

ABSTRACT

Title of dissertation: A Search for Neutrinoless Double-Beta Decay
with EXO-200

Simon Slutsky, Doctor of Philosophy, 2013

Dissertation directed by: Professor Carter Hall
Department of Physics

This work presents a search for neutrinoless double beta decay of ^{136}Xe using data from the EXO-200 detector collected between 2011 and 2012. Neutrinoless double beta decay ($\beta\beta 0\nu$) is a hypothetical nuclear decay possible only if the neutrino is massive and is a Majorana particle. Observation of this process would constitute a measurement of the absolute neutrino mass scale, which is known to be non-zero from neutrino oscillation experiments. EXO-200 is a liquid xenon time projection chamber located at the Waste Isolation Pilot Plant (WIPP) in Carlsbad, NM. The detector is filled with 200 kg of liquid xenon isotopically enriched to 80.6%, used as both detection medium and decay source. Spectral fits based on detailed Monte Carlo simulations are used to constrain the number of $\beta\beta 0\nu$ events in the data. The analysis finds no evidence for $\beta\beta 0\nu$ in ^{136}Xe , placing a lower limit on the half-life of $1.6 \cdot 10^{25}$ yr at 90% confidence level. This implies an upper limit on the effective Majorana neutrino mass between 0.14-0.38 eV, one of the most stringent limits ever set on $\beta\beta 0\nu$.

A Search for Neutrinoless Double-Beta Decay with EXO-200

by

Simon Slutsky

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2013

Advisory Committee:

Professor Carter Hall, Chair/Advisor

Professor Elizabeth Beise

Professor Rabindra Mohapatra

Professor Hassan Jawahery

Professor Alice Mignerey, Dean's Representative

Acknowledgments

This thesis was written with the the support of exceptional colleagues and loving family and friends. I felt privileged to work with the EXO-200 collaboration. Between time spent at WIPP and Stanford, I have learned a great deal and met many remarkable people. Special thanks to the analysis team that all worked so hard, especially Phil, Igor, Kevin, Russell, Mike, Ryan Manuel, Guillaume, and Steve. I will always appreciate Liang and Haiyan, who kindly let me share their home. Special thanks also to Michelle, Jon and everyone else who helped with the hard work on the xenon compressors.

I thank my advisor, Carter Hall for encouraging me to fully explore the opportunities of working with EXO-200. My life would be much less rich had he not suggested I temporarily move to WIPP. I also thank him for his continuous insights and his patient instruction, which never fail to impress me. Yung-Ruey, Clayton, Atilla, Doug, Tom, and Colleen all contributed to an enjoyable work environment at Maryland, as well.

I share the happy occasion of my graduation with my parents, who gave me every opportunity in life and always wished me the best. Finally, thanks to Adriana, who has loved me, encouraged me and reminded me to believe in myself every day.

Table of Contents

| | |
|--|----|
| List of Abbreviations | vi |
| 1 Introduction | 1 |
| 2 Neutrino mass and double-beta decay | 3 |
| 2.1 Neutrino Masses | 3 |
| 2.1.1 Tritium beta spectroscopy | 4 |
| 2.1.2 Cosmological bounds | 5 |
| 2.2 Majorana neutrinos | 6 |
| 2.3 Double beta decay | 8 |
| 2.3.1 Nuclear matrix element calculations | 11 |
| 2.3.2 Previous limits | 13 |
| 2.3.3 Neutrino mass hierarchy | 14 |
| 2.3.4 $\beta\beta$ experimental signature | 16 |
| 3 The EXO-200 detector | 19 |
| 3.1 Liquid xenon physics | 19 |
| 3.2 Installation | 21 |
| 3.3 The EXO-200 liquid xenon TPC | 22 |
| 3.3.1 Charge Collection and Electrostatics | 26 |
| 3.3.2 Scintillation Collection | 27 |
| 3.3.3 Dead region | 30 |
| 3.4 TPC vessel and xenon handling system | 31 |
| 3.5 Cryostat | 35 |
| 3.6 Shielding | 39 |
| 3.7 Muon Veto | 40 |
| 3.8 Calibration system | 41 |
| 3.9 DAQ/Electronics | 41 |
| 3.10 Material selection | 45 |
| 3.10.1 Xenon source material | 47 |
| 4 Event Reconstruction | 50 |
| 4.1 Event characteristics | 50 |
| 4.2 Pattern Recognition | 56 |
| 4.2.1 Matched filters | 56 |
| 4.2.2 U-Wire shaping time variance | 61 |
| 4.2.3 Unshaping | 61 |
| 4.3 Parameter Estimation | 63 |
| 4.4 Clustering | 64 |

| | | |
|-------|---|-----|
| 5 | Monte Carlo | 67 |
| 5.1 | EXOSim | 67 |
| 5.1.1 | Pixelation | 69 |
| 5.1.2 | $\beta\beta$ generators | 70 |
| 5.1.3 | Detector model | 74 |
| 5.2 | Digitization | 79 |
| 5.2.1 | APD digitization | 79 |
| 5.2.2 | Wire digitization | 81 |
| 5.2.3 | Difficulties using scintillation information | 87 |
| 5.3 | Unused Monte Carlo capabilities | 87 |
| 5.4 | Notes on PDF generation | 88 |
| 6 | Calibration | 89 |
| 6.1 | Calibration Runs | 90 |
| 6.2 | U-Wire response | 90 |
| 6.3 | Electron lifetime | 95 |
| 6.3.1 | Peak method | 97 |
| 6.3.2 | Resolution method | 99 |
| 6.3.3 | Electron lifetime during the physics run | 103 |
| 6.4 | Shielding grid inefficiency | 109 |
| 6.5 | Charge energy scale | 117 |
| 6.5.1 | Potential β - γ energy scale discrepancy | 117 |
| 6.6 | Scintillation spatial dependence | 119 |
| 6.7 | Summary of corrections | 122 |
| 6.8 | Combined charge and scintillation calibration | 126 |
| 6.8.1 | Rotation parameter | 127 |
| 6.8.2 | Initial rotated energy scale calibration | 130 |
| 6.8.3 | Rotated energy resolution | 130 |
| 6.8.4 | Energy calibration fitting bias | 133 |
| 6.8.5 | Final rotated energy scale calibration | 133 |
| 6.8.6 | Effect of bias on β energy scale | 135 |
| 7 | Fitting | 139 |
| 7.1 | Culling of charge clusters | 139 |
| 7.2 | Event selection | 141 |
| 7.2.1 | Low-background data set | 141 |
| 7.2.2 | Data masking | 142 |
| 7.2.3 | Fiducial cut | 142 |
| 7.2.4 | Muons tagged by veto panel | 143 |
| 7.2.5 | Muons tagged by TPC | 144 |
| 7.2.6 | Correlated noise | 145 |
| 7.2.7 | Scintillation/ionization ratio | 145 |
| 7.2.8 | Selection results | 147 |
| 7.3 | Efficiencies | 149 |
| 7.4 | PDF generation | 151 |

| | | |
|-------|---|-----|
| 7.4.1 | Applying energy resolution | 151 |
| 7.4.2 | Applying energy scale | 154 |
| 7.5 | Monte Carlo validation | 155 |
| 7.5.1 | Spectral Shape | 155 |
| 7.5.2 | Multiplicity | 156 |
| 7.5.3 | Source rate | 157 |
| 7.6 | Background model | 158 |
| 7.6.1 | Radioactive impurities in detector components | 159 |
| 7.6.2 | Cosmogenic Activation | 160 |
| 7.6.3 | Backgrounds external to the cryostat | 160 |
| 7.6.4 | Dissolved Rn | 160 |
| 7.6.5 | ^{85}Kr | 162 |
| 7.6.6 | Constraints | 162 |
| 7.7 | Maximum likelihood fit to low background data | 163 |
| 7.7.1 | β energy scale determination | 164 |
| 7.7.2 | Fit results | 164 |
| 7.7.3 | Fitting systematics | 167 |
| 7.8 | $\beta\beta 0\nu$ half-life limit | 168 |
| 7.9 | Systematic uncertainties | 174 |
| 7.9.1 | Fiducial volume/source agreement | 174 |
| 7.9.2 | β scale systematic | 175 |
| 7.9.3 | Systematics with floating fit parameters | 181 |
| 7.10 | Comparison with previous results | 184 |
| 7.11 | Prospects for $\beta\beta 0\nu$ searches | 184 |
| 8 | Conclusions | 189 |
| | Bibliography | 190 |

