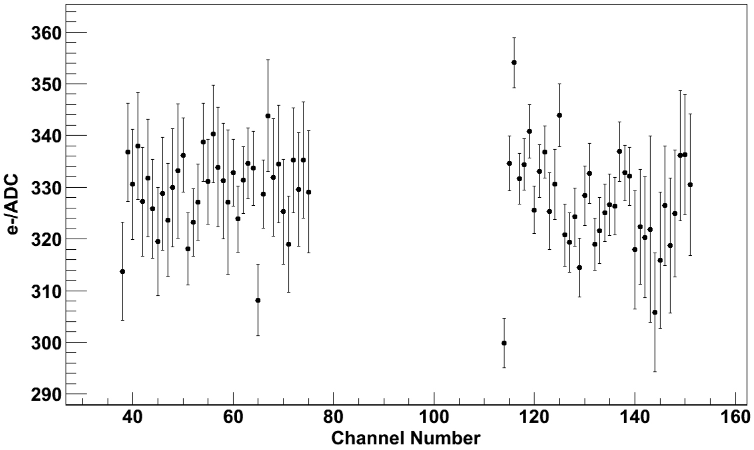

Figure 4.3: V-wire gains measured for each channel with errors.

Channel type	Stage type	Run 1 Time Constant (μs)	Run 2 Time Constant (μs)
APDs			
	Int	3	3
	Int	3	3
	Diff	10	10
	Diff	10	10
	Diff	300	300
U-Wires			
	Int	3	1.5
	Int	3	1.5
	Diff	10	40
	Diff	10	40
	Diff	60	Measured
V-Wires			
	Int	3	3
	Int	3	3
	Diff	10	10
	Diff	10	10
	Diff	60	60

Table 4.1: Shaping times relevant to the transfer functions of different channel types. In Run 2 the final differentiation stage for the u-wires is measured from fits to charge calibration pulses.

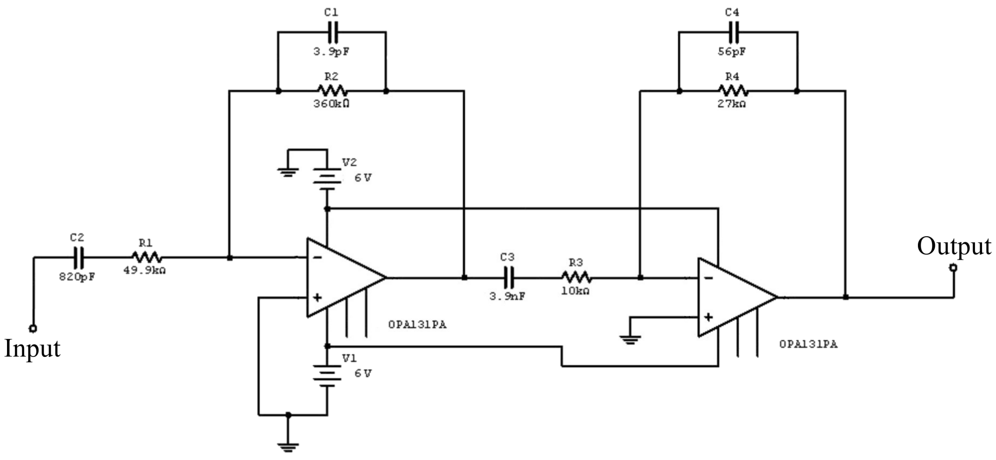

Figure 4.4: A schematic of the first four shaping stages for the u-wires.

the APDs and for the v-wires, but as these channels are more aggressively shaped (the differentiation times used are shorter and the integration times are longer) the effect of the final differentiation is smaller. Figure 4.5 shows the shaping of a prototypical detector pulse, in this example a u-wire pulse is shown. As shall be discussed in section 5.3, this model is used to fit detector signals to extract amplitudes and timing information.

An Analytical Pulse Shape Model

The pulse shaping of the electronics is modeled using a Laplace transform to calculate the effects of the shapers. For a function $f(t)$, its Laplace transform $F(s)$ is

$$F(s) = \int_0^\infty f(t)e^{-st}\,\mathrm{d}t \tag{4.3}$$

This is particularly convenient for representing shapers as the transforms are

$$F(s) = \frac{a}{s+a}(integrator) \tag{4.4}$$

$$F(s) = \frac{s}{s+b}(differentiator) \tag{4.5}$$

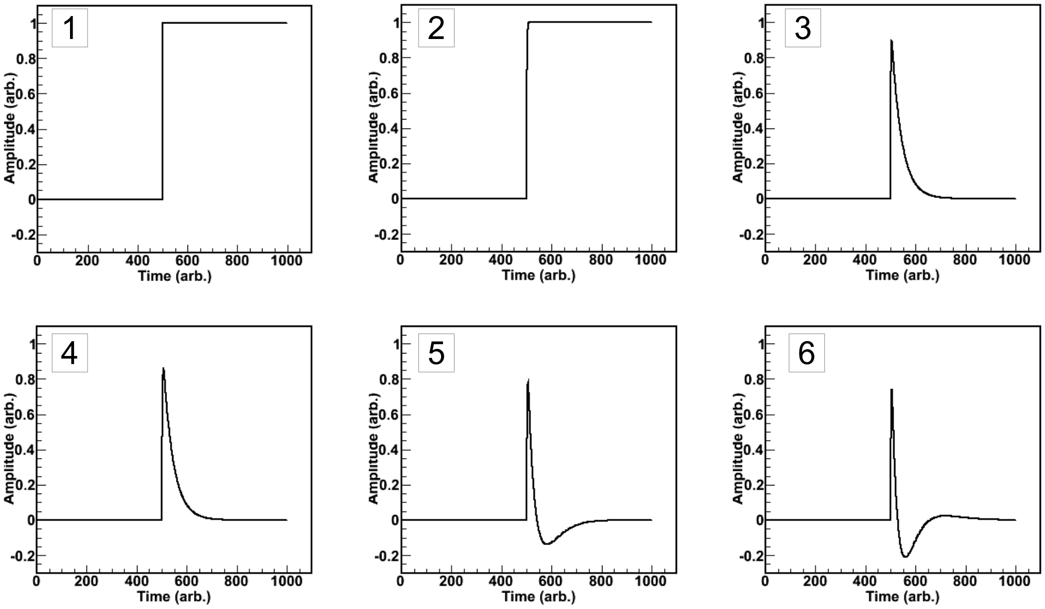

Figure 4.5: Pulse shaping of a u-wire signal. A raw pulse (1) is a step function. This pulse is shaped with a 1.5μs integrator (2), a 40μs differentiator (3), a second 1.5μs integrator (4), a second 40μs differentiator (5), and finally a 60μs differentiator to arrive at the final pulse shape (6).

where $a = \frac{1}{\tau_1}$ and $b = \frac{1}{\tau_2}$. In total there are two integrators which are assumed to have the same time constant $\frac{1}{a}$, and three differentiators, two of which are assumed to have the same time constant $\frac{1}{b}$ and a third assumed to have a different time constant $\frac{1}{c}$. For the nominal shaping times $a = \frac{1}{1.5\mu s}$, $b = \frac{1}{40\mu s}$, and $c = \frac{1}{60\mu s}$. Combining all five of these shapers and multiplying by the transform of a step function, $\frac{1}{s}$,

$$F(s) = \frac{a^2 s^2}{(s+b)^2 (s+a)^2 (s+c)} \tag{4.6}$$

The inverse transform is

$$f(t) = a^2 \left(-\left(\frac{b\left(b^2+b\,a-2\,a\,c\right)}{(b-a)^3\,(b-c)^2\,e^{b\,t}} \right) + \frac{a\left(a^2+b\,(a-2\,c)\right)}{(b-a)^3\,(a-c)^2\,e^{a\,t}} + \frac{c^2}{(b-c)^2\,(a-c)^2\,e^{c\,t}} - \frac{b^2\,t}{(b-a)^2\,(b-c)\,e^{b\,t}} - \frac{a^2\,t}{(b-a)^2\,(a-c)\,e^{a\,t}} \right) \tag{4.7}$$

where t is the time in microseconds, with t=0 at the onset of the pulse. This model can be fit to calibration pulses with the various shaping times allowed to float, hence measuring the true shaping times of the electronics.

Real charge collection pulses have a 2-3 μs rise time associated with the drift of charge between the v and u wires and this is missing from the analytical model. This does not affect the APD signals or charge injection pulses.

U-Wire Shaping Times

U-wire calibration pulses from the on-board calibrator are fit to the analytical with of the shaping times, τ_1 and τ_2, fixed at their nominal values and the nominal 60 μs differentiation time, τ_3, allowed to float. Figure 4.6 shows a charge injection pulse fit to two different models: τ_3 is fixed and one where it it allowed to float. When fixing this shaping time, the model cannot fit both the peak and the undershoot. Not using the correct transfer function for a channel can affect the measurement of the charge deposit energy, as this comes from a fit of the waveform to a pulse shape model (see section 5.3). These shaping times have been measured for all u-wire channels and are used for the pulse signal fits.

The other shaping times on the u-wire channels also have uncertainties of 5% and the two differentiation times are highly correlated in the fit. The fit for the nominal

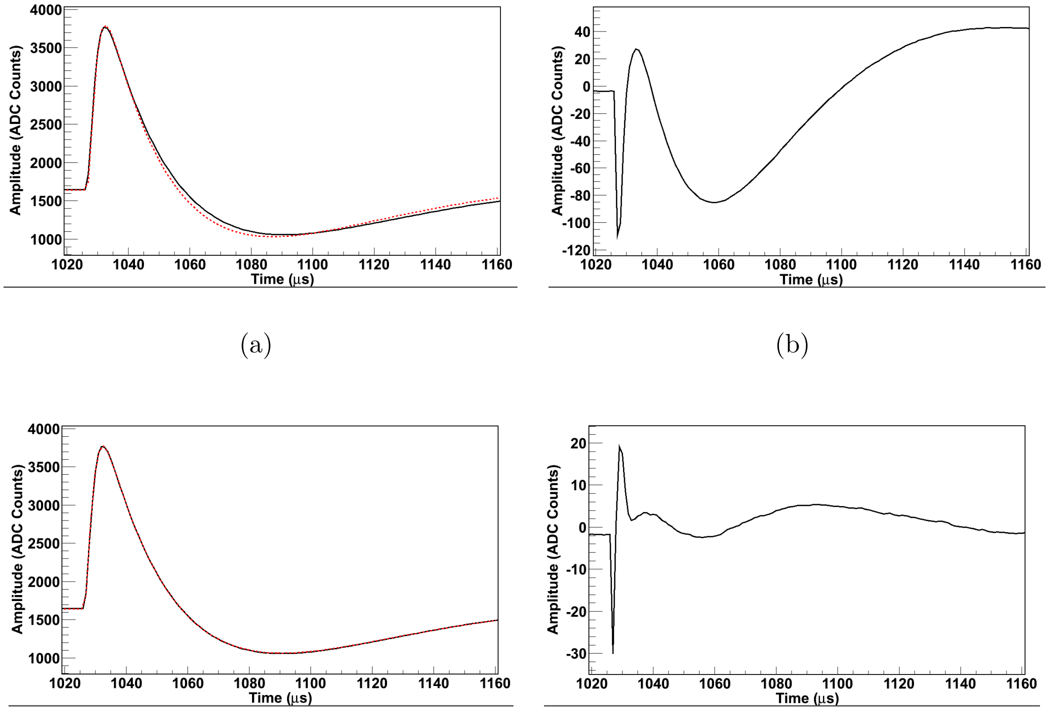

(a) (b)

(c) (d)

Figure 4.6: (a) Fit of a charge injection pulse (black) to a pulse shape model where the 60μs differentiation time is fixed (red) with the difference between the model and signal shown in (b). (c) Fit of a charge injection pulse (black) to a pulse shape model where the 60μs differentiation time is allowed to float (red) with the difference between the model and signal shown in (d). The true value of τ_3 for this channel is measured to be 77.4μs.

60 μs shaping time is really determining an effective value which is affected by several parameters and is only the best approximation of the pulse transfer function. Figure 4.7 shows the measured differentiation times for the various channels. Channels on the same FEC tend to have similar components and stray capacitances, leading to similar τ_3.

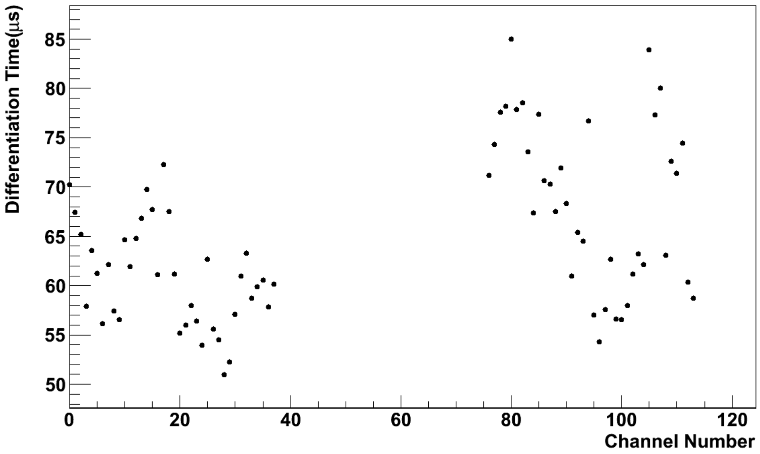

Figure 4.7: U-wire differentiation times for each u-wire channel.

4.2 DAQ

The TEM can trigger on individual channels or the sum of channels. Typically this means the TEM triggers on a single u-wire channel, or on the sum of the signals from all the APDs on one side of the detector. The DAQ can handle a trigger rate of up to 50Hz in one of three modes:

- The physics trigger is used for low background data taking and has a set thresholds for the individual wire channels, individual APDs channels, and sum of the APD signals.
- The prescale trigger level n records every $(n+1)^{th}$ trigger.
- The solicited trigger records events at a rate specified by the user.

The prescale and solicited triggers are necessary for source calibration runs as their event rates are too high for the DAQ. As will be explained in more detail in section 7.5, the solicited trigger source calibration runs are necessary for validation of the detector Monte Carlo. Solicited trigger runs are less efficient at recording events than the prescale trigger as some recorded triggers contain no charge deposits in the detector. The solicited trigger is also used for detector livetime studies, as all physics data runs use a .1 Hz solicited to determine the DAQ livetime.

The ADCs sample at 1 MHz and each waveform is 2048 μs long. The trigger time is set to 1024 μs so that each waveform contains 1 ms before, and 1 ms after, the trigger. A total of 152 wire channels (76 u-wires and 76 v-wires) and 71 APD channels are stored for each event. There are three gangs of APD which had to be disconnected due to unacceptable contributions to the noise.

Chapter 5

Event Reconstruction

5.1 Overview

Event reconstruction progresses in four stages: signal finding, pulse fitting and parameter extraction, signal clustering, and position determination. A pattern recognition algorithm goes through every waveform to find the signals and these signals are fit to a pulse shape model to extract amplitude and time information. The extracted parameters are used to determine the properties of individual charge deposits, referred to as charge clusters, and the parameters of the charge clusters are applied to determine three dimensional position information.

5.2 Signal Finding

5.2.1 Matched Filter

One of the challenges of event reconstruction is to find signals with as low a threshold as possible. Although $Q_{\beta\beta}$ of ^{136}Xe is 2458 keV [42], many of the high energy events in the detector contain several signals which must be combined. If small amplitude signals are missed, the energy of the event is underestimated and the event may be identified as signal site when the event contains multiple charge deposits, leading to an increase in the background.

Consider an inner product of the signal waveform with a filter. The waveform, w(t) is a combination of the signal, s(t), and an additive noise, n(t): w(t) = s(t) + n(t). Assume Gaussian white noise with a spectral density of $\frac{N_0}{2}$. The signal-to-noise ratio for a given filter, h(t), is

$$SNR = \frac{\mid s(t_d) \mid^2}{\mid n(t) \mid^2} = \frac{\mid \int_{-\infty}^{\infty} H(f) S(f) e^{2\pi i f t_d} \, \mathrm{d}f \mid^2}{\frac{N_0}{2} \int_{-\infty}^{\infty} \mid H(f) \mid^2 \, \mathrm{d}f} \tag{5.1}$$

where H(f) and S(f) are the Fourier transforms of the signal and the filter respectively, and t_d is the time of a putative pulse. Applying the Cauchy-Schwarz inequality

$$SNR = \frac{\mid \int_{-\infty}^{\infty} H(f) S(f) e^{2\pi i f t_d} \, \mathrm{d}f \mid^2}{\frac{N_0}{2} \int_{-\infty}^{\infty} \mid H(f) \mid^2 \, \mathrm{d}f} \leq \frac{\int_{-\infty}^{\infty} \mid H(f) \mid^2 \, \mathrm{d}f \int_{-\infty}^{\infty} \mid S(f) \mid^2 \, \mathrm{d}f}{\frac{N_0}{2} \int_{-\infty}^{\infty} \mid H(f) \mid^2 \, \mathrm{d}f} = \frac{\int_{-\infty}^{\infty} \mid S(f) \mid^2 \, \mathrm{d}f}{\frac{N_0}{2}} \tag{5.2}$$

with maximal signal-to-noise ratio in the case of equality. This occurs when H(f) and S(f) are linearly dependent, i.e. for some constant k

$$H(f) = k S^*(f) \tag{5.3}$$

The optimal filter is the signal shape itself. By scanning over time in the waveform, the algorithm can also find the time offset of the pulse. Figure 5.1 shows an example of a u-wire pulse and its matched filter output. Both peak at the same time and the matched filter is a clean signal with a higher signal-to-noise ratio. The matched filter algorithm is applied to find signals on the u-wires, v-wires, and APDs.