List of Abbreviations

| | |
|-------------------|--|
| APD | Avalanche Photodiode |
| $\beta\beta$ | Double Beta Decay |
| $\beta\beta 0\nu$ | Neutrinoless Double Beta Decay |
| $\beta\beta 2\nu$ | Double Beta Decay with two neutrino emission |
| DAQ | Data Acquisition |
| erfc | Complementary Error Function |
| FEC | Front End Card |
| GD-MS | Glow-Discharge Mass Spectroscopy |
| GPM | Gas Purity Monitor |
| ICP-MS | Inductively-Coupled Plasma Mass Spectroscopy |
| LAAPD | Large Area Avalanche Photodiode |
| LXe | Liquid Xenon |
| MC | Monte Carlo |
| MS | Mass Spectroscopy |
| NAA | Neutron Activation Analysis |
| PCD | Pixelated Charge Deposit |
| PDF | Probability Distribution Function |
| ppb | parts per billion |
| scfh | standard cubic feet per hour |
| RMS | Root Mean Square |
| TPC | Time Projection Chamber |
| QE | Quantum Efficiency |
| WIPP | Waste Isolation Pilot Plant |

Chapter 1

Introduction

In this thesis, we describe a search for neutrinoless double beta decay ($\beta\beta 0\nu$) of ^{136}Xe with the EXO-200 experiment. Neutrinoless double beta decay ($\beta\beta 0\nu$) is a theorized nuclear transition with fundamental consequences. Observation of $\beta\beta 0\nu$ would show the neutrino to be the first known Majorana particle, identical with its own antiparticle. $\beta\beta 0\nu$ would also be the first ever process observed to violate conservation of total lepton number, with possible implications for the matter-antimatter asymmetry in the universe. Current limits place the half-life of $\beta\beta 0\nu$ above 10^{25} yr for some isotopes.

Observation of $\beta\beta 0\nu$ would also constitute a measurement of the neutrino mass scale. That neutrinos are massive is known from neutrino oscillation experiments. However, these only measure the mass difference between the different types of neutrinos and not the overall scale. Understanding this mass scale is crucial, as it may be related to the scale of Grand Unified Theories of physics beyond the Standard Model if the neutrino is Majorana [1].

Constraints on the mass of the neutrino are derived from several types of experiments, both laboratory-based and cosmological, placing the mass at $\lesssim 0.3 - 1$ eV. Currently, searches for $\beta\beta 0\nu$ provide leading constraints on the neutrino mass. The oscillation results suggest that it may be possible to observe $\beta\beta 0\nu$ in the near

future, and multiple experiments are currently underway or being planned to search for $\beta\beta 0\nu$ [2, 3].

EXO-200 is a prototype detector in the EXO (Enriched Xenon Observatory) program to search for $\beta\beta 0\nu$ in ^{136}Xe . The goal of EXO is to build a detector sensitive to neutrino masses of ~ 0.01 eV. EXO-200 consists of a low background, liquid xenon (LXe) time projection chamber (TPC) using 200 kg of enriched xenon as both detector medium and decay source. Significant efforts were undertaken to select low background detector materials and to improve the energy resolution. EXO-200 was the first experiment to observe the related decay, $\beta\beta 2\nu$, of ^{136}Xe [4].

The detector is described in Chapter 3. Event reconstruction and parameter estimation is discussed in Chapter 4. A detailed understanding of signal and background spectral shapes is required in order to extract the rate of $\beta\beta 0\nu$, which is achieved through comparison with Monte Carlo simulations; the latter are discussed in Chapter 5. Detector calibration, including energy scale and resolution characterization, is the topic of Chapter 6. In Chapter 7, we describe the generation of probability distribution functions (PDFs) from simulations, and give the results of maximum likelihoods fits of those PDFs to low-background data. Chapter 8 summarizes the ultimate constraint on $\langle m_{\beta\beta} \rangle$ and compares to others.

Chapter 2

Neutrino mass and double-beta decay

2.1 Neutrino Masses

From experiments searching for neutrino oscillations, it is known that neutrinos are characterized by at least three mass eigenstates, ν_1 , ν_2 and ν_3 , which are distinct from the flavor eigenstates ν_e , ν_μ and ν_τ . The squared mass differences are found to be $\Delta m_{21}^2 = (7.41_{-0.19}^{+0.21}) \cdot 10^{-5}$ eV and $\Delta m_{32}^2 = (2.32_{-0.08}^{+0.12}) \cdot 10^{-3}$ eV [5, 6]. Each flavor eigenstate can be related to a mass eigenstate by a complex-valued mixing matrix, $U_{\alpha j}$:

$$\nu_\alpha = \sum_{j=1}^3 U_{ej} m_j, \quad (2.1)$$

where α indicates one of the flavor states. The matrix can be expressed in terms of three mixing angles, θ_{ij} and three phases, α , β and δ :

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

where $c_{ij} = \cos(\theta_{ij})$ and $s_{ij} = \sin(\theta_{ij})$. The first three matrices each correspond to a pair of mass eigenstates; the fourth contains phases which are only present if the neutrino is a ‘‘Majorana’’ particle, discussed further in Section 2.2. Each mixing

angle has been measured and is non-zero [7, 8, 9], although the phases have not been measured.

The mass differences cannot by themselves determine the absolute neutrino mass scale, leaving several allowed possibilities for the mass hierarchy:

$$\begin{aligned}
 m_1 \sim m_2 \ll m_3 & \quad \text{“normal”} \\
 m_3 \ll m_1 \sim m_2 & \quad \text{“inverted”} \\
 m_1 \sim m_2 \sim m_3 & \quad \text{“degenerate”}
 \end{aligned}
 \tag{2.3}$$

Several experiments probe the absolute scale, including β -decay endpoint measurements, cosmological constraints, and neutrinoless double beta decay.

2.1.1 Tritium beta spectroscopy

A non-zero neutrino mass reduces the endpoint energy of the β -decay spectrum, purely due to kinematics. Studies of the highest energy β s produced in tritium decay are therefore sensitive to the electron neutrino mass squared:

$$m_{\nu_e}^2 = \sum_{i=1}^3 |U_{ei}| m_i^2.
 \tag{2.4}$$

An upper limit of $m_{\nu_e} < 2.3$ eV at 95% confidence level has been obtained by the Mainz experiment, with a similar result from the Troitsk experiment [10, 11]. The KATRIN experiment plans to have a sensitivity of 0.2 eV [12].

2.1.2 Cosmological bounds

Massive neutrinos contribute to cosmological structure formation. Measurements of baryon acoustic oscillations, cosmic microwave background radiation, weak gravitational lensing of galaxies, and supernovae can be combined to yield a limit on the sum of the three neutrino masses [13]

$$\sum_{i=1}^3 m_i < 0.48 \text{ eV} \quad \text{at 95\% C.L.} \quad (2.5)$$

Assuming the degenerate scenario, where all three neutrinos have the same mass, this would yield the constraint $m_i < 0.16 \text{ eV}$. This makes it a tighter constraint than that from tritium decay, although it is dependent on cosmological models.

A more stringent limit on the sum of the masses $\sum m_i < 0.17 \text{ eV}$ was obtained by combining quasar Lyman- α lines with SDSS measurements in [14]. More recent analysis has shown that this combination may not be valid due to differing systematics, and that a more appropriate limit of $\sum m_i < 0.90 \text{ eV}$ from SDSS alone should be used. [15]. Thorough reviews of cosmological neutrino mass bounds can be found in [16, 17].

2.2 Majorana neutrinos

All the massive fermions in the Standard Model are represented by four-component spinors, Ψ , which satisfy the Dirac equation,

$$\begin{aligned}i(\hat{\sigma}^\mu \partial_\mu)\Psi_R - m_D\Psi_L &= 0 \\i(\sigma^\mu \partial_\mu)\Psi_L - m_D\Psi_R &= 0,\end{aligned}\tag{2.6}$$

where Ψ_L and Ψ_R are the left and right-handed components of the spinor, m_D is the Dirac mass, and σ^μ are Pauli matrices. Majorana discovered that for an electrically neutral particle, the four-component spinor could be expressed in terms of a single independent two-component spinor, ψ [18]. In this expression, the right- and left-handed fields, ψ_R and ψ_L independently obey the Majorana equation,

$$\begin{aligned}i(\hat{\sigma}^\mu \partial_\mu)\psi_R - m_R\epsilon\psi_R^* &= 0 \\i(\sigma^\mu \partial_\mu)\psi_L - m_L\epsilon\psi_L^* &= 0,\end{aligned}\tag{2.7}$$

where $\epsilon = i\sigma_y$ and m_R and m_L are the left and right-handed masses. When the full 4-component spinor Ψ is written in terms of ψ , it has the unique property that it is an eigenstate of the CP transformation; the particle and anti-particle are identical. Such a particle is called *Majorana*.

A neutral fermion can have both types of mass terms. The full Lagrangian

mass term, \mathcal{L} , can be written as the sum

$$\begin{aligned}\mathcal{L} &= \mathcal{L}_{\mathcal{D}} + \mathcal{L}_{\mathcal{M}} \\ &= -m_D(\overline{\Psi}_L\Psi_R + H.c.) - \frac{m_L}{2}[(\overline{\Psi}_L^c)\Psi_L + H.c] - \frac{m_R}{2}[(\overline{\Psi}_R^c)\Psi_R + H.c],\end{aligned}\tag{2.8}$$

where $\mathcal{L}_{\mathcal{D}}$ and $\mathcal{L}_{\mathcal{M}}$ are the Dirac and Majorana mass terms, respectively.

We can see the consequences of a Majorana mass term on the neutrino mass scale if we take Ψ to represent a neutrino and define a combined spinor,

$$n_L = \begin{bmatrix} \Psi_L \\ (\Psi_R)^c \end{bmatrix}.\tag{2.9}$$

We can then rewrite Equation 2.8 as [3]

$$\mathcal{L} = -\frac{1}{2}(\overline{n}_L^c)\mathcal{M}n_L + H.c.,\tag{2.10}$$

where the mass matrix, \mathcal{M} is

$$\mathcal{M} = \begin{pmatrix} m_L & m_D \\ m_D & m_R \end{pmatrix}.\tag{2.11}$$

In some theories with a large mass scale [19, 20], a heavy, unobserved right-handed neutrino is introduced with mass at that scale, m_R . This leads to the relation $m_R \gg m_D \gg m_L$. Diagonalizing \mathcal{M} yields the mass eigenvalues $m_{light} \approx -\frac{m_D^2}{m_R}$ and $m_{heavy} \approx m_R$, showing that m_{light} is suppressed by a factor of $\frac{1}{m_R}$. This explanation

of the lightness of the observed neutrino masses is referred to as the “see-saw” mechanism.

Neutrinos, being electrically neutral, are candidate Majorana particles. The existence of neutrino mass opens the possibility of experimentally distinguishing a Dirac neutrino from a Majorana neutrino. For massless particles, there is no distinction, as can be seen in the limit of 0 mass in Equations 2.6 and 2.7. Any measurable differences that do exist must then be suppressed by the (unknown) very small mass.

2.3 Double beta decay

Double beta decay ($\beta\beta$) is a second-order, weak process in which a nucleus with charge Z and mass number A decays into one with charge $Z + 2$ and the same mass number, A , accompanied by the emission of two electrons. Depending on the decay mode, other particles are produced. We will focus on the two neutrino mode ($\beta\beta 2\nu$) and the neutrinoless mode ($\beta\beta 0\nu$).

The two-neutrino mode is predicted by the Standard Model and proceeds as

$$(Z, A) \rightarrow (Z + 2, A) + 2e^- + 2\bar{\nu}_e, \quad (2.12)$$

shown schematically in Figure 2.2a. The decay has been observed in twelve nuclei¹ with half-lives ranging from $7 \cdot 10^{18}$ for ^{100}Mo to $1.9 \cdot 10^{24}$ yr for ^{128}Te [23].

Half-lives for $\beta\beta$ are much longer than for single β -decay, meaning the latter

¹ ^{148}Ca , ^{76}Ge , ^{82}Se , ^{96}Zr , ^{100}Mo , ^{116}Cd , ^{128}Te , ^{130}Te , ^{136}Xe ([22]), ^{150}Nd , ^{238}U , and ^{130}Ba

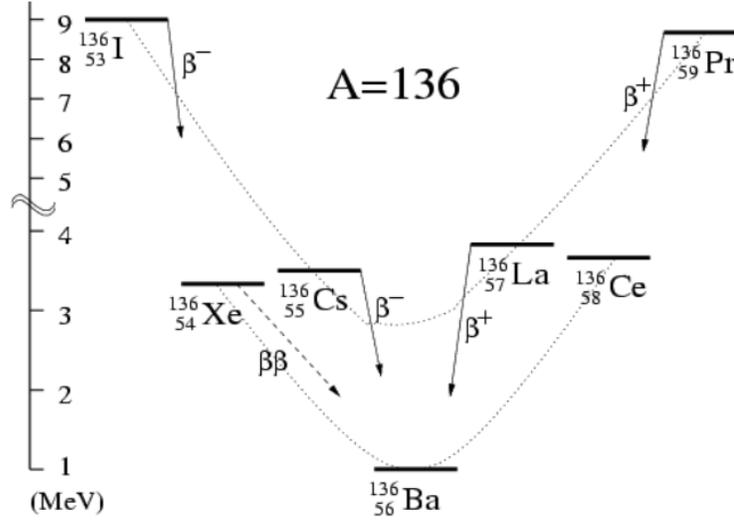


Figure 2.1: Nuclear binding energy level diagram for mass number 136, illustrating the energetic suppression of ^{136}Xe to the heavier ^{136}Cs . Figure from the Table of Isotopes [21].

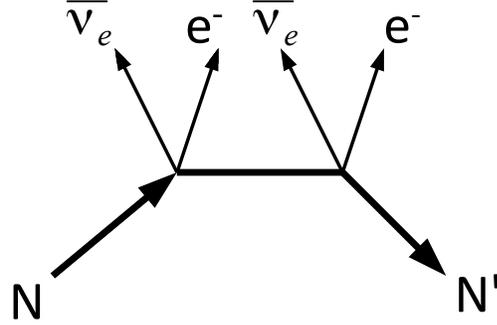
must be suppressed if $\beta\beta$ is to be observed. This condition is satisfied for several even-even nuclei. In this case, the potential daughter nucleus for β decay is more massive than the parent due to nuclear pairing effects, so β decay is energetically forbidden. This situation is illustrated in Figure 2.1, the binding energy scheme for ^{136}Xe . The β daughter, ^{136}Cs , is heavier than ^{136}Xe , while $\beta\beta$ to the lighter ^{136}Ba is allowed.

In the neutrinoless mode, the two electrons are the only products:

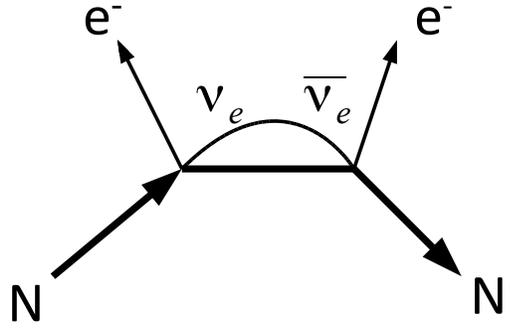
$$(Z, A) \rightarrow (Z + 2, A) + 2e^{-}. \quad (2.13)$$

The neutrinos are absorbed in an internal process, as shown in the schematic in Figure 2.2b. $\beta\beta 0\nu$ violates the conservation of total lepton number, and hence is forbidden in the Standard Model. Furthermore, $\beta\beta 0\nu$ is possible only if the neutrino

is a Majorana particle, and conversely any theory in which the neutrino is Majorana will predict the occurrence of $\beta\beta 0\nu$ at some rate. Thus, observing $\beta\beta 0\nu$ would have profound implications.



(a)



(b)

Figure 2.2: Schematics of $\beta\beta 2\nu$ (a) and $\beta\beta 0\nu$ (b). N and N' are the nuclei (Z, A) and $(Z + 2, A)$.