However, the assumption of Gaussian white noise is not quite correct. The noise in EXO-200 is colored because it goes through shapers, and there may also be components of noise that have specific frequency spectra. To obtain the optimal signal filter, the signal is first passed through a whitening filter. Let R(f) be the frequency spectrum of the waveform including signal and noise, and the noise spectrum be N(f). The filtered spectrum, R'(f) is

$$R'(f) = \frac{R(f)}{|N(f)|} \tag{5.4}$$

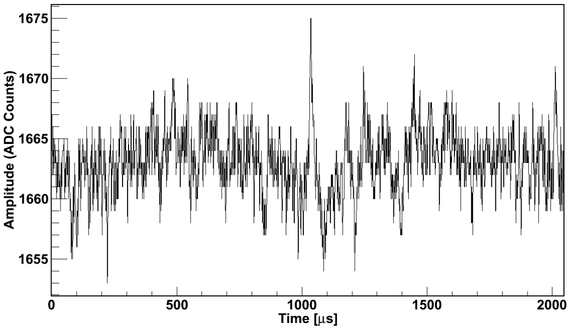

(a)

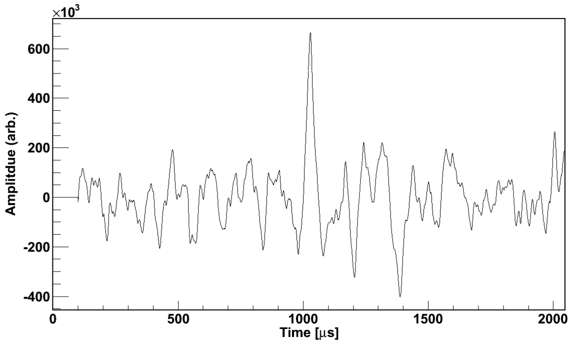

(b)

Figure 5.1: (a) A u-wire pulse from a detector event. (b) the matched filter output for this waveform. The $\frac{peak}{RMS}$ is 3.6 for the raw waveform and 8.2 for the matched filter output.

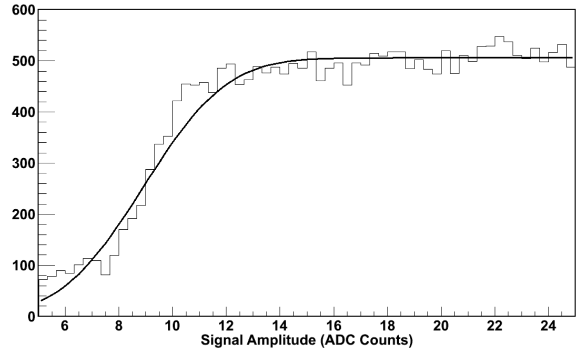

Figure 5.2: Fit of the low energy tail of reconstructed charge clusters to an error function to determine charge reconstruction threshold.

This whitening filter must also be applied to the H(f),

$$H'(f) = \frac{H(f)}{|N(f)|} \tag{5.5}$$

so that the optimal filter, H”(f), is

$$H''(f) = \frac{S^*(f)}{|N(f)|^2} \tag{5.6}$$

The whitening filter is applied to APD signal finding, but has been found to degrade the signal-to-noise ratio of u-wire signal finding and is not applied for the u-wires.

5.2.2 Charge Reconstruction Threshold

The charge detection threshold on the u-wires using the matched filter can be determined using source calibration data. Multi site events with greater than 500 keV of energy were used and the distribution of the lowest amplitude, shown in figure 5.2, is fit to an error function to model the charge pulse detection threshold. Above 14 ADC counts (∼80 keV), 98% of charge clusters are reconstructed.

5.2.3 Waveform Unshaping

Figure 5.3 shows a trace from a u-wire channel which contains two pulses, but the matched filter is only able to find one. The second plot shows the matched filter output for this trace with a red line indicating the threshold. There is a bump in the matched filter output corresponding to the earlier of the two pulses, but it is not discernible from the main peak and would be missed if reconstruction depended solely on the matched filter to find signals. Instead reconstruction uses the matched filter to determine if there was a signal on a given channel, and a sub-trace of the waveform is unshaped (as shown in the third panel of figure 5.3) then reshaped with a 2μs triangular filter to create distinct peaks.

Figure 5.4 shows several tests of the reshaping algorithm on simulation versus reconstructed energy. In thorium source calibration simulations 10% of single site events are moved to multi site in most energy bins, but this fractions gets as high as 50% near the full energy peak. In contrast, less than 1% of $2\nu\beta\beta$ events in any energy window are moved from single site to multi site.

Figure 5.5 shows this test repeated on thorium source calibration data. The fraction of events moved from the single site to the multi site is lower than in simulation, ranging from 5% to 25% depending on the energy window. This is likely related to differences in the noise and knowledge of the shaping times and baselines. Although the noise in simulation can be chosen to have the right amplitude, the frequency power spectrum may not match that of data. Also, the unshaping process is sensitive to the shaping times of the pulses and the waveform baselines, which are known exactly in the simulation, but only approximately in the data.

From source calibration data there is also evidence that the energy of some events is underestimated when not using the reshaping algorithm. Figure 5.6 shows the ratio of the reconstructed energy when using the reshaping algorithm to the reconstructed energy when not using the reshaping for events moved from single site to multi site by the reshaping. The reconstructed energy when using the reshaping is on average 3% higher. As shall be discussion in section 5.3, the energy of the event is measured using fits of a pulse model to the waveforms, and if the fit model contains fewer pulses than are in the data, an inaccurate energy measurement may result. The fraction of

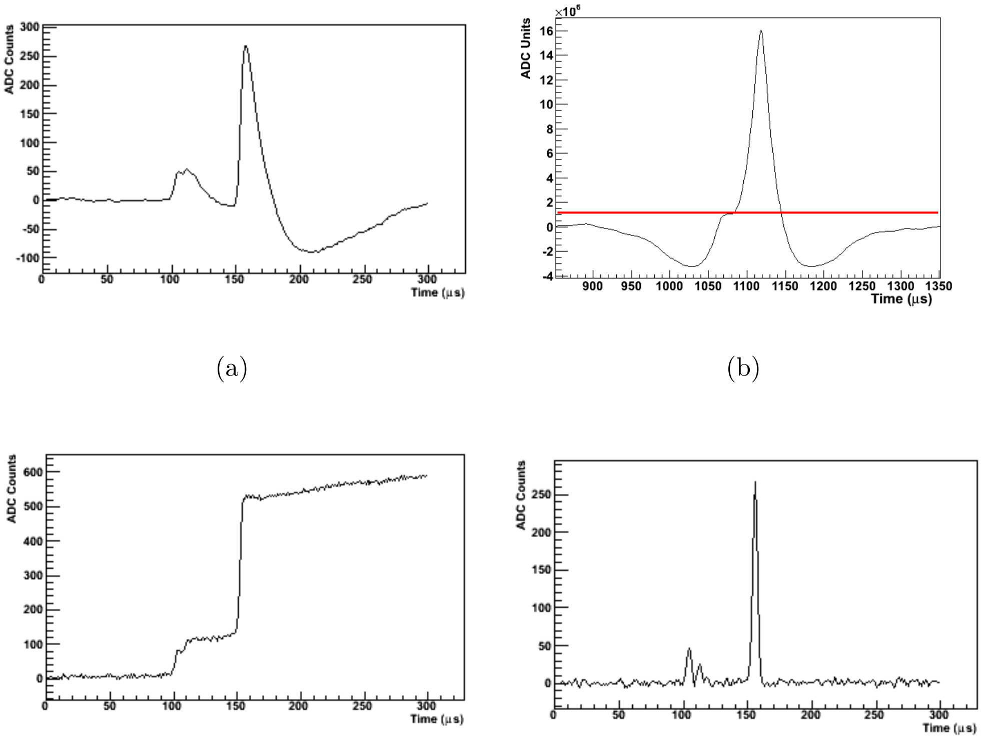

(a) (b)

(c) (d)

Figure 5.3: (a) a raw waveform from a u-wire channel with two pulses (b) the matched filter output for this waveform (c) an unshaped sub-trace of this waveform (d) the sub-trace after application of a 2μs triangular filter.

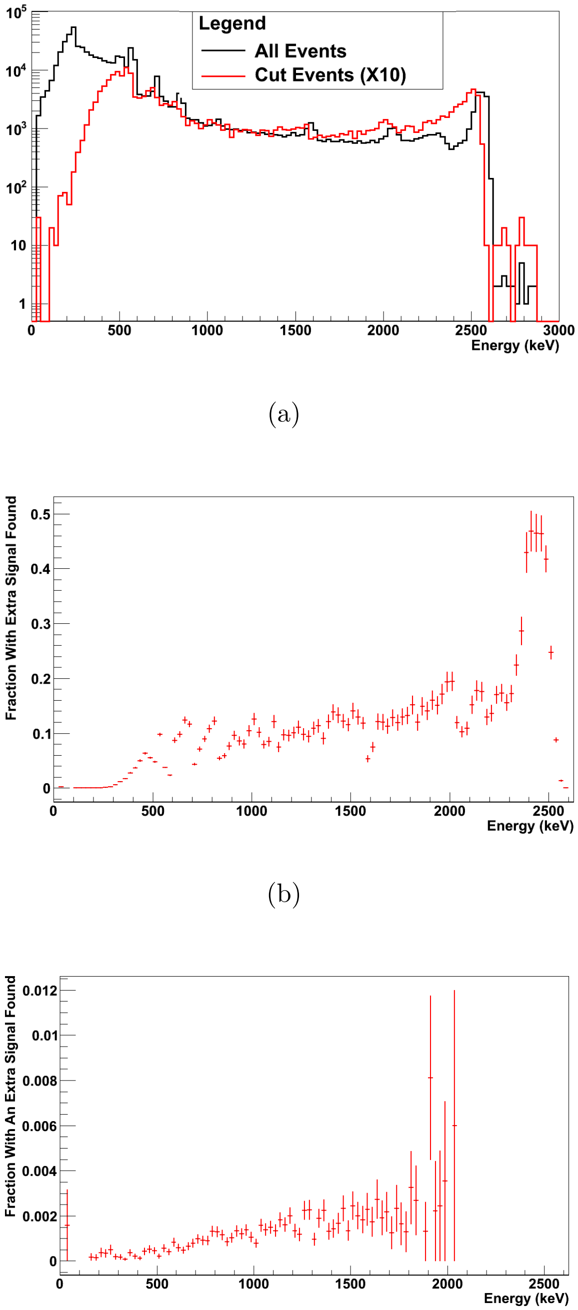

(a)

(b)

(c)

Figure 5.4: (a) Single site spectrum for thorium source calibration simulation (black) without the reshaping algorithm and the energy spectrum of events moved from single site to multiple site by the algorithm scaled by a factor of ten (red). (b) Fraction of events moved from single site to multi site versus energy for thorium source calibration simulation. (c) Fraction of events moved from single site to multi site in $2\nu\beta\beta$ decay simulation versus energy. In all cases the energy is the energy is taken to be the energy measured without the reshaping algorithm.

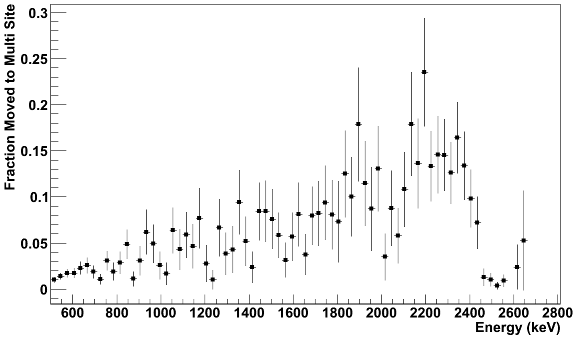

Figure 5.5: The fraction of events moved from single site to multi site by the triangular reshaping algorithm in thorium calibration data versus the energy of the event measured without the reshaping algorithm.

events moved from single site to multi site is higher just below the full energy peak because there are true multi scatter full energy deposition events that have their energy underestimated due to this effect. A 2615 keV gamma from ^{208}Tl which has its energy underestimated by 3% will be reconstructed with 2537 keV, within 80 keV of $Q_{\beta\beta}$ of ^{136}Xe. This is an important point when considering the effects of the reshaping technique on the low background physics data, which will be discussed in section 7.5.6.

5.3 Pulse Fitting and Parameter Extraction

To derive a numerical model of the pulse shapes for physics events, a step function is shaped with the appropriate integration and differentiation times of the electronics (this is exactly what is done in figure 4.5). The step function includes the rise time due to charge drift between the v and u wires (see section 4.1.2) as modeled by MAXWELL [50]. U-wire and APD channels where one or more signals are found are fit to this model for parameter extraction. The v-wires channels are not fit, the magnitude of the matched filter amplitude is taken as the signal amplitude. The signal finding output provides initial guesses for the number of pulses, and their times and amplitudes which are allowed to float in the fit. Both individual APD channels and the sum of the APD waveforms for each plane are fit. For the u-wires, the measured

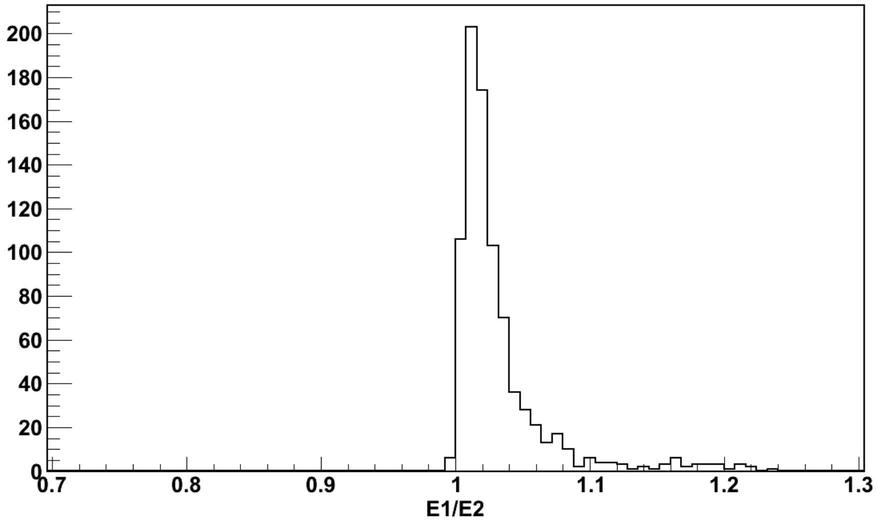

Figure 5.6: The ratio of the reconstructed energy when using the reshaping algorithm, E_1, to the reconstructed energy when not using the reshaping algorithm, E_2 from ^{228}Th source calibration data for event moved from single site to multi site.

shaping times for each channel are used (see section 4.1.2).

In some cases, a waveform will be fit to multiple signals and the fit will favor one true signal, and this prompts the need for a pulse culling procedure. An individual pulse is culled if its fitted parameters meet any of the following criteria:

- $\frac{Amplitude}{AmplitudeError} < 5$ (6) for APDs (u-wires)
- The pulse amplitude is less than 5 ADC counts
- The pulse has an error of >15 ADC counts and there is a second pulse within 10μs. This criteria only applies to u-wires.

The culling is an iterative process: pulses are removed from the fit and the waveform is refit repeatedly until no remaining pulses meet any of the criteria for culling.

5.4 Signal Clustering

For each u-wire signal, the event reconstruction looks for the corresponding v-wire signals using a matched filter at a fixed time 2 μs earlier. For a given charge deposit, there can be several v-wire signals as the induction signal can reach over a distance of a few wires, even if it is only collected on one wire. If no v-wires signals are found, the charge cluster will have no associated v-signals.

After finding the appropriate v-wire signals, multiple u-wire signals may need to be combined into a single cluster. If a charge deposit occurs between two u-wires, it is possible for both u-wires to collect some of the charge and identifying this as a multi site event would be incorrect. Whenever there are signals on adjacent u channels, or on the same u channel, within 3.5 μs of each other, these two signals are combined into a single charge cluster. Events containing more than two wire channels are classified as multi site.

Each charge cluster is associated with the closest scintillation signal in time, requiring that the scintillation signal occur first and that the time difference between them is not longer than the longest possible drift time. If no scintillation signals meets these criteria, no scintillation signal is associated with the charge cluster.

5.5 Position Determination

Based on which u-wires had signals for a given charge cluster, a U-position for the event is determined. If there is more than one u-wire in a given charge cluster, the U-position is the energy weighted mean of the U-positions of all channels. In the V-direction the position is also calculated as the energy weighted mean of the v-wire signals. As discussed in section 3.1.1, the U and V coordinates are not orthogonal, there is a 60° angle between them, and these coordinates are translated into orthogonal coordinates X and Y. The X-direction points away from the cryostat door and the Y-direction points toward the ceiling of the clean room. If a charge cluster does not have any v-wire signals, no X-Y coordinates are assigned.

As explained in section 3.1.1, in a time projection chamber the Z coordinate is determined using the time difference between the scintillation and charge collection signals. If a charge cluster has no associated scintillation signal, no Z coordinate is assigned. Determining Z from the drift time requires a measurement of the drift velocity in the detector at the given electric field, which can be determined by measuring the maximum drift time and taking advantage of the fact that the maximum drift distance is known from detector geometry.

5.5.1 Drift Velocity Measurements

Using the maximum drift time, dt, and the known cathode to u-wire distance of 198.6mm, the drift velocity can be calculated. Fig. 5.7 shows one of a histograms of charge cluster drift times from ^{60}Co calibration data on the left with a zoomed in view on the falling edge on the right, at a drift field of 475 V/cm. A maximum drift time of 110 μs is measured which corresponds to a drift velocity of 1.80 mm/us. In both Run 1 and Run 2, a cathode voltage of -8 kV is used, with the v-wires floating at -780 V, resulting in a drift field of 376 V/cm.