For $\beta\beta 0\nu$ mediated by a massive neutrino, we can express the decay rate in terms of an effective, or “Majorana”, neutrino mass,

$$\langle m_{\beta\beta} \rangle = \left| \sum_k m_k U_{ek}^2 \right|. \quad (2.14)$$

All three mass eigenstates appear by virtue of coupling to the electron neutrino

through the mixing matrix U . The half-life is then

$$[T_{1/2}]^{-1} = G_{0\nu}(Q, Z)|M_{0\nu}|^2\langle m_{\beta\beta}\rangle. \quad (2.15)$$

$G_{0\nu}(Q, Z)$ is an exactly calculable phase space factor depending on the initial nuclear charge and mass difference between the initial and final nuclei, Q . The latter quantity is also known as the ‘‘Q-value’’. Values used to compute $G_{0\nu}$ can be found in Table 6 of [24]. For the transition of ^{136}Xe to ^{136}Ba , $G_{0\nu}$ is $5.5 \cdot 10^{24} \text{ y}\cdot\text{eV}^2$ [25]. $M_{0\nu}$ is the transition matrix element between the initial and final nuclei; the theoretical value may be uncertain to a factor of 2 or 3, as different models lead to different values. Accurate matrix element calculations are crucial for inferring the effective neutrino mass from a measured half-life.

2.3.1 Nuclear matrix element calculations

Calculating the transition matrix elements, $M_{0\nu}$, is challenging, since each candidate isotope contains a large number of nucleons. A number of different approximations can be made to render the calculations tractable, the choice of which leads to one of several models. These include the generating coordinate method (GCM) [26], interacting shell model (ISM) [27], quasiparticle random phase approximation (QRPA) [28], and renormalized quasiparticle random phase approximation (RQRPA) [29]. Values of $M_{0\nu}$ for ^{136}Xe are given in Table 2.1.

The properties of these models are discussed in [25, 30, 3]. Generally, these models employ mean field methods, taking a single potential to represent the average

Table 2.1: Values of the nuclear transition matrix element, $M_{0\nu}$, for ^{136}Xe , as calculated using various nuclear models. The matrix elements are dimensionless. Where a range is given, different methods of accounting for short-range interactions yield results spanning that range. Acronyms are defined in the text.

| Technique | $M_{0\nu}$ | Reference |
|-----------|------------|-----------|
| GCM | 4.20 | [26] |
| ISM | 2.12-2.81 | [27] |
| QRPA | 1.80-3.11 | [29] |
| RQRPA | 1.67-2.89 | [29] |

influence of all the nucleons on a single bound nucleon. The single nucleons are taken to be in eigenstates of a harmonic oscillator potential. A number of nucleons occupy core states, which are fully occupied in the initial and final states of the decay. States outside this core, in “valence shells”, may participate in the transition. Two-nucleon and three-nucleon correlated interactions are added to the mean field potential and tuned to match experimental results. Models can differ in the choice of which states belong to the valence space and the method for incorporating the correlations.

QRPA/RQRPA The quasiparticle random phase approximation (QRPA) and renormalized quasiparticle random phase approximation (RQRPA) use collective state methods to calculate the single-nucleon energies. In QRPA, nucleon-nucleon interactions are treated using the BCS approximation, which treats the nucleons as distant correlated pairs. The expectation value of the BCS ground state is used to approximate the nucleon states. RQRPA similarly deals with distant nucleon pairs, but uses the expectation value of the ground states found using QRPA. In both cases, the interactions strengths are scaled multiple times by factors of ~ 1 to match known nuclear mass differences and β^+ and $\beta\beta 2\nu$ rates. A strength of the theory is

that a large number of valence nucleons from multiple shells can be used. However, since $\beta\beta 2\nu$ rates are used to tune the theory, it loses some predictive power.

ISM In the interacting shell model (ISM), nucleon-nucleon interactions are taken as the residual forces after subtracting off the N-body mean field. The interactions are essentially confined to the outermost shell, so only a small number of valence nucleons can be included, Some effects from marginal nucleons will be omitted, with possibly effects up to 25% on the resulting matrix $M_{0\nu}$ [3]. However, the small number means it is possible to do calculations with realistic nucleon-nucleon forces, which are renormalized to match measured nuclear properties.

GCM When using the generating coordinate method (GCM), the initial and final nuclear states are expressed in terms of axially symmetric states with differing quadrupole moments. Deformed nuclei are therefore taken into account intrinsically. The valence space is quite large (~ 10 shells). However, in this formalism it is not possible to calculate the two-neutrino matrix element without further approximations.

2.3.2 Previous limits

Limits on the half-life of $\beta\beta 0\nu$ ($T_{1/2}$) have been obtained for multiple isotopes; leading limits are shown in Table 2.2. The largest lower limit comes from the Heidelberg-Moscow collaboration [31], with $T_{1/2} > 1.9 \cdot 10^{25}$ yr at 90% confidence level. This was achieved using five high-purity enriched germanium detectors with

an exposure of 47.7 kg yr. The limit on the Majorana mass is also the best among isotopes other than ^{136}Xe , at $\langle m_{\beta\beta} \rangle < 0.350$ eV using the QRPA matrix element [31]. In ^{136}Xe , the leading limit prior to the data taken in this work was achieved by the Kamland-Zen collaboration using liquid scintillator doped with enriched xenon; the limit was placed at $T_{1/2} > 5.7 \cdot 10^{24}$ yr, corresponding to $\langle m_{\beta\beta} \rangle < 0.3 - 0.6$ eV depending on the matrix element.

A part of the Heidelberg-Moscow collaboration [32] has claimed an observation of $\beta\beta 0\nu$ of ^{76}Ge , with a half-life of $T_{1/2} = (2.22_{-0.31}^{+0.44}) \cdot 10^{25}$ yr at $\sim 4 \sigma$ significance. The spectrum in the analysis window near the Q-value is shown in Figure 2.3, depicting the putative $\beta\beta 0\nu$ peak at 2039 keV and several background peaks from ^{214}Bi γ s. The result is not universally accepted, as not all background peaks are identified, and independent analysis has questioned the significance of the peaks. The authors of [33] determine the rate of ^{214}Bi decay using much stronger lines outside the analysis window of [34], and find this to be too low to explain the rates inside the window. In [35], the peaks' significance is shown to depend strongly on the size of the analysis window, becoming unimportant when the size is increased.

2.3.3 Neutrino mass hierarchy

The effective mass $\langle m_{\beta\beta} \rangle$ can be combined with the mass differences measured by neutrino oscillations to provide insight into the neutrino mass hierarchy. In isolation, $\langle m_{\beta\beta} \rangle$ cannot fix the mass of any single eigenstate, due to the unknown phases in the mixing matrix, $U_{\alpha j}$. Figure 2.4 shows the allowed values of $\langle m_{\beta\beta} \rangle$,

Table 2.2: Lower limits on the $\beta\beta 0\nu$ half-life ($T_{1/2}$), prior to the data in this work. Limits are given at 90% confidence level.

| Isotope | Technique | $T_{1/2}$ (yr) | Reference |
|-------------------|--|-----------------------|-----------|
| ^{48}Ca | CaF ₂ scintillating crystal | $> 5.8 \cdot 10^{22}$ | [36] |
| ^{76}Ge | High purity Ge detector | $> 1.9 \cdot 10^{25}$ | [31] |
| ^{82}Se | Thin-foil tracking calorimeter | $> 3.6 \cdot 10^{23}$ | [37] |
| ^{96}Zr | Thin-foil tracking calorimeter | $> 9.2 \cdot 10^{21}$ | [38] |
| ^{100}Mo | Thin-foil tracking calorimeter | $> 1.1 \cdot 10^{24}$ | [37] |
| ^{116}Cd | enriched $^{116}\text{CdWO}_4$ crystal scintillators | $> 1.7 \cdot 10^{23}$ | [39] |
| ^{130}Te | TeO ₂ bolometers | $> 3.0 \cdot 10^{24}$ | [40] |
| ^{136}Xe | Enriched Xe doped scintillator | $> 5.7 \cdot 10^{24}$ | [41] |
| ^{150}Nd | Thin-foil tracking calorimeter | $> 1.8 \cdot 10^{22}$ | [42] |