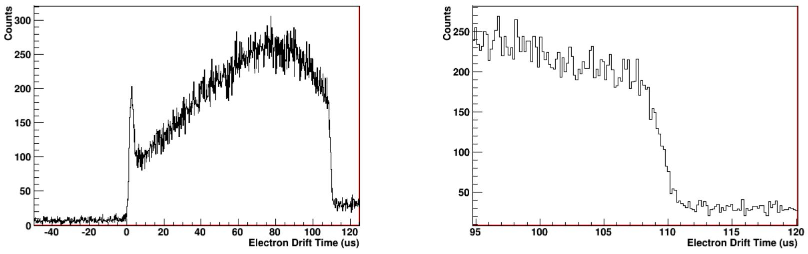

(a) (b)

Figure 5.7: Electron drift time in TPC2. The full spectrum is shown in (a) whereas (b) is a zoom in on the falling edge. There are some events with drift times longer than the measured maximum drift time, these are events where a charge cluster has been associated with the wrong scintillation signal. There is also a peak close to a drift time of zero. This peak exists because some charge deposits occur between the APD plane and the u-wires. This is a 6 mm gap between and drift times of these clusters can be a few μs. This peak contains both charge clusters from between the u-wires and the v-wires and from between the u-wires and the APD plane.

All the drift velocities calculated for different drift fields are summarized in fig. 5.8. In black are drift velocities measured by this method while the red points show drift velocity measurements performed using muons passing through both halves of the TPC. The two are consistent to within errors.

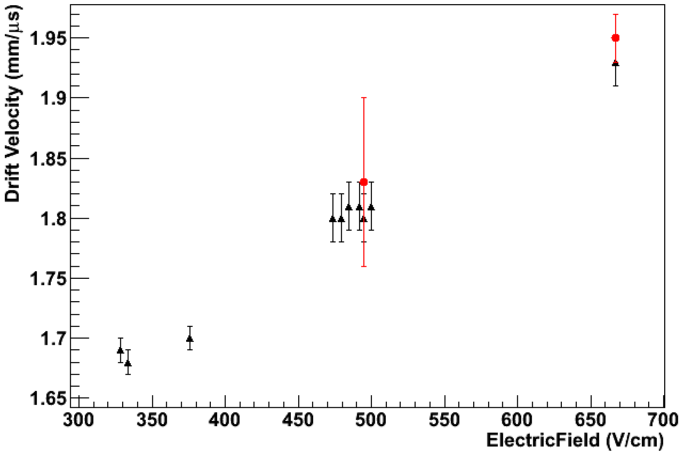

Figure 5.8: Electron drift velocities based on this drift time end point based method (black) and using muons which pass through both halves of the TPC (red).

5.6 Pulse Rise Times

Pulse rise times were studied as a possible tool to discriminate between beta events, which are more point-like, and gamma events, which can contain multiple scatters and be more spread out. If the electron cloud is larger in the Z-direction, then it should take longer to collect all the corresponding charge. As the detector drift velocity is 1.71 mm/μs, a charge deposit which is spread out over 4 mm will take more than 2 μs to collect and this information should be in the pulse rise time. Unfortunately, as the pulse rise times are already on this order or larger, using rise time information does not improve beta versus gamma discrimination over what can already be done with the triangular reshaping algorithm discussed in section 5.2.3.

Here, the rise time is defined as the time to go from 10% to 90% of the maximum of the of fast rise exponential decay pulse coming out of the FET. To get this pulse shape, a subsection of the waveform around a found signal has two 40 μs differentiators and two 1.5 μs integrators unshaped. An example of this unshaping is shown in figure

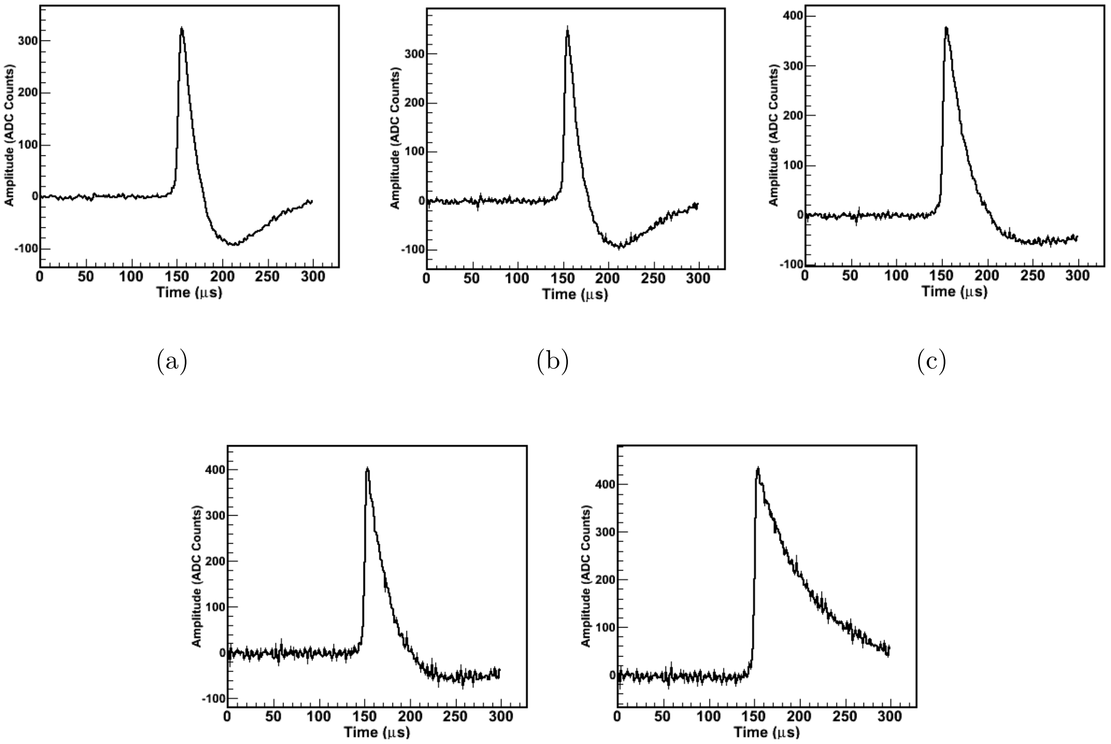

(a) (b) (c)

(d) (e)

Figure 5.9: The stages of unshaping of a raw u-wire pulse. (a) is the raw pulse, stages (b)-(d) are the unshaping of the 1.5 μs integration, 40 μs differentiation, 1.5 μs integration again, and 40 μs differentiation again. The resulting pulse at (e) is a fast rise exponential decay and is used for rise time measurement.

5.9. Figure 5.10 shows a comparison of rise times of events from thorium source calibration data (red) to low background data (green) and charge injection pulses (black). Although the low background data does contain some gamma events, 90% of these events are $2\nu\beta\beta$. Typical u-wire signals have rise times of 2-6 μs, which is 0.5-4.5 μs longer than the rise times of the charge injection pulses. This is due to the 2-3 μs rise time associated with the drift of the electron cloud between the v and u wires mentioned in section 4.1.2. The pulse height in this case is in ADC counts and is uncalibrated. For reference, one ADC count on a u-wire channel corresponds to 5-6 keV, and this number is channel dependent as different channels have different gains.

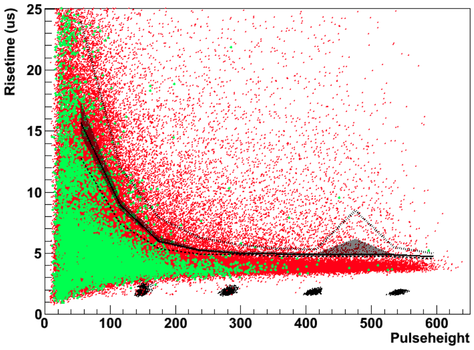

Figure 5.10: Shows rise time versus pulse amplitude for u-wire pulses from thorium source calibration data (red), low background data (green), and u-wire charge injection runs (black). The amplitude is in ADC counts. To translate this to an energy, in the charge one ADC counts corresponds to 5.23 keV after calibration on average. For example, a 2458 keV charge cluster would have an amplitude of roughly 470 ADC counts. This number is channel dependent due to gain variations

More importantly, there are many thorium events with much larger rise times. One example of such an event is shown in figure 5.11 which has two signals on the same channel. This is the same issue as presented in figure 5.3 and is addressed by the triangular reshaping algorithm. If one ignores these events with anomalously long rise times, the difference between the low background events and the thorium events is not enough to separate the two populations.

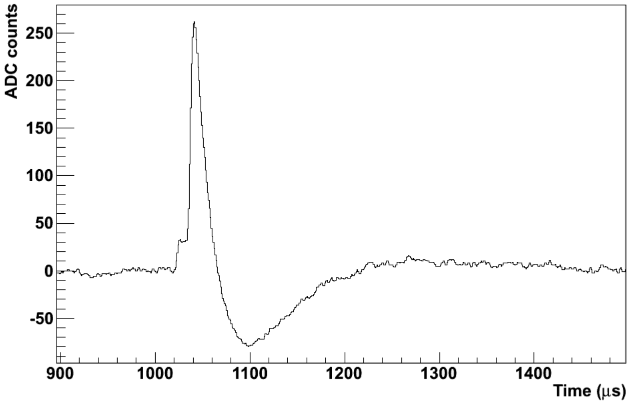

Figure 5.11: Example of a raw waveform with large rise time. In this example a two charge deposits are apparent, one roughly 20μs earlier than the other.

Chapter 6

PDF Generation and Background Model

6.1 Overview of PDF Generation

Probability distribution functions (PDFs) for fitting to low background physics data were derived from Geant 4 [51] Monte Carlo simulation. The experimental geometry is replicated in the simulation, including the detector, the HFE, the cryostat, and the lead and the walls of the salt mine. PDF generation proceeds in three steps: particle tracking, digitization and reconstruction, and energy spectrum generation.

6.1.1 Particle Tracking

Particle tracking includes propagation of radiation events to termination and recording charge deposits. The detector geometry is pixelized into .167 mm X .167 mm X .167 mm cubes, and for a given event each pixel in which energy was deposited has its position, total deposited energy, and the u-wire on which its charge is collected (if any) recorded. The Monte Carlo also has an optical simulation of the detector, but this is not used for this analysis.

6.1.2 Digitization and Reconstruction

The drift of individual pixelated charge deposits to u-wires in the weight field of the wire planes, calculated using MAXWELL [50], to determine the induction and charge collection signals for a given event. Signals for both the wires and the APDs are digitized using the appropriate electronics shaping times. The delay between scintillation and ionization pulses for the drift length is reproduced using the detector drift velocity. The digitized signals are then run through the same reconstruction algorithm as data.

6.1.3 Energy Spectrum Generation

The reconstructed output is used to determine event position and to classify events as single site or multi site. Because the same reconstruction algorithm used on data is applied, the simulation reproduces the effects of the discrete nature of the wires on the ability to discriminate between single site and multi site events (see section 7.5.3). The energy of the simulated event is calculated using the pixelated charge deposit information. For each charge deposit in the event, it is counted as part of the total detected energy if and only if its corresponding u-wire signal passes all analysis cuts. Some charge deposits are ignored either because the corresponding u-wire signal is below threshold or the charge is deposited in the dead region of the detector.

One modification to the default reconstruction in the simulation is required to make the effective threshold for v-wire detection match between simulation and data. In data, the threshold for finding v-wire signals is higher than in simulation, as shown in figure 6.1. The threshold matching is enforced stochastically: for each charge cluster found in simulation, there is some probability to have the cluster flagged as having no v-wire signal even if one is found. In this way the effective v-wire threshold in simulation can be made to match the threshold in data.

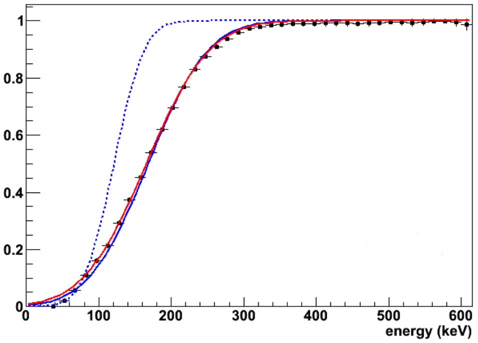

Figure 6.1: V-wire detection threshold is determined by taking the fraction of charge clusters with given energy that have a detected v-wire signal for both simulation (dashed blue line) and data (black points). The red line is a fit a sigmoid function to the data. V-wire detection is fully efficient above 350 keV. The blue line shows the effective detection threshold in simulation after a stochastic application of the threshold from data.

6.2 Backgrounds Model

6.2.1 Backgrounds from Copper Vessel and TPC Components

^{238}U, ^{232}Th, and ^{40}K are all common long lived radioactive isotopes and are typical backgrounds for low radioactivity experiments. In EXO-200 ^{40}K is only a background for $2\nu\beta\beta$, as it is a beta emitter with an endpoint of 1460 keV, well below the $Q_{\beta\beta}$ for ^{136}Xe. The ^{232}Th decay chain the 2615 keV gamma from ^{208}Tl and the ^{238}U chain contains the 2448 keV gamma from ^{214}Bi. All three of these isotopes can be present in the copper vessel, the APD planes, and the signal cables.

There are several cosmogenically activated isotopes of copper, including ^{54}Mn, ^{65}Zn, and ^{60}Co. Although ^{54}Mn and ^{65}Zn are β emitters with end points well below $Q_{\beta\beta}$ of ^{136}Xe, ^{60}Co has a sum peak at 2506 keV. This peak will be mainly multi site, but the detector position resolution has some granularity, meaning some of these events will be reconstructed as single site.

6.2.2 Backgrounds from Outside the Copper Vessel

According to daily measurements of the air in the clean room, the radon activity in the air around the cryostat is between 6 and 8 B/m^3. The volume of air between the cryostat and the lead has a volume between .191 m^3 and .252 m^3, for a total activity of between 1.2 and 2.0 Bq, resulting in $5 - 10 \times 10^5$ 2448 keV gammas a year from ^{214}Bi. There is also a possible background from the ^{238}U content of the lead, and these two backgrounds are spectrally degenerate.

6.2.3 Backgrounds Dissolved in the Xenon

The xenon is constantly repurified (see section 3.4.3) so internal cosmogenics which are not xenon are removed. Because the two most common isotopes of xenon in the detector are ^{134}Xe and ^{136}Xe, the two most likely xenon isotopes created by neutron capture are ^{135}Xe and ^{137}Xe. As shall be discussed later, there is some evidence for the presence of ^{135}Xe in the data, but little evidence for ^{137}Xe in the detector data.

There is also ^{222}Rn dissolved in the liquid xenon, which is a progeny of ^{238}U and has a decay chain including ^{214}Bi. Part of the ^{222}Rn decay chain is shown in figure 6.2. The radon content of the xenon has been measured using three different methods:

- β-decay of ^{214}Bi produces ^{214}Po which α decays with a 164 μs half-life, which is shorted than the 1 ms buffer recorded after an event trigger. 98.5% of of ^{214}Po decays occur within 1 ms of the ^{214}Bi decay, meaning one can look for events which contain a β followed by an α and thus measure the ^{222}Rn content of the detector. The α's can be distinguished from β's exploiting their lower charge to light ratio in liquid xenon. Figure 6.3 shows this discrimination in ^{214}Bi to ^{214}Po coincidences.

- In addition to the ^{214}Bi-214Po β-α coincidence, there is an α-α coincidence between ^{222}Rn and ^{218}Po with a relevant half-life of 3.05 minutes.

- The preferred method is to use a scintillation only analysis, referred to as alpha spectroscopy, to measure the number of 5.59 MeV and 6.11 MeV α's in the xenon (corresponding to the decays of ^{222}Rn and ^{218}Po in the xenon). This has

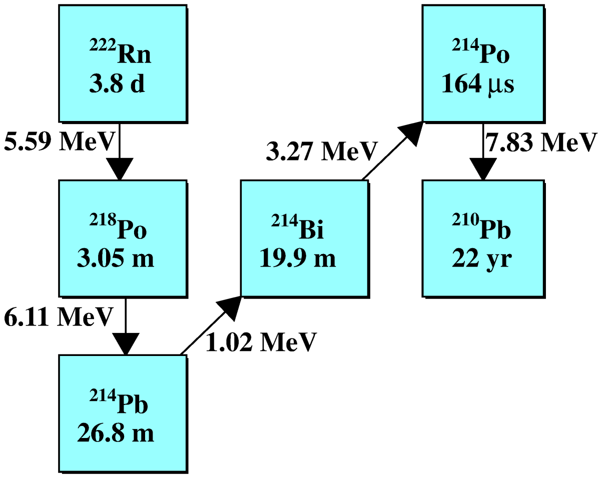

Figure 6.2: The ^{222}Rn decay chain through ^{210}Pb. Vertical arrows indicate α decays and diagonal lines indicate β decays. Half-lives and decay Q-values are shown.

lower systematic uncertainties because it does not requirement reconstruction of the weak charge signal from α's which often lack v-wire signals.

From the alpha spectroscopy analysis, there is a measured activity of 3.65±.37 μBq/kg of ^{222}Rn in the xenon. In order to calculate the overall effect of this concentration of radon, there are three cases to consider. Each requires a separate PDF:

- Radon in the bulk xenon
- ^{214}Bi on the cathode
- Radon in the inactive xenon

The rates of all three of these PDFs are constrained by the measured radon content in the xenon.

Radon in the bulk active xenon

^{222}Rn and its daughters decay in the bulk xenon, creating positive ions which drift toward the cathode. Studies of ^{214}Bi to ^{214}Po coincidences in the liquid xenon show that

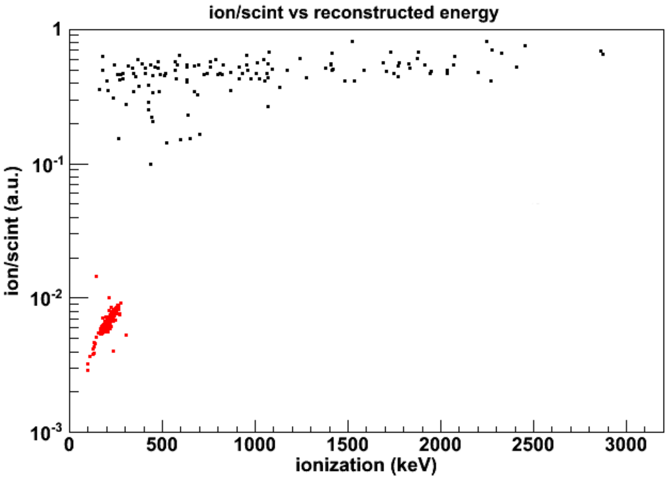

Figure 6.3: Ratio of ionization to scintillation for β's (black) and α's (red) taken from ^{214}Bi to ^{214}Po coincidences.

70% of these coincidences occur on the cathode. For decays occurring on the surface of the cathode, half of the ^{218}Po alphas will deposit their energy in the cathode and won't be detected. Because the detector doesn't see half of these coincidences, which means that actually 83% of ^{214}Bi decays occur on the cathode, with the remaining 17% in the bulk. This corresponds to a total activity of 78.2±7.8 μBq.

^{214}Bi on the Cathode

As discussed, 83% of the ^{214}Bi decays in the active liquid xenon occur on the surface of the cathode. This corresponds to an activity of 381.9±38.2 μBq. Half these decays will be vetoed as a ^{214}Bi to ^{214}Po coincidence.

Radon in the Inactive Xenon

There are 30.2 kg of inactive xenon, yielding a total activity of 110.3±11.0 μBq. Gammas emitted from this region can interact in the fiducial xenon, but as charge deposits in the inactive xenon can are not seen by the detector, these events cannot be rejected using coincidences.

^{220}Rn in the Liquid Xenon

Studies of alpha-alpha coincidences in the liquid xenon have determined that the ^{220}Rn content of the liquid xenon, if any, is negligible. ^{220}Rn decays with a half-life of 55.6 s to ^{216}Po which has a half-life of only 145 ms. The short half-life of ^{220}Rn makes its absence in the xenon plausible, as it would have to diffuse into the xenon from the outside, and this takes time (^{222}Rn, on the other hand, has a half-life of 3.8 days). The short half-life of ^{216}Po admits a search for ^{220}Rn in the xenon by looking for a coincidence of two alpha events within 1.4 s of each other (this corresponds to almost ten half-lives of the decay).

Alpha events can be distinguished from gamma or beta event because their light to charge ratio is much higher. Using the Run 1 data, 3 such candidate coincidences were found in a sample of 24.8 live days. After considering fiducial cuts on both events in each coincidence, the total detection efficiency for these coincidences is 38.8%. This results in a 90% confidence level limit of .04 μBq/kg ^{220}Rn decays per day in the xenon, which is negligible.

Chapter 7

Fits to Data

7.1 Analysis Cuts

As the final result comes from fits of PDFs generated from simulation to the data, the cuts applied to both simulation and data are made as similar as reasonably possible. The cuts described below are derived mainly from the need to achieve good fits to the source calibration data (see section 7.5), which is used to verify the PDF generation process. The final result comes from maximum likelihood fits of PDFs to the data, and the reliability of this procedure is dependent on ensuring that the detector behavior in simulation matches that in data.

7.1.1 Fiducial Volume

As explained in section 3.3.1, part of the detector has unknown charge collection efficiency, and charge deposits in this dead layer are ignored. This is different than a fiducial cut, as the event is not thrown out. Instead, if an event has two charge deposits, one of which occurs in the dead region of the detector and one of which occurs deeper inside the detector, the deposit in the dead layer is ignored and the event is treated as single site.

This same cut is also applied to the simulation data. This will cause some true multi site events to be moved into the single site spectrum or have their energy underestimated, but as long as the same cut is applied to the simulation, this effect is taken into account. What would not be acceptable would be to ignore this effect and still fit PDF generated from simulation to the data. Then some true multi site events in the data would look like single site events, while the analogous events in the simulation would not.

The resulting fiducial volume must mimic the true charge collection geometry, which is hexagonal, not circular. In the X-Y plane the fiducial area is a hexagon with an apothem of 163 mm. The teflon reflector has an inner-radius of 183 mm, and an additional radial cut requires all clusters occur inside this radius. A cut in the Z-direction to removes events which occur within 5 mm of the cathode, and within 10 mm of the anodes, where large shielding grid corrections would be necessary (see 7.3.1). While the fiducial cuts in the X-Y plane simply ignore clusters rather than disqualifying events from the analysis, this is not the case with the Z direction cut. For example, if there is an event with two charge deposits, one 5 mm away from one of the anodes, and one in the bulk of the detector, far away from any boundary, the entire event would be eliminated by the fiducial cut.

7.1.2 Three Dimensional Position Reconstruction

In order to apply a fiducial volume cut, all clusters in an event must have three-dimensional position reconstruction, which requires both a scintillation signal and a v-wire signal associated with every charge cluster. The scintillation detection is fully efficient above 700 keV, but a multi site event may contain a charge cluster below the v-wire reconstruction threshold even if the event energy is rather high.

Efficiency

In ^{228}Th source calibration ($0\nu\beta\beta$) an efficiency for three-dimensional reconstruction of 71% (70%) is calculated from Monte Carlo. As a crosscheck, the three dimensional reconstruction efficiency is compared between simulation and data for single site events in figure 7.1. The efficiency is calculated by taking the fraction of events passing all other cuts which also pass the three dimensional reconstruction cut. The definition of single site here needs to be made clear: if an event has one charge deposit which does not have a v-wire signal, and the that event is not being cut for not having three-dimensional reconstruction, that cluster must be culled as it cannot be gauntleted to be fiducial. There is agreement to within errors except in the ^{208}Tl full energy deposition peak. Otherwise there is a general trend toward lower efficiency at higher energies for ^{228}Th source data and low background data.

7.1.3 Coincidence Cuts

There are three veto coincidence cuts relevant for this analysis:

- Events within 25ms of a veto panel hit are removed from the data, which has an associated dead time of .58%.

- If any two detector events occur within 1 s of each other, both events are vetoed. Such coincident events are most likely to be ^{214}Bi to ^{214}Po coincidences in the xenon or cosmic ray induced, where some cosmogenically activated isotope has to undergo multiple short half-life decays to reach a stability. This cut incurs a dead time of 3.3%.

- The position reconstruction of the detector allows for reconstruction of muons traversing the TPC. After such an event, there is some risk of exotic cosmogenically activated isotopes. Events within 60 s after a muon traversing the detector are vetoed, which leads to a 5.0% dead time.

In total these three cut result in a dead time of 8.6%.

7.1.4 Charge to Light Ratio Cut

Figure 7.2 shows a two dimensional plot of the energy of events from ^{228}Th runs, with the ionization energy on the x-axis and the scintillation energy on the y-axis. There is a population of events which have a large scintillation signal and a weak ionization signal. These events are mainly due to gammas which multi scatter and deposit some energy in the dead region of the xenon and some in the fiducial xenon.

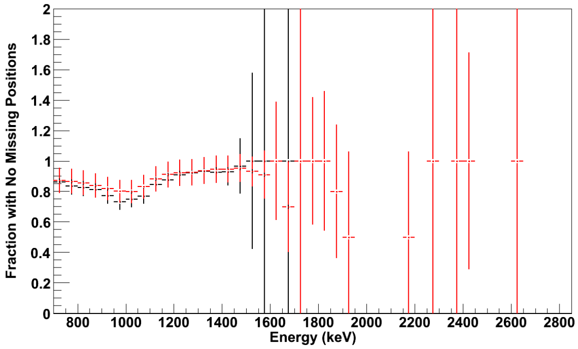

(a)

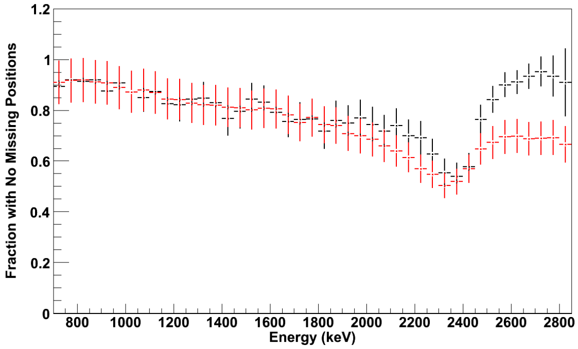

(b)

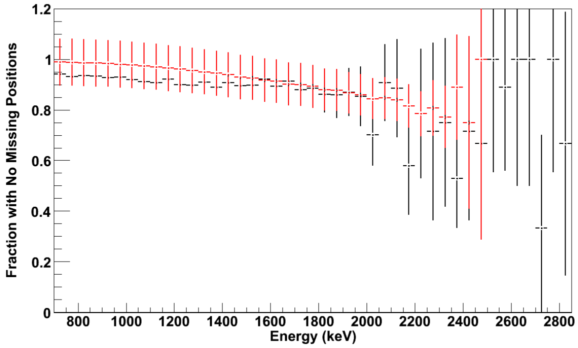

(c)

Figure 7.1: The fraction of events passing all other cuts that have full three dimensional reconstruction for single site events in data (black) and simulation (red) for (a) ^{60}Co, (b) ^{228}Th, and (c) low background physics data, which is compared to $2\nu\beta\beta$ simulation. 90% of single site low background events are $2\nu\beta\beta$, so this last comparison is reasonable. While the error bars on the data are purely statistical, the errors bars on the simulation include a 9.4% uncertainty determined from calibration source agreement studies (see section 7.5.2).

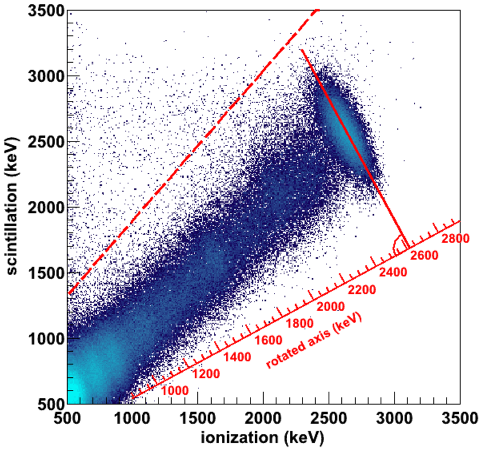

Figure 7.2: Scintillation signal versus ionization signal for ^{228}Th events. The dashed red lines shows the threshold for the light to charge ratio cut. The red axis represents the the energy scale for the calibrated energy after combining the scintillation and ionization signals (see section 7.4).

All the scintillation light is detected, but not all of the charge. Any event which is more than 4σ away from the average light to charge ratio for a given energy is cut. This cut removes .4% of events in ^{228}Th source calibration data.

This cut also removes a second class of events as shown in figure 7.3. This is a similar plot, with ionization energy on the x-axis and scintillation energy on the y-axis, but for low background physics data and on a log-log scale. The dashed red line represents the cut determined from ^{228}Th calibration data. There are two separate groups of events: the alphas, which are above the red line and have more light and less charge, and the betas and gammas, which are below the red line and have less light and more charge. Although alphas generally have much higher energies than $Q_{\beta\beta}$ of ^{136}Xe, seeing that they are in fact there and behave as expected remains important.

7.1.5 Other Cuts

The other cuts which are applied are:

- Events with multiple scintillation signals are cut. This is targeted to address ^{214}Bi to ^{214}Po coincidences.

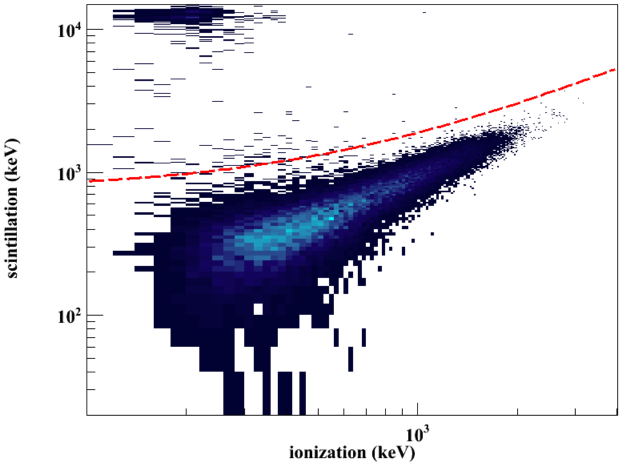

Figure 7.3: Scintillation signal versus ionization signal for low background physics data. The dashed red line represents the light to charge ratio cut.

- Any event tagged as a muon traversing the detector is cut from the data.
- Similar to looking for muons traversing the detector, there are noise events which are identified by the data processing chain, these are cut from the data.
- Any event with a saturated channel is cut from the analysis. These events are typically due to muons or electronic noise.

7.2 Masking Procedure

Initially, only the first ten minutes of every half hour of run time for each run were available for analysis. Not until the reconstruction algorithm, simulations, determination of systematic uncertainties, and calibrations were finalized were the remaining two-thirds of the data made available. This was not a blinding scheme, a half-life limit for $0\nu\beta\beta$ could be calculated from the information available. However, it did act as a safeguard against some forms of bias, as the analyzers did not know the final result before finalizing analysis cut and calibrations. There was one change to the background model after unmasking: there was evidence that the ^{137}Xe content on the detector was being overestimated, which would cause an overestimate of the background for $0\nu\beta\beta$.

7.3 Position Dependent Corrections

Position dependent energy corrections are made both to the ionization and scintillation signals. Once charge and light signals are combined to determine the event energy, no further corrections are made. There are three known affects for which a correction is made to the charge signal:

- A drift time (Z) dependent correction for the effects of the electron lifetime.
- A Z dependent correction for the imperfect shielding of the v-wires.
- A U dependent correction for the gains of the u-wire channels.

These corrections are all determined using source calibration data. The light signal has one correction applied for the light collection efficiency of the detector, which is a phenomenological model also determined from source calibration data.

7.3.1 Ionization Signal Corrections

The gains of the u-wires are measured using source calibration. In thorium calibration data, the ^{208}Tl double escape peak is prominent in the energy spectra for events where charge was collected on only one u-wire. This peak has a different position for each wire, and this corresponds to the difference in gains between the wires. Figure 7.4 shows fits to this peak for two different u-wire channels, showing the difference in peak position between them. From these fits a relative correction factor is determined.

As described in section 3.1.1, the xenon contains impurities. The longer the charge has to drift before being collected, the more likely a given unit of charge is to be captured by an electronegative impurity. There are source calibrations runs almost every day used to measure the electron lifetime by dividing the detector into several bins in the Z direction, and in fitting the full energy deposition peak. The position

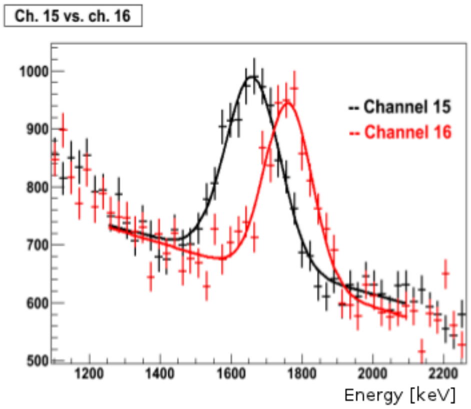

Figure 7.4: Fits to the ^{208}Tl double escape peak for two different u-wires channels.

of this peak versus Z is fit to an exponential to extract the electron lifetime as shown in figure 7.5. The electron lifetime versus time is fit to a fourth order polynomial to determine the purity correction of low background data runs, which occur in between the source calibration runs.

Because the v-wires do not perfectly shield the u-wires from the charge deposited, the drift of the positive charge can effect the measured charge collected. This effect is described in [52] and has a $\frac{1}{d}$ dependence, where d is the distance from the charge deposit to the v-wire plane. Thorium calibration data is fit to this model to determine the correction as shown in figure 7.6.

7.3.2 Scintillation Signal Corrections

The correction light collection efficiency is determined using the charge collection signal to isolate full energy deposition peak events for different regions of the detector. The scintillation response of these full energy deposition events in the charge spectrum is used to produce a light collection efficiency map for the detector. This is applied to correct the scintillation signal based on the events position.

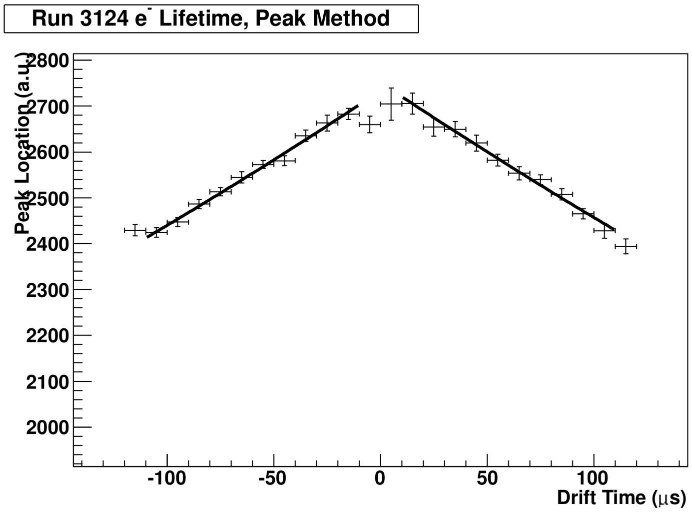

Figure 7.5: Fitted peak position versus Z for ^{232}Th source calibration, which is fit to an exponential to extract the electron lifetime.