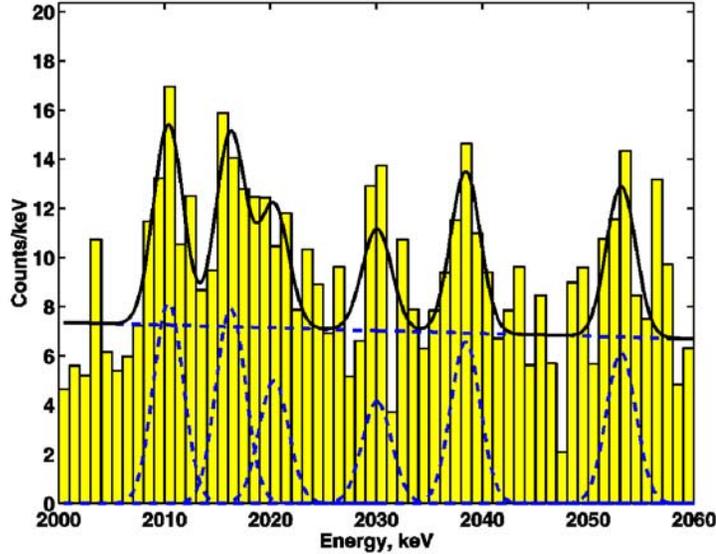


Figure 2.3: Energy spectrum leading to a claimed observation of $\beta\beta 0\nu$ in ^{76}Ge , from [34]. A fit is made to a linear background plus six gaussians, corresponding to four known γ lines from ^{214}Bi at 2011, 2017, 2022 and 2053, the $\beta\beta 0\nu$ Q-value of 2039 keV, and an unidentified line at 2030 keV. The data were obtained with an exposure of 71.7 kg·yr.

using only the measured mixing angles as input. $\langle m_{\beta\beta} \rangle$ is given in terms of the lightest neutrino mass, $m_{min} = m_1$ (normal hierarchy) or m_3 (inverted hierarchy). In the case of the (non-degenerate) inverted hierarchy, bounds can be set of 10^{-2} eV $\lesssim \langle m_{\beta\beta} \rangle \lesssim 5 \cdot 10^{-2}$ eV can be set in the case of the inverted hierarchy. In the normal hierarchy, the mass scale must be below $5 \cdot 10^{-3}$ eV if m_{min} is small ($< 10^{-3}$ eV); no lower bound can be set, as cancellations between the complex phases of $U_{\alpha j}$ can reduce the value $\langle m_{\beta\beta} \rangle$ to below m_{min} .

2.3.4 $\beta\beta$ experimental signature

The spectra of $\beta\beta 0\nu$ and $\beta\beta 2\nu$ are illustrated in Figure 2.5. A double beta decay signal appears as the sum energy of the two β s produced. In $\beta\beta 2\nu$, the energy is shared with the neutrinos, leading to a continuous distribution. For $\beta\beta 0\nu$, the only products are the two β s, so the spectrum is a peak at the Q-value of the decay. Due to the detector energy resolution, this peak is widened, and overlaps with the $\beta\beta 2\nu$ spectrum. Thus, $\beta\beta 2\nu$ forms a background for $\beta\beta 0\nu$, and an understanding of the spectral shape is important for separating the two. For ^{136}Xe , the Q-value is 2.458 MeV.

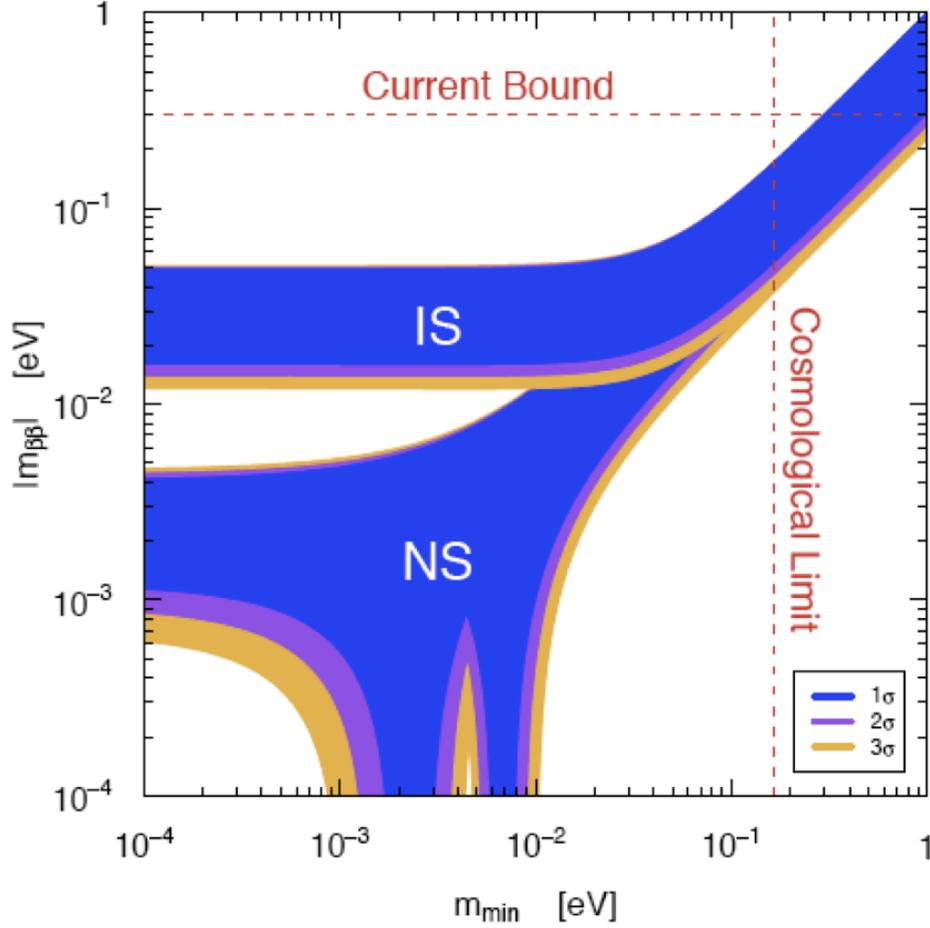


Figure 2.4: The allowed effective neutrino mass $\langle m_{\beta\beta} \rangle$ as a function of the lightest neutrino mass, m_{min} . The allowed range differs for the normal hierarchy (NS) and inverted hierarchy (IS). Figure from [1]. Also shown are the upper bounds on $\langle m_{\beta\beta} \rangle$ from Heidelberg-Moscow [31] and on m_{min} from cosmological constraints, assuming $m_{min} = \frac{1}{3} \sum_i m_i$, i.e., the degenerate mass hierarchy. A clear lowest $\langle m_{\beta\beta} \rangle$ is possible in the inverted hierarchy, but in the normal hierarchy, cancellations allow $\langle m_{\beta\beta} \rangle$ to be below the neutrino mass.

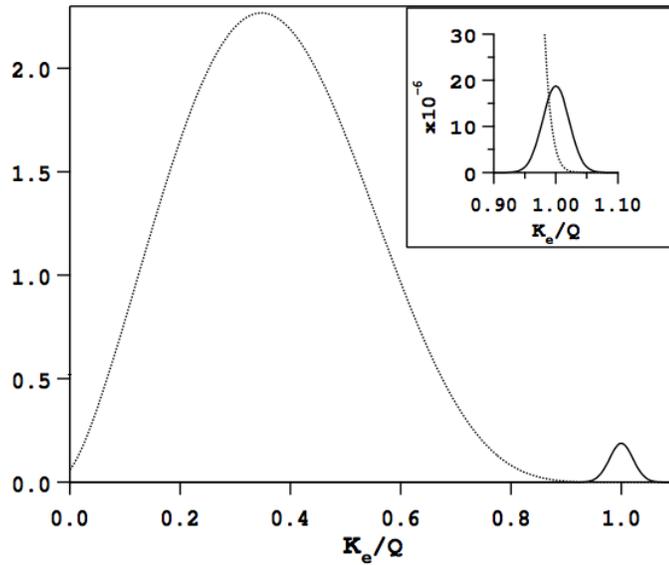


Figure 2.5: The energy spectra of $\beta\beta 0\nu$ and $\beta\beta 2\nu$ is illustrated, with the energy variable normalized to the Q-value of the decay. The spectra are normalized so that the number of $\beta\beta 0\nu$ events is 10^{-2} (10^{-6} in the inset) times the number of $\beta\beta 2\nu$ events. A 5% gaussian energy resolution has been included to illustrate detector effects. The inset shows that this causes $\beta\beta 2\nu$ to become a background for $\beta\beta 0\nu$. From [43].