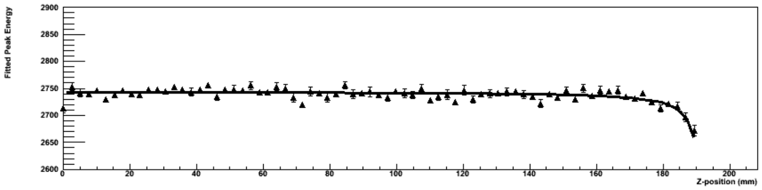

Figure 7.6: ^{232}Th source calibration versus Z after the purity correction with a fit to the $\frac{1}{d}$ dependence for shielding grid inefficiency.

7.4 Calibrations

The detector is calibrated using all three calibration sources: ^{60}Co, ^{137}Cs, ^{228}Th. This admits a total of four gamma lines for calibration: 661.7 keV from ^{137}Cs, 1173.2 keV from ^{60}Co, 1332.5 keV from ^{60}Co, and 2615.5 keV from ^{208}Tl, which is a part of the ^{228}Th decay chain. There are separate calibrations for single site and multi site events. For each category of event the ionization and scintillation measured are combined to determine the total event energy using the following formula

$$E_0 = Icos(\theta) + Ssin(\theta) \tag{7.1}$$

Where I is the measured ionization signal, S is the measured scintillation signal, and θ is chosen to optimize the detector energy resolution at 2615.5 keV. There are separate θ's for single site events and multi site events, and each is calibrated with a second order polynomial ($E = a * E_0^2 + b * E_0 + c$). The residuals from the calibration and the energy resolution as a function of energy for both single site and multi site events is shown in figure 7.7. The resolution at $Q_{\beta\beta}$ of ^{136}Xe is 1.67% for single site events and 1.84% for multi site events. Included are fits to an energy resolution model. The one sigma uncertainty on the energy can be parametrized as

$$\sigma_{tot}^2 = p_0^2 E + p_1^2 + p_2^2 E^2 \tag{7.2}$$

So that the energy resolution, R(E), versus energy is

$$R(E) = \frac{\sigma_{tot}^2}{E} = \sqrt{\frac{p_0^2}{E} + \frac{p_1^2}{E^2} + p_2^2} \tag{7.3}$$

The first term represents a $1/\sqrt{E}$ from Poisson statistical fluctuations, the second represents a constant width from noise, and the third term is any width which grows proportional to energy. This third term can represent an uncorrected spatial or time dependence of the energy calibration. Currently the second term, due to noise, dominates the energy resolution at $Q_{\beta\beta}$ of ^{136}Xe.

Figure 7.8 shows the variation in the detector energy calibration over time from

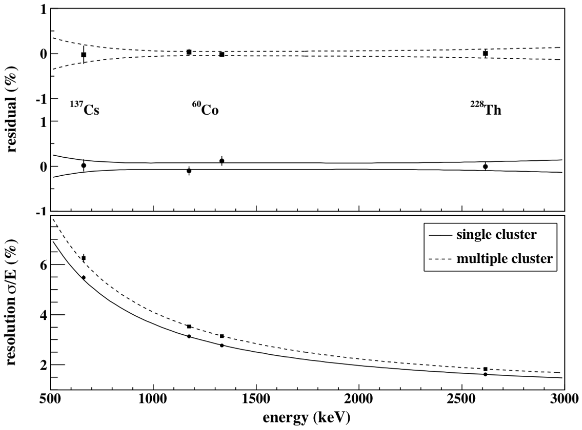

Figure 7.7: Top: Residuals for energy calibrations for single site events (circles, solid lines) and multi site events (boxes, dashed lines). The lines show the one sigma error band for each for each calibration. Bottom: Energy resolution as a function of energy for single site (circles, solid lines) and multi site (boxes, dashed lines).

daily calibration runs. Most runs use the 2615 keV line from weak ^{232}Th source and some use the 1332 keV line from the weak ^{60}Co source. The detector calibration is stable with time so no time dependent correction to the event energy is made.

7.4.1 Pair Production Events

One way to study the ionization signal in more detail is to look at pair production events from the 2615 keV gamma from ^{208}Tl. In a pair production event, a gamma, in the presence of a nucleus, generates an e^{+}-e^{-} pair. The total energy of this pair is the full energy of the gamma minus twice the mass of the electron. Because the e^{+} will quickly annihilate with an electron in the bulk xenon, the events consists of one charge cluster which has 1593 keV of energy (2615 keV minus twice the mass of the electron) and possibly two or more others corresponding to the scatters of the two 511 keV gammas given off in the annihilation. The two 511 keV gammas are back-to-back, which means such pair production events where both 511 keV gammas full energy deposit will have exactly three charge clusters that will be collinear. Using this geometric constraint, one can look for golden pair production events. This provides

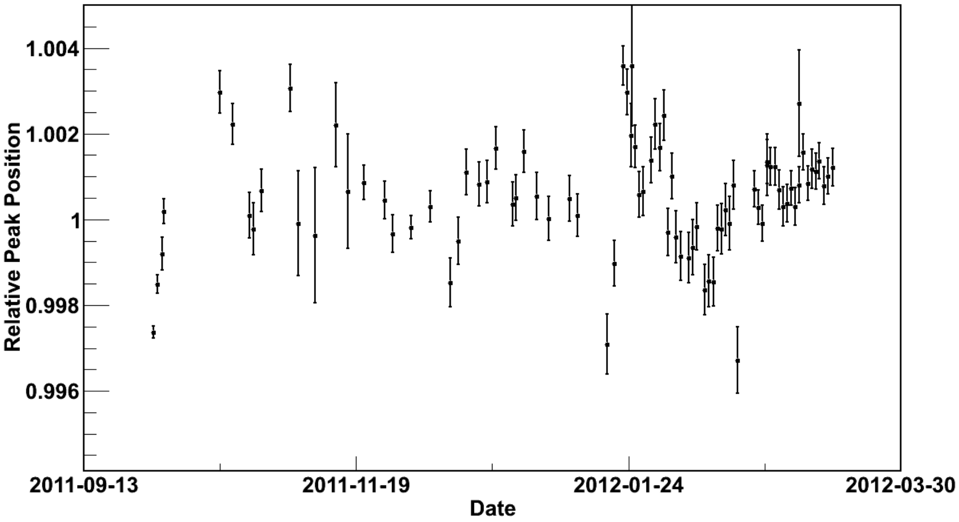

Figure 7.8: Measured peak position using the detector calibration divided by true peak energy over time. Deviation is less than 0.4% over all of Run 2.

two additional calibration peaks for the charge only calibration: one at 511 keV and one at 1593 keV. Fits to each of these peaks are shown in figure 7.9.

The energy of the 1593 keV peak does not follow the energy calibration of gammas in the detector. Figure 7.10 shows the residuals for the charge only calibration including the 1593 keV peak from these golden pair production events. The energy scale for these events is offset from the rest of the ionization energy scale. This may indicate that the charge only energy scale of more point-like events, such as the pair production cluster or $2\nu\beta\beta$ events, may differ from gamma events. As a result, the β energy scale is a parameter determined from fits to the low background physics data and its uncertainty is absorbed as a systematic (see section 7.7.2).

7.5 Source Agreement

PDFs generated from the simulation can't be used unless the simulation is verified to accurately reproduce the detector performance. Fitting to source calibration data is critical to the analysis, as it is the simplest way to compare the simulation to the

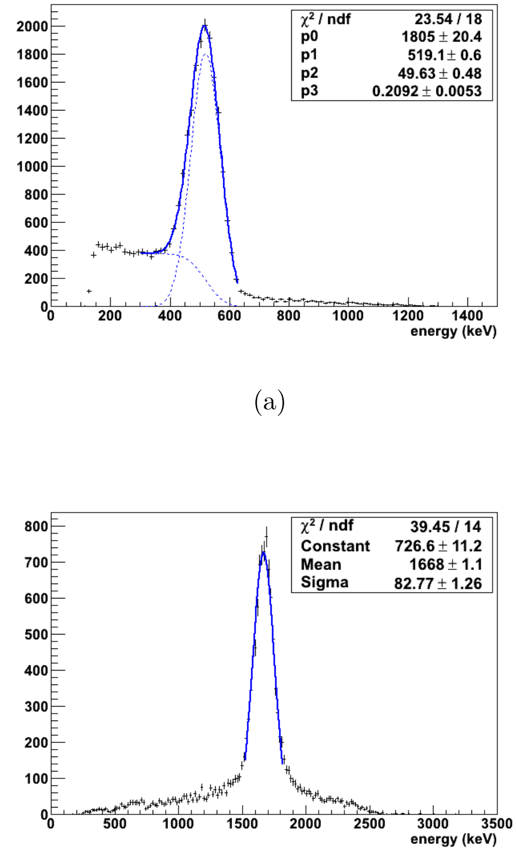

(a)

(b)

Figure 7.9: Fits to the (a) 511 keV and (b) 1593 keV charge only peaks from pair production events.

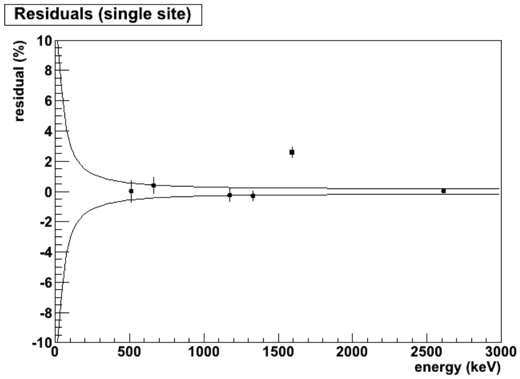

Figure 7.10: Residuals for ionization signal only energy calibration including the 1593 keV and 511 keV peaks from pair production events.

data. There are three critical parameters for these source agreement studies: shape agreement, rate agreement, and single site fraction agreement. This means the PDFs must reproduce the energy spectra of signals in the detector, must accurately calculate the detection efficiency of different signals, and must correctly calculate the fraction of all events which are single site for a given signal.

Only weak source runs are used for source agreement studies. This is designed to ensure that any change in detector perform due to high trigger rates is minimized. Only runs using the solicited trigger are applied to source agreement studies. The prescale trigger has a bias toward multi site events that adversely affects singe site fraction studies (see section 7.5.3). Only the ^{60}Co and ^{228}Th sources are used for source agreement, as the ^{137}Cs source full energy deposition peak is below the analysis threshold of 700 keV. Only source calibration runs at the cathode are used for source agreement studies.

7.5.1 Shape Agreement

Figures 7.11 and 7.12 show fits of PDFs, generated from simulation, to data for ^{228}Th and ^{60}Co source calibration runs. The energy calibration is applied to the data and the PDF is smeared with the measured energy resolution of the detector. The energy offset and the energy resolution parameters are allowed to float in the fit.

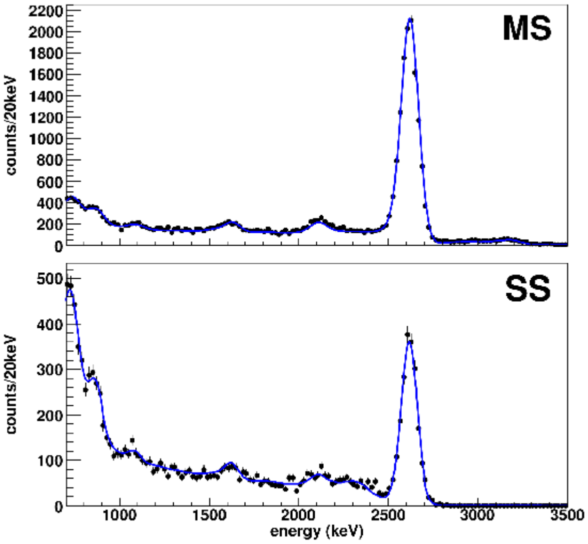

Figure 7.11: Maximum likelihood fit of a ^{228}Th source calibration PDF (blue line) to calibration data (black points).

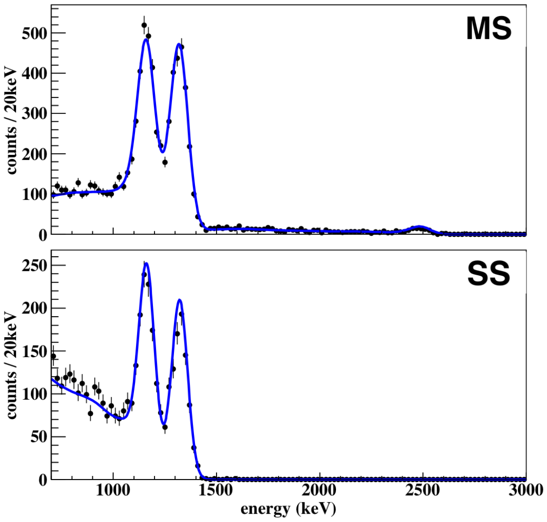

Figure 7.12: Maximum likelihood fit of a ^{60}Co source calibration PDF (blue line) to calibration data (black points).

Table 7.1: χ^2/NDF goodness of fits for the single site and multi site spectra of ^{60}Co and ^{228}Th source agreement.

Source	Single Site	Multi Site
^{60}Co	.7505	1.0247
^{228}Th	.8405	1.4494

Table 7.2: Percentage of toy Monte Carlos for each PDF which have a $Kolmogorov-Smirnov$ distance greater than or equal that of the data from the given PDF.

Source	Single Site	Multi Site
^{60}Co	2.2%	3.6%
^{228}Th	85%	75%

The goodness of fit is determined using two different methods. The first is a χ^2 metric, summarized in table 7.1. Although these reduced χ^2 figures are all below 1.5, the χ^2 goodness of fit test assumes Gaussian errors and requires all bins to have more than five events to be valid. Therefore a second statistical test is required to verify the goodness of fit, so a $Kolmogorov - Smirnov$ test is implemented as an additional check. The $Kolmogorov - Smirnov$ test compares two distribution using a distance defined as the maximum difference between their cumulative distribution functions, which follows a known distribution under certain assumptions. One of these assumptions is that the two distributions be continuous, and while the PDF is continuous, the data points clearly are not. There is a second requirement that the test distribution is fully specified, which means none of its parameters can be estimated from the data. This is not the case if one is comparing a best fit PDF to the data. This means that the typical $Kolmogorov - Smirnov$ distribution cannot be used, so instead the distribution corresponding to each PDF (in this case ^{60}Co and ^{228}Th) is calculated using a series of $\sim$ 1000 toy Monte Carlos. These results are summarized in table 7.2. The goodness of fit is reasonable, showing both distributions are within 2-3 sigma.

7.5.2 Rate Agreement

Comparing the activities of the calibration sources as measured by the detector to their known activities is an important crosscheck on the analysis. A discrepancy could point to systematics in the analysis, and one of the requirements for the analysis is to get better than 10% agreement between the measured calibration source activities and their known activities. The source activity, γ is calculated by

$$\gamma = \alpha \frac{N}{\epsilon T} \tag{7.4}$$

where α is a correction to account for pileup. There is a cut in the analysis that removes events with more than one scintillation signal, but in source calibration data these are just accidental coincidences resulting from the high rate of the source. α is the inverse of the fraction of all scintillation signals which come from single scintillation signal events. N is the number of events in data above threshold as determined by the fit, ϵ is the detection efficiency from Monte Carlo, and T is the livetime of the calibration run. These runs use the random trigger, so the livetime is the number of triggers times the effective time length of the waveform trace. The rate discrepancy, defined as $\frac{Measured-Expected}{Measured}$ is 9.4% for ^{228}Th and 8.9% for ^{60}Co.

The discrepancy is similar for both sources, which points to some error by an overall factor. One possible explanation for this would be if the calibration sources were not quite at the correct locations, i.e. if they were closer to the detector or farther away than expected. The source calibration simulations were repeated with the calibration sources at three positions: the nominal source position, 1 cm closer to the detector, and 1 cm farther away. The results are summarized in table 7.3. The result of moving the source in or out affects the measured rate by roughly ±10%. The calibration source is not likely to be 1 cm closer to the detector than its nominal position, but this position is a critical parameter for source rate agreement studies. Future analyses may require a more precise determination of the source location to improve upon this.

Table 7.3: Rate disagreement as a function of source distance from the detector for ^{60}Co and ^{228}Th sources. Percentage is $\frac{Measured-Expected}{Measured}$

Position	^{60}Co	^{228}Th
1 cm closer	-1.5%	2.6%
nominal	9.4%	8.9%
1 cm farther	21.5%	21.8%

7.5.3 Single Site Fraction Agreement

As EXO-200 uses single site versus multi site discrimination as a background rejection technique, there is a systematic uncertainty incurred if the simulation predicts a different single site fraction for a given signal than what exists in the data. However, the trigger can introduce a bias. Consider two events, a single site events with 1 MeV of energy, and a two site event consisting of two 500 keV charge deposits. Suppose the trigger is a prescale level one (which means every other trigger will be recorded to disk) then the single site event has a 50% chance of being recorded. The two site event, will always have one of the two charge deposits trigger a record, so this event will be recorded 100% of the time. Because a multi site event has multiple charge clusters above the trigger threshold, it is more likely to be recorded by the prescale trigger than a single site event. As solicited trigger runs don't use the event information to decide when to record to disk, there is no trigger bias toward multi site events.