Chapter 3

The EXO-200 detector

In this chapter, we describe a detector designed to measure the energy of $\beta\beta$ signals. The detector uses liquid xenon as both an ionizable medium and as the source of the decay being monitored. Liquid xenon produces scintillation photons when ionized, providing a second channel for energy measurement along with the ionization electrons. The detector records both these charge and light signals, leading to accurate reconstruction of the position and time of the energy deposit. In addition, it is seen that there is an anti-correlation between those two types of signals which can be exploited to improve the energy resolution of the detector.

In order to achieve the desired sensitivity, radioactive impurities near the detector must be minimized. All materials comprising and surrounding the detector have been carefully selected for low radioactive content. In addition, components are designed to minimize the amount of material used.

3.1 Liquid xenon physics

High-energy radiation impinging on liquid xenon produces a cascade of electrons, resulting in pairs of Xe^+ ions and electrons. In addition, some neutral Xe is excited into an excimer state, Xe^* . By applying an electric field across the xenon, some of these electrons can be drifted to an electrode and collected. Electrons that

are not collected recombine with Xe^+ ions, producing scintillation photons and some Xe^* . Relaxation of the Xe^* , whether produced by the initial radiation or during recombination, yields further scintillation. The total scintillation spectrum is peaked in the VUV at 178 nm, with a width of 14 nm [44]. The amount of scintillation produced depends on the amount of recombination, which in turn depends strongly on the applied electric field and the ionization density of the deposit. In particular, alpha particles ionize small regions as compared to betas and gammas, leading to higher light yields and allowing for particle discrimination [45]. Both the number of electrons and number of photons produced is proportional to the energy of the incident radiation.

Liquid xenon exhibits large event-to-event fluctuations in the amount of charge created. This can be characterized in terms of the variance, σ_e^2 ,

$$\sigma_e = \sqrt{F_e N_e}, \quad (3.1)$$

where σ_e is in units of electron charge, N_e is the number of electron-ion pairs created, and F_e is termed the Fano factor. The Fano factor describes deviations from Poisson statistics, which result from the special case $F_e = 1$. Liquid xenon has been found to have an unusually large Fano factor, greater than 20 [46], limiting the ultimate resolution that has been obtainable with detectors using this medium.

EXO has found that this enhanced fluctuation in charge production corresponds to a similarly enhanced fluctuation in scintillation, and that the two modes are strongly anti-correlated [46]. Thus, by taking an appropriate combination of

the energy measured in the charge and light channels, it is possible to achieve a higher energy resolution than is possible in either channel separately. This effect is illustrated in Figure 3.1, which shows the charge and light spectra of a ^{232}Th calibration source (2615 keV), as measured by the EXO-200 detector. Near 2615 keV on both axes, a diagonal “island” of events appears, corresponding to the full-energy γ deposits and showing the anticorrelation of light and charge.

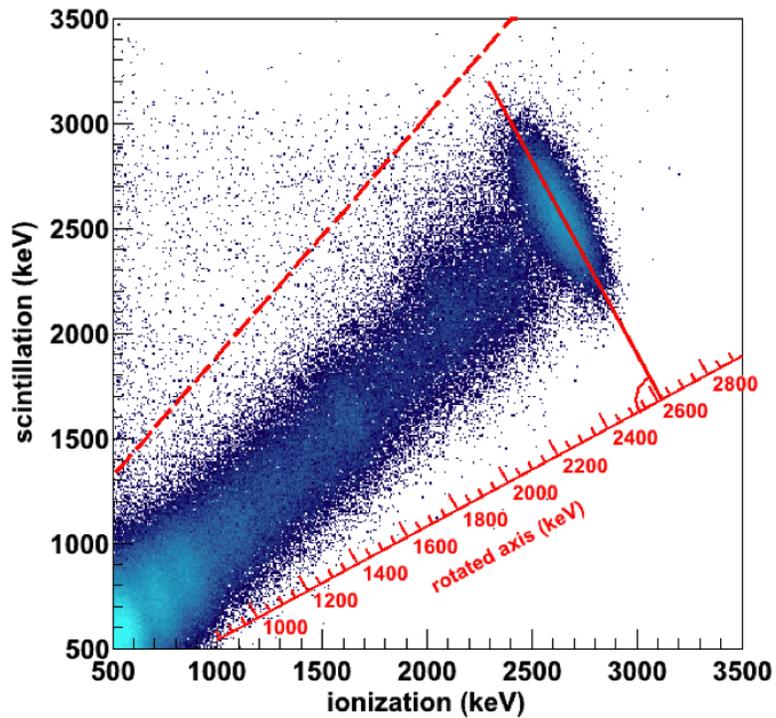


Figure 3.1: A two-dimensional spectrum of charge and light measured in the EXO-200 detector, showing the anticorrelation between the two. The data was taken with a ^{232}Th calibration source (2615 keV). The “rotated” axis indicates the linear combination that optimizes the energy resolution.

3.2 Installation

The EXO-200 detector is installed at the Waste Isolation Pilot Plant (WIPP), a transuranic waste disposal facility located in an underground salt deposit near

Carlsbad, NM, USA. The facility is 2150 ft below the earth’s surface, corresponding to an overburden of 1600 m.w.e. This provides some reduction in the number of cosmic rays (but see Section 3.7). The detector is installed in a class 1000 clean room with a dedicated air handling system and HEPA filters. The clean room is temperature- and humidity-controlled.

3.3 The EXO-200 liquid xenon TPC

EXO-200 employs the concept of a Time Projection Chamber (TPC) for measuring the energy of ionizing radiation. A schematic of a generic TPC is shown in Figure 3.2. Two electrodes, typically planar, are placed in an ionizable medium as an anode and a cathode. An electric field is established by setting a relative voltage between anode and cathode. Electrons produced by radiation are drifted to the anode and collected. A plane of shielding wires may be placed in front of the anode to screen the anode from induction of positive ions that do not recombine before the signal is measured. The TPC can be augmented, as shown, with photodetectors for recording scintillation.

EXO-200 uses two TPCs with liquid xenon as the active medium. The two TPCs are arranged back-to-back with a common central cathode, as seen schematically in Figure 3.3. As compared to two separate drift chambers, the common cathode reduces the number of wires and cables needed. This configuration also reduces drift distance, relaxing requirements on electronegative impurities as compared to a single chamber. We will refer to each TPC generically as a “half-TPC”,

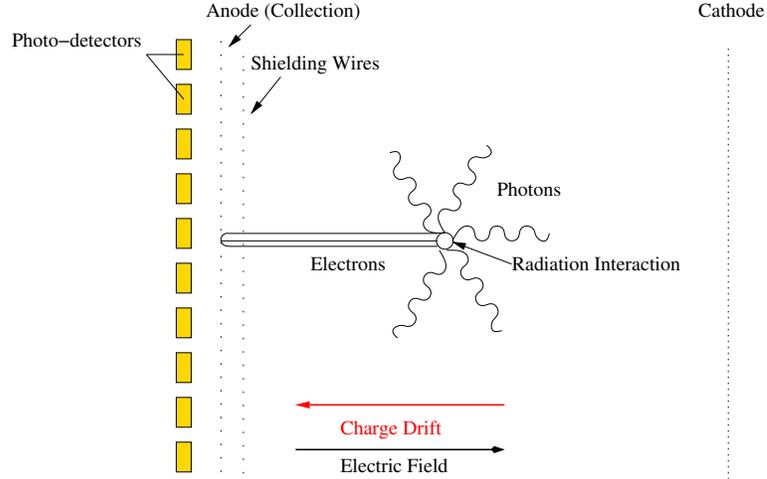


Figure 3.2: A schematic of a generic TPC. Charge and scintillation are produced during an ionizing interaction. An electric field is used to drift the charge to an anode, where it is collected and the amount measured. Typical drift velocities are $\sim 1 \text{ mm}/\mu\text{s}$. A secondary set of wires may be used to shield induction effects, as described in the text. The TPC can be supplemented by photodetectors to collect scintillation photons, which arrive promptly after the interaction.

individually as “TPC 1” and “TPC 2”, and the whole detector with its housing as “the TPC”. A cut-away view of the EXO-200 TPC and its containing vessel can be seen in Figure 3.4. A single half-TPC is pictured in Figure 3.5.

Each TPC has the geometry of a right circular cylinder; the drift region is 19.8 cm from anode to cathode and has a radius of 18.3 cm. At the end opposite the cathode of each TPC, there are two sets of wires, an anode/collection grid (“U-wires”) behind a shielding/induction grid (“V-wires”). The U-wires collect charge while the V-wires shield the U-wires from induction effects. In addition, these wires allow for position measurements of the primary charge deposit. Charges passing the V-plane induce signals on nearby wires, and charges collected a U-wire indicate the closest U-wire to the deposit location. The two sets of wires take the shape of a

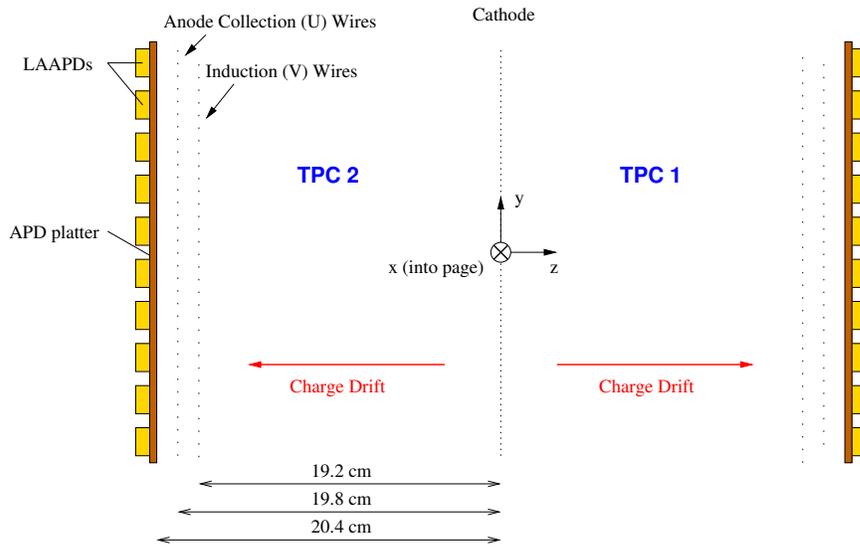


Figure 3.3: A schematic of the EXO-200 detector, which uses two back-to-back TPCs. (Not to scale.)

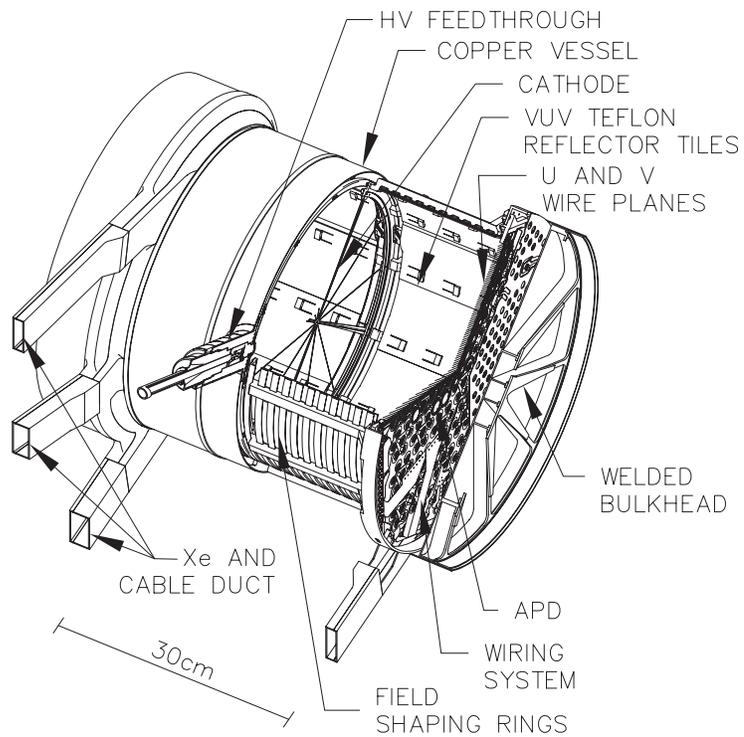


Figure 3.4: Cutaway view of the TPC, showing critical detector components and the containing copper vessel. From [22].

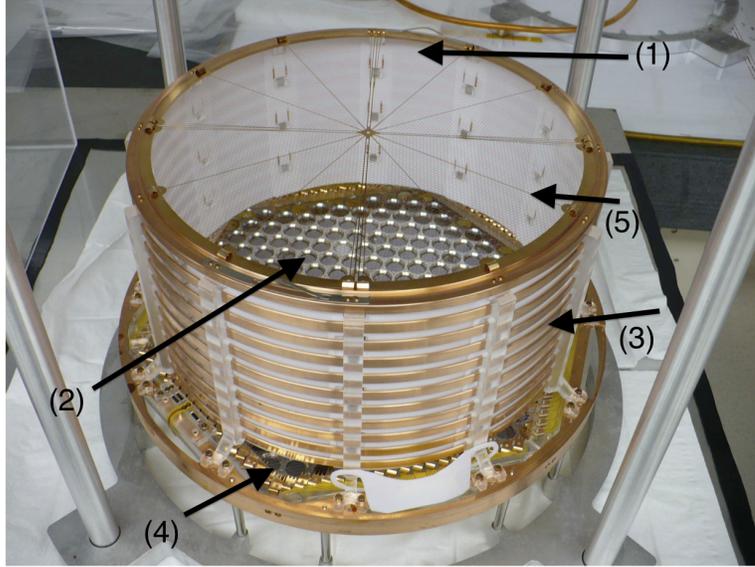


Figure 3.5: Overhead view of one of the assembled half-TPCs. Indicated are (1) the VUV-reflective PTFE tiles; (2) the LAAPD platter, with LAAPDs installed on the reverse; (3) the copper field shaping rings; (4) the anode wires and associated cabling; and (5) the hexagonal mesh cathode. From [47].

hexagon with apothem¹ 17.5 cm.

Large Area Avalanche Photodiodes (LAAPDs) are placed on a platter behind the readout wires for collection of scintillation. The high electric field in the drift region precludes the possibility of placing LAAPDs around the body of the cylinder of the TPC, which would otherwise be desirable for more efficient light collection. The scintillation signal is prompt, so the time difference between the scintillation and charge collection gives the drift time of the deposited charge. When combined with knowledge of the charge drift velocity, this allows for determination of the distance of the deposition of energy deposits from the anode. Thus, with the U- and V-wire measurements, the full 3+1-dimensional space and time position of the deposit can be reconstructed.

¹The distance from the center of the hexagon to the midpoint of one of the sides.

3.3.1 Charge Collection and Electrostatics

The cathode and anode consist of wires photoetched from 0.13 mm thick plates of phosphor bronze. For convenience in construction, the cathode is a hexagonal mesh. The U- and V-wires are rows of parallel, square profile wires spaced 3 mm apart, allowing for position discrimination. The U- and V- wires are “ganged” into groups of three; thus, the readout pitch is 9 mm. This is done to reduce the number of cables required, necessary for space and for purity for radioactivity. In total, each half-TPC has 38 U and V channels. Each end of a wire gang is mounted to one of six acrylic beams; the beams form a hexagon inside a larger copper support ring, as pictured in Figure 3.6 and schematically shown in Figure 3.7. The U- and V-wire planes are separated by 6 mm and are oriented at 60 degrees to each other due to the hexagonal detector geometry. Signals are read out on flexible copper-clad polyimide cables, after which the U-wire signals are integrated to yield a measurement of the total charge collected.

The cathode is held at -8 kV, while the V-wires are biased at -745 V and the U-wires are kept at virtual ground. Additionally, the APD platters are biased, the platter in TPC 1 at -1420 V and the platter in TPC 2 at -1390 V. To create uniformity in the electric field, ten copper rings of 37.4 cm diameter are mounted at the perimeter of each half-TPC. The rings are coupled via 900 M Ω resistors, stepping down the high voltage between rings. This arrangement establishes a 375 V/cm electric field between the cathode and the V-wires and a 750 V/cm field between the U- and V-wires, and it ensures that the V-wires are electrically transparent to

the drifting charge. To confirm this, the configuration of wires and voltages was simulated using MAXWELL, with the result that the field lines in the bulk of the detector do terminate on the U-wires and not the V-wires.