Figure 7.13 shows the single site fraction discrepancy between simulation and data for several energy bins for both ^{60}Co and ^{228}Th. The single site fraction is defined as $\frac{\#\, single\, site\, events}{\#\, single\, site\, events + \#\, multi\, site\, events}$ and the discrepancy is defined as $\frac{Data-Simulation}{Simulation}$. The largest discrepancy for any bin for either source is 8.5%, and this is quoted as the single site fraction systematic uncertainty.

7.5.4 Agreement for Different Fiducial Volumes

A check on the source agreement studies is to look for agreement with different fiducial cuts. Figure 7.14 shows source shape comparisons for several different fiducial volumes, defined in terms of the apothem of the hexagon, starting from the nominal

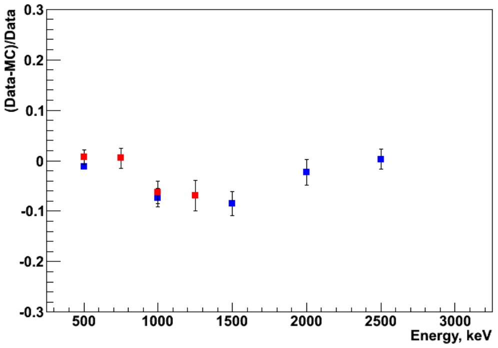

Figure 7.13: Single Site fraction discrepancy between simulation and data for ^{60}Co and ^{228}Th as a function of energy bin. The largest discrepancy observed is 8.5%.

Table 7.4: Source rate and single site fraction agreement versus hexagonal fiducial cut. The Z cuts here are on the absolute value of Z.

Apothem Cut	Z cut	Rate Discrepancy	Single Site Fraction Discrepancy
163mm	5mm to 182mm	9.4%	-6.9%
133mm	5mm to 152mm	13.3%	-2.7%
103mm	5mm to 122mm	13.6%	1.4%
73mm	5mm to 92mm	12.0%	4.9%
43mm	5mm to 62mm	12.6%	2.1%

fiducial cut of 163mm with each cut 30mm deeper into the detector. Each successive fiducial cut also removes 30mm from each anode, but not from the cathode. The data and the PDF have the same fiducial cuts applied. This check shows that the detector performance is understood both on the outside of the detector as well as deep inside the detector.

Table 7.4 summarizes the source rate and single site fraction agreement as a function of fiducial cuts. The single site fraction discrepancy in this case is over all energies from 700 keV to 3500 keV. Source rate discrepancies range from 9.4% to 13.6%. This spread has to do with the fiducial volume uncertainty. The single site fraction discrepancy is worst for the nominal fiducial volume at -6.9% and improves to as small in absolute value as 1.4%, so good single site fraction agreement is seen for all possible fiducial volumes.

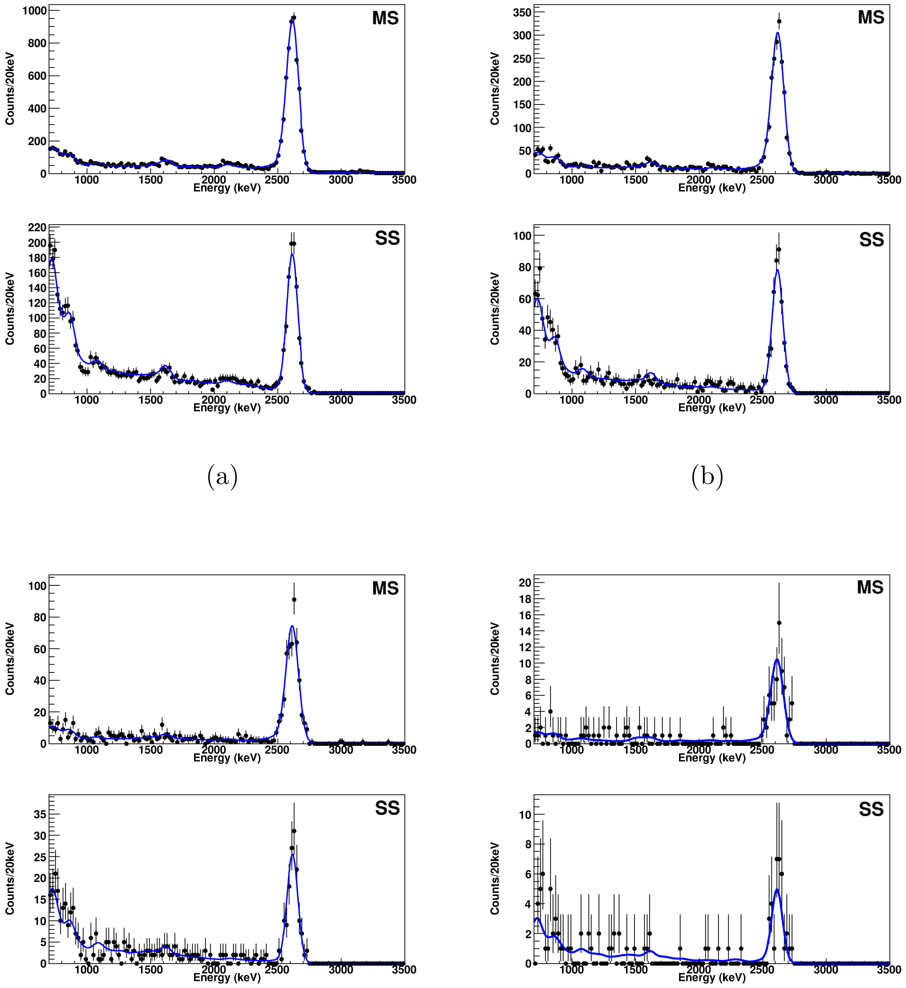

(a) (b)

(c) (d)

Figure 7.14: Fits of ^{228}Th PDFs to data for different fiducial cuts. The fiducial cuts are hexagonal, and defined by the apothem of the hexagonal (the radius of the circle inscribed in the hexagon). The different cuts shown are (a) 133mm, (b) 103mm, (c) 73mm, and (d) 43mm.

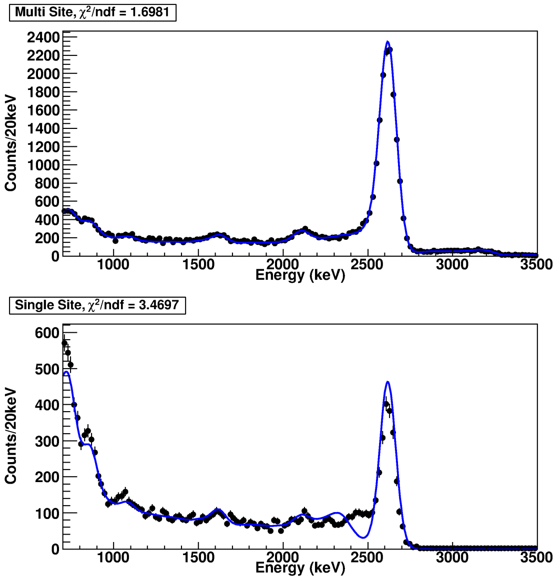

Figure 7.15: Source agreement plots where the three dimensional reconstruction cut is not applied in simulation (blue curve) or data (black points).

7.5.5 Effect of Three Dimensional Reconstruction Cut on Source Agreement

The requirement that all clusters have three dimensional reconstruction is particularly critical for source shape agreement. Figure 7.15 shows the ^{228}Th source calibration fit without the three dimensional reconstruction cut. The cut was changed both in the generation of the PDF and in the analysis of the data. In both cases, if an event contains a cluster that is lacking three dimensional position reconstruction, this cluster must be culled, as it cannot be certified as a fiducial cluster. The multi site agreement is not affected, but the single site spectrum shows particularly large discrepancies between 2200 and 2500 keV, and this interval includes $Q_{\beta\beta}$ of ^{136}Xe. The resulting discrepancy seen in figure 7.15 is likely due to events with one low energy Compton scatter in the dead region followed by a second scatter in the bulk xenon. Any single site event spectrum which falls between the Compton shelf and the full energy deposition peak must either be a true multi scatter events or involve a gamma which scatter outside the active volume before entering the detector. This

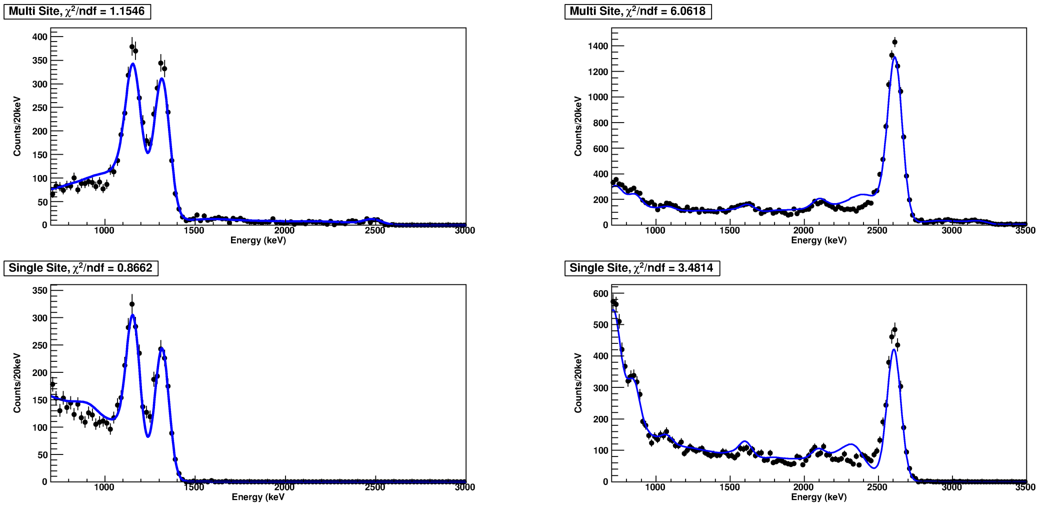

(a) (b)

Figure 7.16: Source agreements plots for (a) ^{60}Co and (b) ^{228}Th without the triangular reshaping algorithm applied in reconstruction. The reduced χ^2's are included.

is why the discrepancy is particularly obvious in this part of the spectrum.

7.5.6 Effect of Triangular Reshaping on Source Agreement

Addressing the issue of multiple u-wire signals on the same channel is also critical for source shape agreement. This issue is currently addressed in the signal finding algorithm using a triangular reshaping, as discussed in section 5.2.3. Figure 7.16 shows fits of PDFs to the data for source calibration without the triangular reshaping stage in signal finding. This means that in these plots both the simulation and the data have been reprocessed without the reshaping stage, with no other aspect of the analysis changed.

The ^{60}Co fits still have reduced χ^2's below 1.5, but the ^{228}Th fits are both larger than 3. In the ^{228}Th fits most of the disagreement happens above the gamma lines from ^{60}Co. As was seen in figure 5.5 in section 5.2.3, the triangular reshaping has a greater effect on events at energies above 1.5 MeV than below. This likely explains why the ^{60}Co is not as heavily affected by this modification to the signal finding in reconstruction. In the ^{228}Th fits, the agreement is worst between 2200 and 2500 keV, near $Q_{\beta\beta}$ of ^{136}Xe. The single site low background spectrum above 2200 keV is heavily affected by the additional of the reshaping algorithm. The single site fraction discrepancy over the entire energy range is 35.2% for ^{60}Co and 32.9% for ^{228}Th.

7.6 Data Taken Between 4/3/12 and 4/15/12

A change to the DAQ on 4/3/12 caused the data from one of the u-wire channels to not be recorded. The detector was still triggering on this channel, but its waveforms were not being saved to disk. This issue continued through May 11th, however no data after April 15th is used for the Run 2 analysis. Only a total of 6.3 live days out of the total 120.7 days of detector livetime in Run 2 were affected and only one out of 76 u-wires channels was lost, reducing the effective exposure by a negligible amount.

Figure 7.17 shows an additional crosscheck, source shape agreement was revisited in an attempt to understand the effect of this data on the fits. The black points show the thorium source energy spectrum for data before April 3rd. The blue line is a weighted average of the source data from before April 3rd and source data runs from between April 3rd and May 11th. The weighting is based on the overall fraction of the livetime, so the time between April 3rd and April 15th only makes up $\frac{6.3}{120.7}$=5.2% of the blue curve. Little if any change is noticeable, and this can be extended to the fits to the low background data: the overall effect of this missing channel is negligible and does not affect the final result.

7.7 Fits to Low Background Data

Figure 7.18 shows the binned maximum likelihood fits to the low background spectrum. The total livetime was 120.7 at 98.5 kg fiducial mass, for a total of 32.6 kg-yrs. of data. The background model includes all PDFs discussed in the previous chapter, except ^{137}Xe, as there was no conclusive evidence for its presence in the fit.

The individual PDFs are smeared with the measured energy resolution function

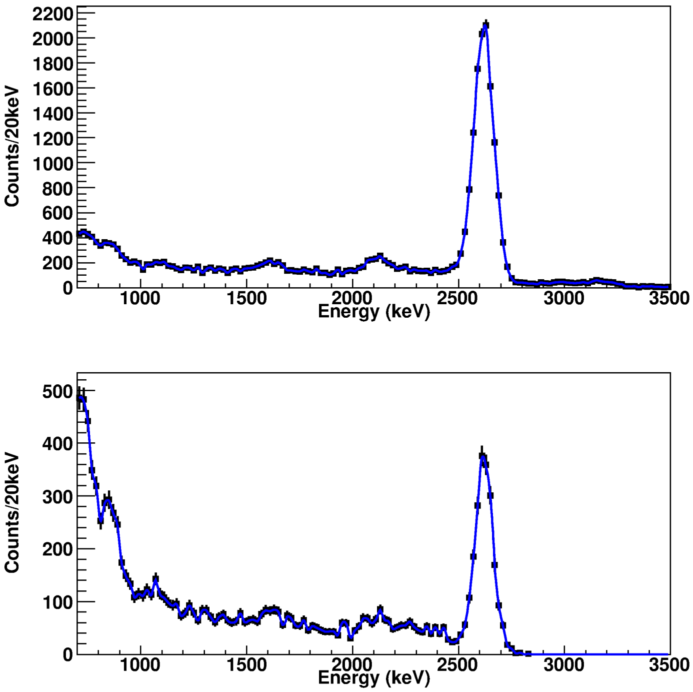

Figure 7.17: Comparison of ^{228}Th source data before April 3rd to a weighted average of source data over the course of Run 2.

of the detector and the energy resolution parameters are allowed to float in the fit with a Gaussian constraint determined by the uncertainties of the energy resolution measurement. The single site fraction is also allowed to float for each PDF individually, with a Gaussian constraint with a mean determined by the predicted value from Monte Carlo and a one sigma error of 8.5% determined from source agreement studies. If the single site fraction of $2\nu\beta\beta$ is allowed to float without constraint, it would have discrepancy of 5.8%, which is less than the 8.5% error quoted based on source agreement studies. A half-life of $2\nu\beta\beta$ of ^{136}Xe is $(2.23 \pm .017(\text{stat}) \pm .22(\text{sys})) \times 10^{21}$ yr, which is in agreement with [46] and [15].

7.7.1 ^{135}Xe

Figure 7.19 shows a profile likelihood for ^{135}Xe in the low background physics data. Each point represents a fit to the low background data where the amount of ^{135}Xe is fixed and all other parameters are allowed to float, with a different likelihood determined for each fit. Here the negative log likelihood (NLL) is used, which is defined as NLL=-log(L), where L is the likelihood. The values plotted are the difference between the NLL for a given fit and the minimum NLL. The total measured number of ^{135}Xe events in the data is 635 ± 135, which is 4.7σ different from zero.

80.6% of the xenon is ^{136}Xe and 19.4% of the xenon is ^{134}Xe, with all other isotopes present only in trace amounts. There are four times as many ^{136}Xe atoms in the detector than ^{134}Xe atoms, so if the dominant process for the production of ^{135}Xe and ^{137}Xe in the detector is thermal neutron capture, there should be more ^{137}Xe than ^{135}Xe. However, the ^{135}Xe is spectrally degenerate with most beta decays with an end point close to 1 MeV, meaning it might be acting as a spectral surrogate for another beta emitter in or near the xenon. There is a possibility that the ^{135}Xe is produced in the recirculating gaseous xenon outside the detector, which is outside the lead shielding and the HFE, so the flux of high energy neutrons is higher than inside the detector. The half-life of ^{135}Xe is 9.1 hours whereas the half-life of^{137}Xe is only 3.8 minutes, therefore not all the ^{137}Xe created in the recirculating gas will make it back to the detector before decaying, whereas almost the ^{135}Xe will.

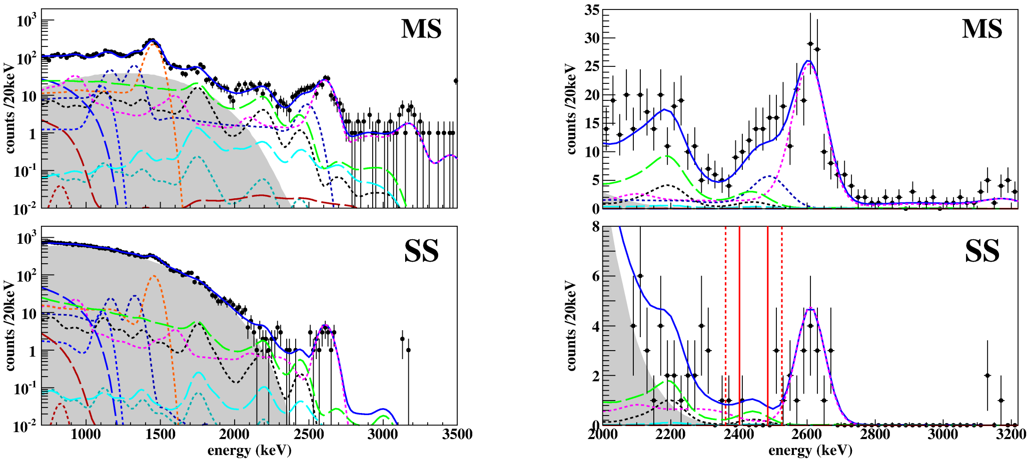

(a) (b)

Figure 7.18: (a) Fit to low background physics data to a background model plus $2\nu\beta\beta$ and $0\nu\beta\beta$. The grey area is the fit to $2\nu\beta\beta$. Other components include: ^{232}Th in the copper vessel (dotted magenta), ^{222}Rn in the air around the cryostat (dashed green), ^{238}U in the vessel (dotted black), ^{214}Bi on the cathode (dashed cyan), ^{222}Rn dissolved in the inactive xenon (dotted turquoise), ^{222}Rn in the active liquid xenon (dashed brown), ^{60}Co in the vessel (dotted blue), ^{54}Mn in the vessel (dotted brown), ^{40}K in the vessel (dotted orange), and ^{135}Xe in the active xenon (dashed blue). (b) A zoom in on the plot in (a) on the area around $Q_{\beta\beta}$ of ^{136}Xe. The solid vertical red line represents the one sigma window around $Q_{\beta\beta}$ whereas the dashed red lines show the two sigma window.

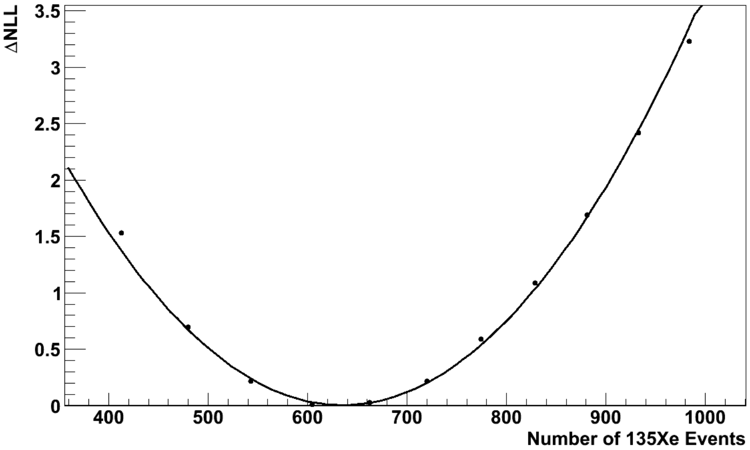

Figure 7.19: Profile likelihood of ^{135}Xe in the fit to the low background physics data with a fit to a parabola. Plot shows the difference between the negative log likelihood (NLL) from the minimum NLL versus number of ^{135}Xe events in the fit. The profile measures a total of 635 ± 135 ^{135}Xe events in the dataset, 4.7σ different from zero.

7.7.2 Systematic Uncertainties

Source Rate Discrepancy

The only systematic which is not floated in the maximum likelihood fit to the data is the 9.4% systematic uncertainty which comes from the source rate agreement discrepancy. This is taken into account after the fit as a correction to the 90% confidence level for $0\nu\beta\beta$ and as an uncertainty for $2\nu\beta\beta$ and is the largest systematic. It addresses the affects of the fiducial volume uncertainty and any unaccounted for uncertainties in the efficiency calculations. This affects the $0\nu\beta\beta$ result by 12.34%. This number is actually larger than 9.4% because the final $0\nu\beta\beta$ result is a 90% confidence limit, not a measurement.

β Scale

There is reason to believe β-like events may have a slightly different energy scale in the detector than gamma events (see section 7.4.1), a calibration parameter for electron events is determined using the fit to the $2\nu\beta\beta$-decay spectrum. All electron-like PDFs

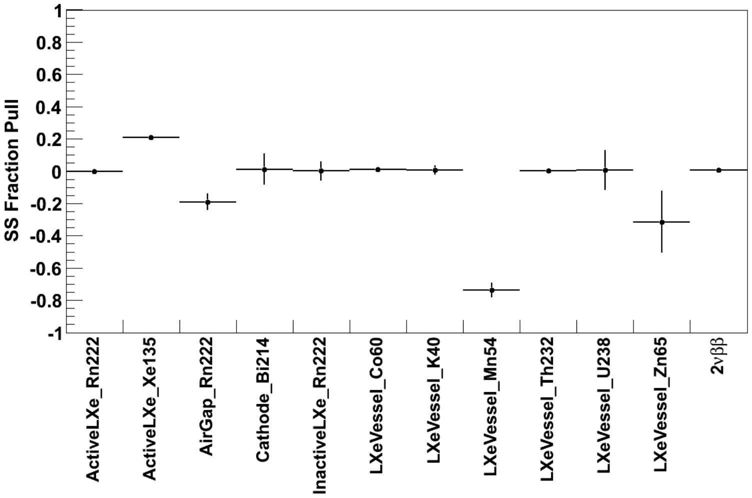

Figure 7.20: The single site fraction pulls for the various PDFs. The pull here is defined as the difference between the fit value for the single site fraction and the predicted single site fraction from Monte Carlo divided by the measured single site fraction.

have a common parameter β which is an overall multiplicative factor to the energy scale. A profile likelihood scan is performed over this parameter to determine its to be $.995 \pm .063$. This has an effect on the final $0\nu\beta\beta$ result of 9.32%.

Single Site Fraction

The single site fraction is floated in the fit for each PDF individually with a Gaussian constraint with a one sigma width of 8.5%. The mean of the Gaussian comes from the Monte Carlo prediction. This systematic affects the final $0\nu\beta\beta$ result by .93%. The pulls for the individual PDFs are shown in figure 7.20. The pull is defined as $\frac{(measured\,SS\,fraction)-(MC\,SS\,fraction)}{measured\,SS\,fraction}$. The PDF for ^{54}Mn in the copper vessel has an unusually large pull, but this PDF also fits to fewer than one total event. Almost all other pulls are within two sigma of zero.

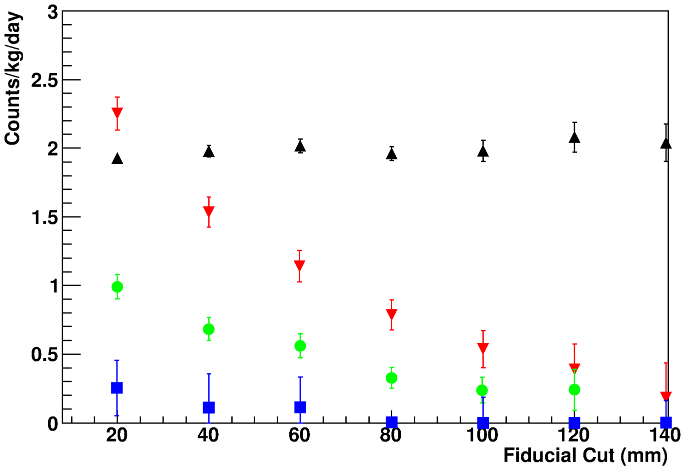

Figure 7.21: Rates versus fiducial cut coming from maximum likelihood fits to the low background physics data for various signals. Shown: $2\nu\beta\beta$ (black triangles), ^{40}K in the vessel X10 (red triangles), ^{232}Th in the vessel X10 (green circles), ^{238}U in the vessel X10 (blue squares).

Other Systematics

Other systematics, in aggregate, change the final result by less than .2%. They include uncertainties on measured rates of the backgrounds, the energy offset, and the energy resolution parameters.

7.8 Rates as a Function of Fiducial Cut

Figure 7.21 shows the rates for $2\nu\beta\beta$ and various gammas versus fiducial cut, which is the hexagonal standoff distance, so a cut of 0 mm removes none of the xenon and a cut of 183 mm removes all of the xenon. The activity for each PDF is taken from the fit to the low background data and normalized to the mass of the given fiducial cut. As expected, the rate of $2\nu\beta\beta$ is constant with choice of fiducial cut, whereas as the rates of various gamma background drops off for smaller fiducial volumes deeper inside the detector.

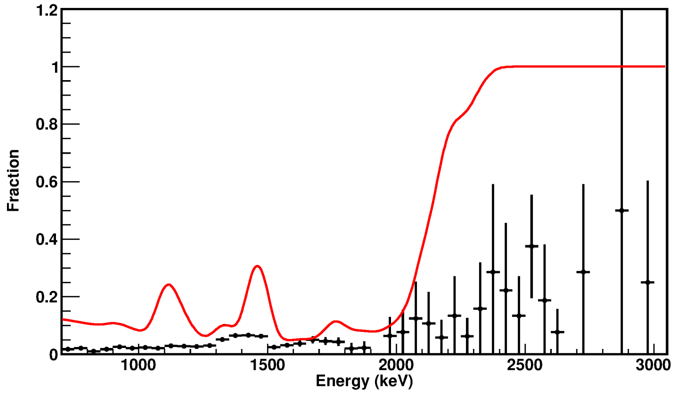

Figure 7.22: Black: faction of single site events that were moved to the multi site spectrum by the triangular reshaping algorithm in signal finding as a function of energy. Red: fraction of masked single site events which are gamma events according to the maximum likelihood fits to the low background data as a function of energy.

7.9 Effect of Triangular Reshaping on Low Background Data

Several events near the $\beta\beta$ decay end point were moved to multi site by the triangular reshaping algorithm discussed in section 5.2.3. Figure 7.22 shows the fraction of events in the low background data that were moved from single site to multiple site by the triangular reshaping algorithm as a function of energy (black). The red curve shows the fraction of single site events which are gammas according to the PDF fits. The triangular reshaping tends to move more single site events near known gamma lines to multi site than elsewhere. Near the end point energy of ^{136}Xe a very large fraction of single site events are moved to multi site and these events are likely 2615 keV ^{208}Tl full energy deposits which had their energy underestimated. The fraction of events moved from single site to multi site is higher where the gamma fraction is higher, as expected.

Table 7.5: Expected number of events from backgrounds near $Q_{\beta\beta}$ of ^{136}Xe from the fit to the low background physics data.

Background	Expected Events (1σ window)	Expected Events (2σ window)
^{222}Rn in cryostat airgap	1.9±0.2	2.9±0.3
^{238}U in the Vessel	0.9±0.2	1.3±0.3
^{232}Th in the Vessel	0.9±0.1	2.9±0.3
^{214}Bi on the cathode	0.2±0.01	0.3±0.02
All others	~0.2	~0.2
Total	4.1±0.3	7.5±0.5
Observed	1.9±0.2	2.9±0.3
Background in $kg^{-1}yr^{-1}keV^{-1}$	$(1.5\pm0.1)\times10^{-3}$	$(1.4\pm0.1)\times10^{-3}$

7.10 $0\nu\beta\beta$ Limit

From the fit, background rates of $4.1 \pm .3$ counts in the one sigma region of interested and $7.5 \pm .5$ counts in the two sigma region of interest are calculated. The measured backgrounds are summarized in table 7.5. These backgrounds rates are the expected background rates based on the total fit results, corresponding to $(1.5 \pm .1) \times 10^{-3}$ and $(1.4 \pm .1) \times 10^{-3}$ counts/kg/yr/keV respectively. The largest background comes from ^{222}Rn in the air around the cryostat. In a there is one observed single site event 1σ window around the $Q_{\beta\beta}$, and five single site events in a 2σ window. A profile likelihood determined a limit on the total number of $0\nu\beta\beta$ events of 2.83 at the 90% confidence level. This takes advantage of the fact that the observed number of events was below the calculated background, which is completely allowed statistically. Toy Monte Carlo studies determined that with the current backgrounds and a data set of this size the experiment will establish a limit as good as 2.83 events or better 6.5% of the time.

After taking into account the 9.4% systematic uncertainty from the source rate discrepancy, this corresponds to a limit on the half-life of $0\nu\beta\beta$ of ^{136}Xe of 1.59×10^{25} yr, which is more than a factor of 2.5 improvement over the previous limit from [15].

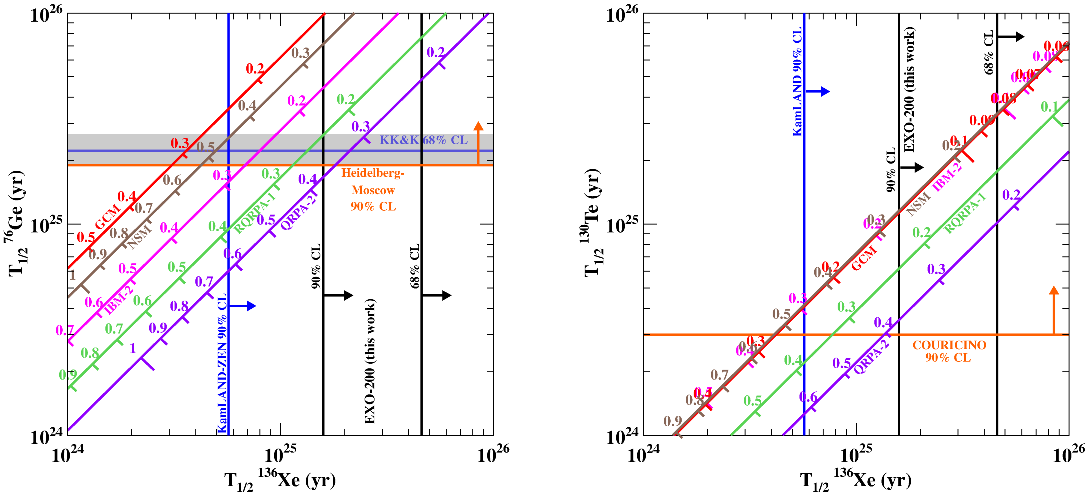

(a) (b)

Figure 7.23: Comparison of $0\nu\beta\beta$ limits for pairs of isotopes for various RQRPA (labeled RQRPA-1) [10], IBM-2 [11], GCM [12], NSM [13], and QRPA-2 [14]. Both plots show the the 90% confidence level limit from KamLAND-ZEN [15]. (a) Comparison of ^{136}Xe half-life limit on the x-axis to ^{76}Ge half-life on the y-axis, including the limit from the Heidelberg-Moscow experiment [2]. Also included is the 68% confidence interval for the claimed discovery by a subset of the Heidelberg-Moscow experiment [1]. (b) Comparison of ^{136}Xe half-life limit on the x-axis to ^{130}Te half-life on the y-axis. Included is the best limit from the COURICINO experiment [3].

7.11 Comparison of Results to Other Experiments

Figure 7.23 summarizes the comparison the EXO-200 result to the best limits in ^{136}Xe, ^{76}Ge, and ^{130}Te. The EXO-200 limit is in conflict with the claimed result from a subset of the Heidelberg-Moscow experiment at the 90% confidence level for most matrix element. Even though the current best half-life limit for ^{130}Te is much shorter than the current limit for ^{76}Ge it remains competitive with ^{76}Ge measurements because it has a favorable combination of phase space factor and nuclear matrix elements. This result places an upper limit on the Majorana mass of the electron neutrino of 140-380 meV depending on the choice of model.

Appendix A

The Compton Telescope Technique

A.1 Principle of a Compton Telescope

A large fraction of the multi-site events in the EXO-200 data are Compton scatters of gammas rays, especially while running with calibration sources. These events contain useful information because they are constrained by the kinematics of the interaction, leading to a well-known relationship between the angle of the scatter and the energy of the incoming and outgoing gamma, given by:

$$\phi = arccos\left[1 - m_e c^2 \cdot \left(\frac{1}{E_\gamma - E_1} - \frac{1}{E_1}\right)\right] \tag{A.1}$$

Where ϕ is the Compton scattering angle, m_e is the mass of the electron, c is the speed of light, E_γ is the energy of the incoming gamma, and E_1 is the energy lost in the Compton scatter, which is the energy deposited in the detector at that site. The detector measures an energy and three spatial coordinates for each charge deposit, as shown in figure A.1. For each Compton scatter event in the detector, the location of the first two scatters and the angle ϕ define a cone pointing back to the source. Taking a collection of such events and looking at the intersection of these cones, one can identify the location of the source. This is referred to as a Compton telescope or a Compton camera. This technique has been previously applied to gamma ray astrophysics, nuclear science, and medical imaging [53] [54] [55] [56].

In a low background liquid xenon experiment, one can use this to track calibration source locations and point to radiation hot spots, helping to identify sources of background. To do this imaging requires the determination of the correct sequence of the scatters within each of these events and there are different algorithms depending on the number of interaction sites in the event. For the purposes of EXO-200, only two and three site events are considered for Compton telescope analysis. The full algorithm for sequencing both two site and three site events used here is duplicated from [57]. Consider first the case of a two site event, where at very worst one can already sequence the event correctly 50% of the time. One can improve upon this by looking at the energy sharing between the two sites. If a gamma deposits a large amount of its energy on the first scatter, it is much more likely to full energy deposit on the second scatter. Thus, the algorithm assigns the higher energy to the first scatter and the lower energy of the two charge clusters to the second scatter. Previous work has shown that this criteria works well above 1 MeV and works better as the event energy increases [57].

To sequence three site events the algorithm takes advantage of the fact that the event will actually contain two Compton scatters, one at each of the first two sites. Fig. A.1 shows a diagram of a three site event. A given sequence of the three sites defines two vectors: one pointing from site one to site two and one pointing from site two to site three. If one knows that site two was a Compton scatter one could calculate the angle between these two vectors in two ways: from geometry and from equation A.1, the Compton scattering formula. For any three site event there are six possible permutations of the sites, and for each one can define a figure of merit, T, by:

$$T = \frac{cos^2\phi_{Com} - cos^2\phi_{geo}}{\sigma^2_{Com} + \sigma^2_{geo}} \tag{A.2}$$

Where ϕ_{Com} and ϕ_{geo} are the angles calculated from equation A.1 and geometrically respectively with uncertainties σ^2_{Com} and σ^2_{geo}. If the correct sequence is chosen, T should be close to zero and thus the sequence with the lowest figure of merit is selected as the correct sequence.

Once the sites are sequenced and the scattering angels calculated, one can image

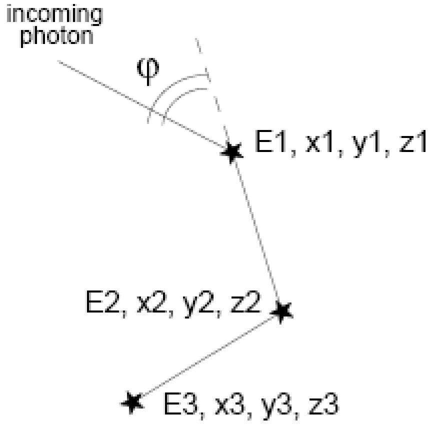

Figure A.1: A three site Compton scatter with the energy and spatial coordinates measured at each site in the detector. In the case of a three site event this defines two vectors in space, between site one and site two, the other between site two and site three. Using the spatial coordinator the angle between these two vectors can be calculated.

all possible sources by looking at the intersection of the cones from each event. The code can project the intersections of the various cones onto an arbitrary surface, though for the current analysis only a cylindrical surface is considered.

A.2 Application to Source Calibration Runs

The Compton imager was applied to identify the location of calibration sources when deployed to the various source positions. An image produced for one calibration source can be seen in figure A.2. This image shows the copper vessel unfolded. The two circles are the end caps of the vessel and the strip in the middle is a wraparound of the vessel. Thus in the circles the coordinates can be thought of as X and Y whereas in the strip the proper coordinates are θ and Z. However, here theta is not measured as an angle but as a distance traveled parallel to the source of the vessel. In this case this is an image of a source at positive Y.

The image is produced by summing all the Compton cones drawn from multi site events. Note that the parameters of these cones have uncertainties, which means the cones themselves have uncertainties. The image takes these uncertainties into account using a Monte Carlo method. For each Compton cone, many proto-cones are drawn which fall within the uncertainties of the main cone. This addition of cones is weighted and assumes Gaussian uncertainties. For this image only two site events were used as this in practice seems to produce a sharper image. Part of the issue is that the energy of the individual charge deposits must be used to when applying the sequencing algorithm for three site events, so good charge only energy resolution is necessary. As the charge only energy resolution suffers from the fact that the detector's drift field is lower than expected, the ability to sequence three site events will suffer some.

One can determine the source position by taking a one dimensional slice and fit to a one dimensional Gaussian. For example, figure A.3 shows a one dimensional slice in the Z/X direction at the intended position in y. The peak is fit to a Gaussian to determine the the Z position of the source. This is repeated for all source positions.

The list of actual positions as determined by the Compton imager is as follows:

- Positive X: Z $= -26.4 \pm 0.5 mm$
- Positive Y: Z $= 40.6 \pm 0.8 mm$
- Negative Y: Z $= -29.3 \pm 1.2 mm$
- Positive Z: X $= 24.4 \pm 0.7 mm$
- Negative Z: X $= -29.1 \pm 0.4 mm$

Of particular note, the positions suggest that the source cable is being insert 1 inch short of its intended location in the source tubing. This is consistent with the event distributions in these runs. For example, figure A.4 shows the Z-distribution for a source calibration run on the positive X side of the detector. The Compton image suggests that the source is displaced to the negative Z direction, and there are

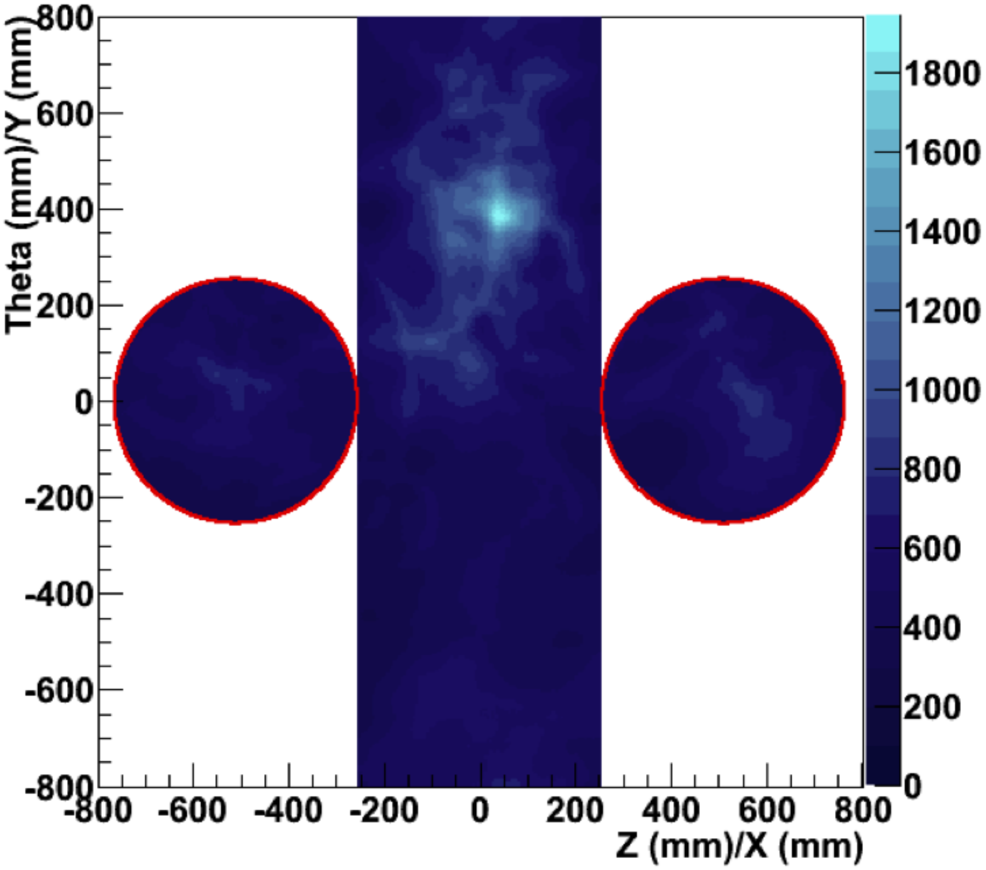

Figure A.2: Compton Telescope image of a calibration source at positive y. These images are used to identify the source locations, which is then used to generate simulated calibration runs for comparison to data.

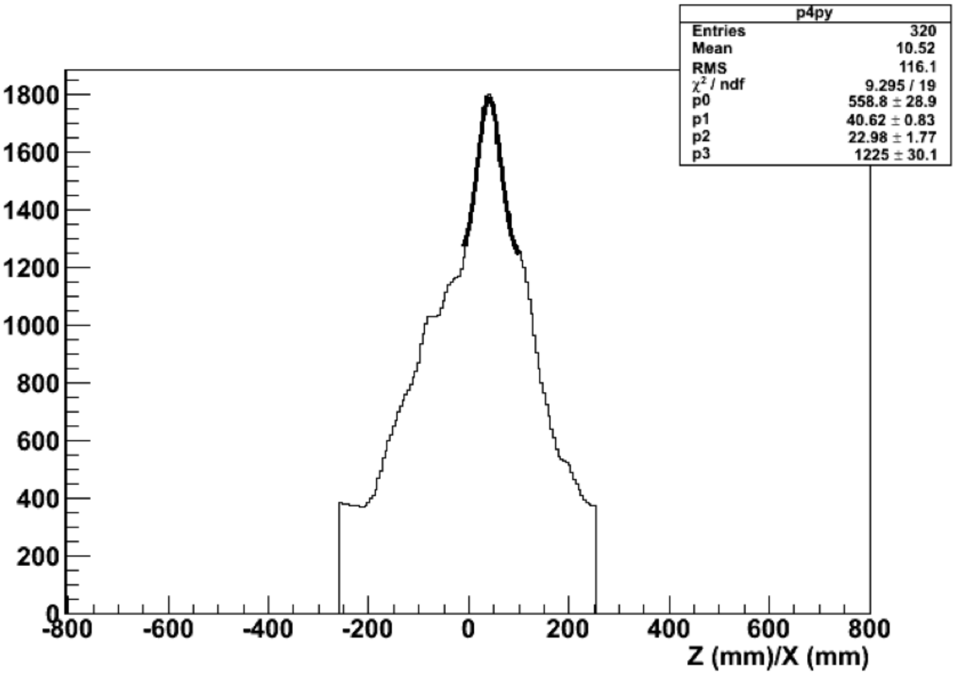

Figure A.3: A horizontal slice of the 2D histogram shown in figure A.2, which is fit to a Gaussian to determine the location of the point-like source.

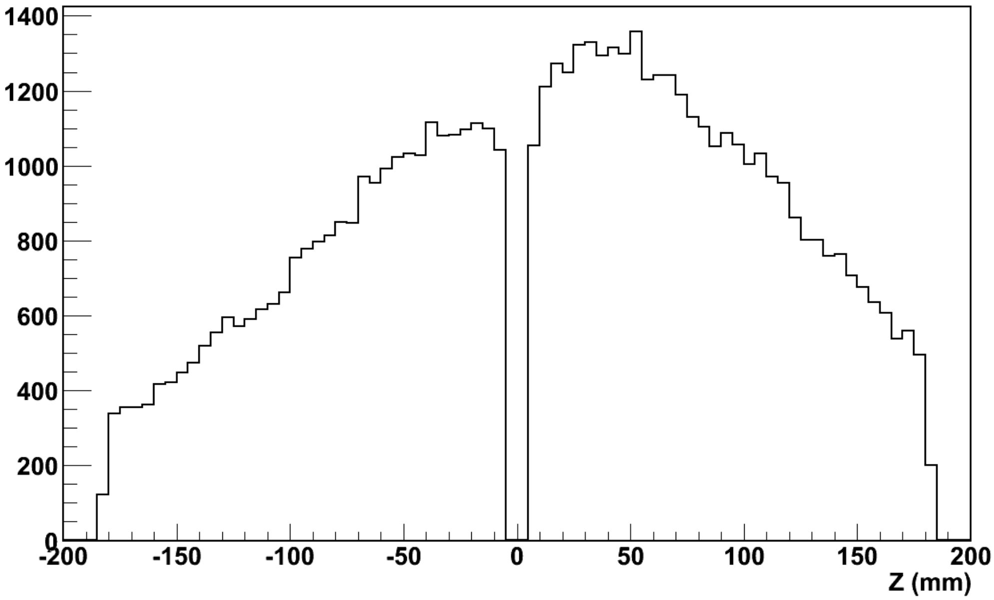

Figure A.4: Z distribution of events for a source calibration run at positive x.

more events in the negative Z half of the detector. Both pieces of evidence tell the same story.

The positions measured using the Compton telescope are used in calibration source simulations to help with source agreement studies and Monte carol verification. The Compton telescope will be applied in the future to look for radiation hot spots or point sources of radioactivity in the detector.

Bibliography

[1] H. V. Klapdor-Kleingrothaus and I. V. Krivosheina, Mod. Phys. Lett. A **21**, 1547 (2006).

[2] H. Klapdor-Kleingrothaus *et al.*, Euro. Phys. J. A **12**, 147 (2001).

[3] C. Arnaboldi *et al.*, Phys. Rev. C **78**, 035502 (2008).

[4] A. Barabash and N. C. Brudanin, V. and, Phys. Atom. Nuc. **74**, 312 (2011).

[5] S. Hannestad, Prog.Part.Nucl.Phys. **65**, 185 (2010), 1007.0658.

[6] S. Bilenky and C. Giunti, Mod.Phys.Lett. **A27**, 1230015 (2012), 1203.5250.

[7] S. R. Elliott and P. Vogel, Annu. Rev. Nucl. Part. Sci. (2002).

[8] L. S. Miller, S. Howe, and W. E. Spear, Phys. Rev. **166**, 871 (1968).

[9] C. Silva, *Study of the reflectance distributions of fluoropolymers and other rough surfaces with interest to scintillation detectors*, PhD thesis, University of Coimbra, 2010.

[10] F. Šimkovic, A. Faessler, H. Müther, V. Rodin, and M. Stauf, Phys. Rev. C **79**, 055501 (2009).

[11] J. Barea and F. Iachello, Phys. Rev. C **79**, 044301 (2009).

[12] T. R. Rodriguez and G. Martinez-Pinedo, Phys. Rev. Lett. **105**, 252503 (2010).

[13] J. MenÃľndez, A. Poves, E. Caurier, and F. Nowacki, Nuc. Phys. A **818**, 139 (2009).

[14] A. Staudt, K. Muto, and H. V. Klapdor-Kleingrothaus, Europhys. Lett. **13**, 31 (1990).

[15] KamLAND-Zen Collaboration, A. Gando *et al.*, Phys. Rev. C **85**, 045504 (2012).

[16] KamLAND Collaboration, A. Gando *et al.*, Phys.Rev. **D83**, 052002 (2011), 1009.4771.

[17] MINOS Collaboration, P. Adamson, Phys. Rev. Lett. **101**, 131802 (2008).

[18] SNO Collaboration, B. Aharmim *et al.*, Phys. Rev. C **81**, 055504 (2010).

[19] The Super-Kamiokande Collaboration, R. Wendell *et al.*, Phys. Rev. D **81**, 092004 (2010).

[20] Daya Bay Collaboration, F. P. An *et al.*, (2012), arXiv:1203.1669v2.

[21] O. Mena, B. A. Reid, L. Verde, and R. JimÃľnez, J. Cosmol. Astropart. Phys. (2010).

[22] S. A. Thomas, F. B. Abdalla, and O. Lahav, Phys. Rev. Lett. **105**, 031301 (2010).

[23] C. Otten, E. W. Weinheimer, Rep. Prog. Phys. **71** (2008).

[24] S. M. Bilenky and S. T. Petcov, Rev. Mod. Phys. **59**, 671 (1987).

[25] F. T. Avignone, S. R. Elliott, and J. Engel, Rev. Mod. Phys. **80**, 481 (2008).

[26] L. F. Li and F. Wilczek, Phys. Rev. D **25**, 143 (1982).

[27] B. Kayser, Phys. Rev. D **26**, 1662 (1982).

[28] R. Arnold *et al.*, Phys. Rev. Lett. **95**, 182302 (2005).

[29] T. Bernatowicz *et al.*, Phys. Rev. C **47**, 806 (1993).

[30] S. Umehara *et al.*, Phys. Rev. C **78**, 058501 (2008).

[31] J. Argyriades *et al.*, Nuc. Phys. A **847**, 168 (2010).

[32] F. A. Danevich *et al.*, Phys. Rev. C **68**, 035501 (2003).

[33] C. E. Aalseth *et al.*, Mod. Phys. Lett. A **17**, 1475 (2002).

[34] F. Feruglio, A. Strumia, and F. Vissani, Nuc. Phys. B **637**, 345 (2002).

[35] Y. Zdesenko, F. Danevich, and V. Tretyak, Phys. Lett. B **546**, 206 (2002).

[36] A. Strumia and F. Vissani, Nuclear Physics B **726**, 294 (2005).

[37] H. V. Klapdor-Kleingrothaus, A. Deitz, H. L. Harney, and I. V. Krivosheina, Mod. Phys. Lett. A **16**, 2409 (2001).

[38] H. Klapdor-Kleingrothaus, I. Krivosheina, and A. Dietz, Mod. Phys. A **16**, 2409 (2002).

[39] H. Klapdor-Kleingrothaus, I. Krivosheina, A. Dietz, and O. Chkvorets, Phys. Lett. B **586**, 198 (2004).

[40] V. Cirigliano, A. Kurylov, M. J. Ramsey-Musolf, and P. Vogel, Phys. Rev. Lett. **93**, 231802 (2004).

[41] J. Schechter and J. W. F. Valle, Phys. Rev. D **25**, 2951 (1982).

[42] M. Redshaw, E. Wingfield, J. McDaniel, and E. G. Myers, Phys. Rev. Lett. **98**, 053003 (2007).

[43] D. Leonard *et al.*, Nuc. Instrum. Meth. A **591**, 490 (2008).

[44] W. Maneschg *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **593**, 448 (2008).

[45] EXO Collaboration, M. Auger *et al.*, To appear in Phys. Rev. Lett. (2012), arXiv:1205.5608.

[46] EXO Collaboration, N. Ackerman *et al.*, Phys. Rev. Lett. **107**, 212501 (2011).

[47] E. Conti *et al.*, Phys. Rev. B **68**, 054201 (2003).

[48] E. Aprile, K. L. Giboni, P. Majewski, K. Ni, and M. Yamashita, Phys. Rev. B **76**, 014115 (2007).

[49] F. Leport *et al.*, Rev. Sci. Instrum. **82**, 105114 (2011), 1104.5041.

[50] h. Ansoft.

[51] S. Agostinelli *et al.*, Nuc. Instrum. Meth. A **506**, 250 (2003).

[52] O. Bunemann, T. E. Cranshaw, and J. A. Harvey, Canadian Journal of Research **27a**, 191 (1949).

[53] E. Aprile *et al.*, Nuc. Instrum. Meth. A **593**, 414 (2008).

[54] S. Kobayashi *et al.*, Adv. Space Res. **37**, 28 (2006).

[55] M. J. Myjak and C. E. Seifert, IEEE Transactions on Nuclear Science **55**, 769 (2008).

[56] K. Giboni *et al.*, J. Instrum. **2**, P10001 (2007).

[57] A. Curioni, *Laboratory and Balloon Flight Performance of the Liquid Xenon Gamma Ray Imaging Telescope (LXeGRIT)*, PhD thesis, Columbia University, 2004.