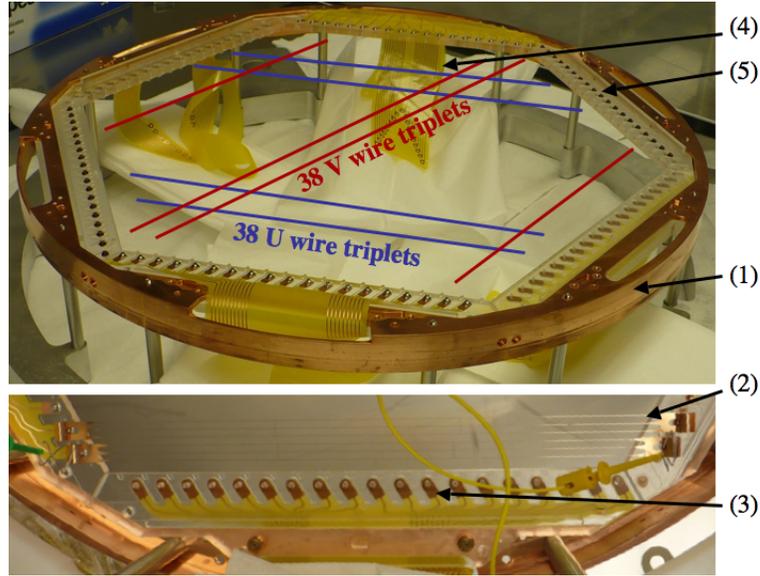


Figure 3.6: The anode support structure. Six acrylic beams (5) arranged in a hexagon are placed in a circular copper support ring (1). Anode wires are mounted on two opposing beams, U-wires (2) on one side of the ring, V-wires (not shown) on the side (3). Cabling (4) provides electrical connections at one end of each wire. The other end of the wire is simply anchored to the acrylic support (5). From [47].

3.3.2 Scintillation Collection

Scintillation photons are collected with 468 silicon Large Area Avalanche Photodiodes (LAAPDs), 234 in each chamber [48]. The LAAPDs were produced by Advanced Photonix (API) for the EXO-200 collaboration. In total, 851 LAAPDs were produced; those used in EXO-200 were chosen based on noise, gain and efficiency characteristics [49]. The quantum efficiency (QE) of each LAAPD was measured with respect to a designated standard LAAPD; 96% were found to have the same

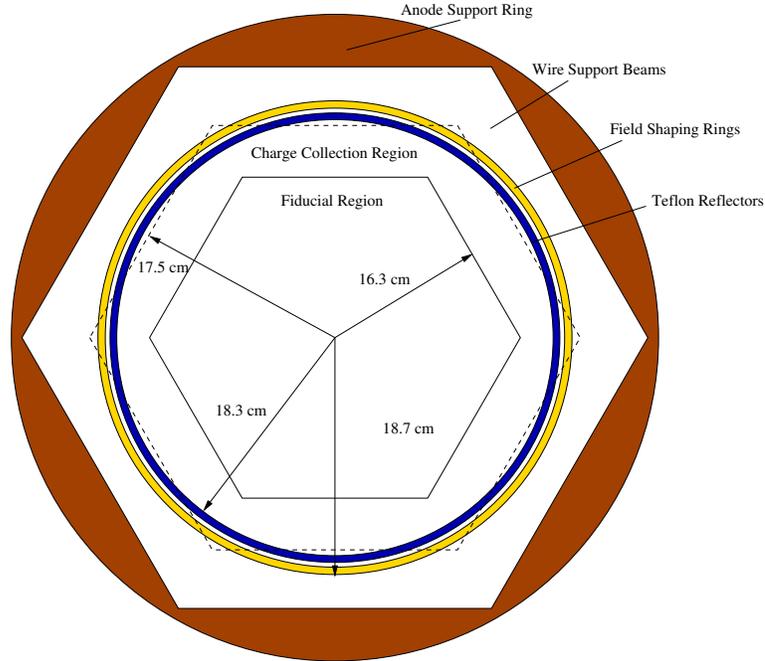


Figure 3.7: Schematic axial view of a half-TPC. The dashed line indicates the boundary between the acrylic beams (white, outside the line) and the charge collection region (white, inside dashed line). The fiducial region used for analysis is also indicated. The teflon tiles and field shaping rings each have a thickness of 1.5 mm.

QE as the standard, within 20%, and LAAPDs with relative QE lower than 80% were rejected. The gain of each LAAPD was measured over a range of different bias voltages at the operating temperature of about 170 K. The voltage required to achieve a gain of 100 was found to vary from 1360 V to 1460 V, and near that voltage, the gain was found to change about 1.5%/V, indicating a large variation in gain between LAAPDs.

The LAAPDs are used without the standard ceramic encapsulation, avoiding radioactive impurities in the ceramic and allowing for a higher packing density. The LAAPDs are mounted on two platters of low-activity copper, allowing for common biasing of the LAAPD cathodes at -1400 V. One platter can be seen in Figure 3.8, partially filled with cabled APDs. On the side facing the drift chamber, the platters

are coated with aluminum and MgF_2 to reflect photons arriving in between the LAAPDs. The LAAPDs are placed on the opposite side of the platter, which is plated with gold to improve the electrical connections.

The LAAPDs are anchored to the platter with a phosphor-bronze spring, or “spider”. The spider also electrically couples seven LAAPDs into one “gang”, as seen in the inset of Figure 3.8. Each gang is read out as a single channel by copper-clad polyimide cabling. Ganging reduces the number of cables needed, saving space and reducing radioactive background. To mitigate non-uniformities in gain, each gang is composed of LAAPDs with similar gain characteristics. Further, each gang can be individually biased via the signal cable. This “trim” voltage ranges from 10 to 60 V, and reduces the spread in LAAPD response to 2.5%.

Scintillation collection is further improved by twelve PTFE [50] tiles around the sides of the TPC. The flat tiles are 0.15 cm thick, made from skived TE-6472 modified PTFE, and They are mounted inside the copper field-shaping rings at an inner diameter of 36.6 cm. PTFE, when immersed in liquid xenon, has been measured to have a high reflectivity to 178 nm xenon scintillation light. Reflectivity near or above 90% is reported in [51] and [52]. Photon collection in the EXO-200 detector has been simulated using GEANT4, and it was found that the PTFE tiles improved collection efficiency by 50% to 150% as compared to non-reflecting walls, depending on the position in the detector.



Figure 3.8: Bare LAAPDs (1) placed on a mounting platter (2); in the TPC, the side shown faces away from the drift region. Phosphor-bronze springs (3) anchor each gang of seven LAAPDs to the platter and make electrical contact between the cathodes of the APDs and the copper traces (4) on the flexible signal cabling (5). Acrylic washers (6) electrically decouple the springs from the platter, which is coupled to the anodes of the LAAPDs. From [47].

3.3.3 Dead region

The schematic in Figure 3.7 shows that the scintillation and charge collection regions are potentially distinct. The scintillation collection region is limited by the inner diameter of the reflective Teflon tiles (inside the yellow circle in the figure). The region of efficient charge collection is limited mainly by the hexagon defined by the wire support beams, since the bulk of the field lines outside the support beams may not terminate on collection wires. It is clear that some parts of the support beams lie at smaller radius than the tiles, leaving a “dead region” of six disjoint circular segments where light will be collected but not charge. Properly accounting for the dead region is discussed in Section 7.1.

3.4 TPC vessel and xenon handling system

The two half-TPCs comprising the detector are enclosed in a cylindrical copper vessel with wall a thickness of 1.37 mm, shown in Figure 3.9. Six rectangular tubes, or “legs”, project from the vessel to carry cabling to the detector. Two of these legs are connected to a xenon handling system and are also used to evacuate the vessel and add or remove xenon. All six legs are welded to the circular copper door used to close the inner region of the cryostat. The only other penetration in the vessel itself is a feedthrough admitting a high-voltage cable, via which the cathode is biased. It can be seen pointing downward toward the viewer in Figure 3.9.

After insertion of the two half-TPCs, the vessel was sealed by TIG welding copper endcaps to the body, one of which can be seen in Figure 3.13. Weld seals were chosen over gasket seals since no gasket material with low enough radioactivity could be found. The endcaps are flared and have varying thickness; restrictions on radioactive content in the endcaps are looser, since these events are more easily rejected as fiducial and are farther from the active detector than the walls.

Trace radioactive impurities in the vessel present an unshielded source of background; this places a constraint on the acceptable thickness. At the chosen thickness of 1.37 mm, the vessel is designed to only withstand modest differential pressures of up to 35 kPa in either direction.

To ensure that this limit is not exceeded, a Labview [53] controlled “slow-control” system was built to continuously monitor and control the xenon pressure inside and the HFE pressure outside the vessel; a schematic is shown in Figure 3.10.

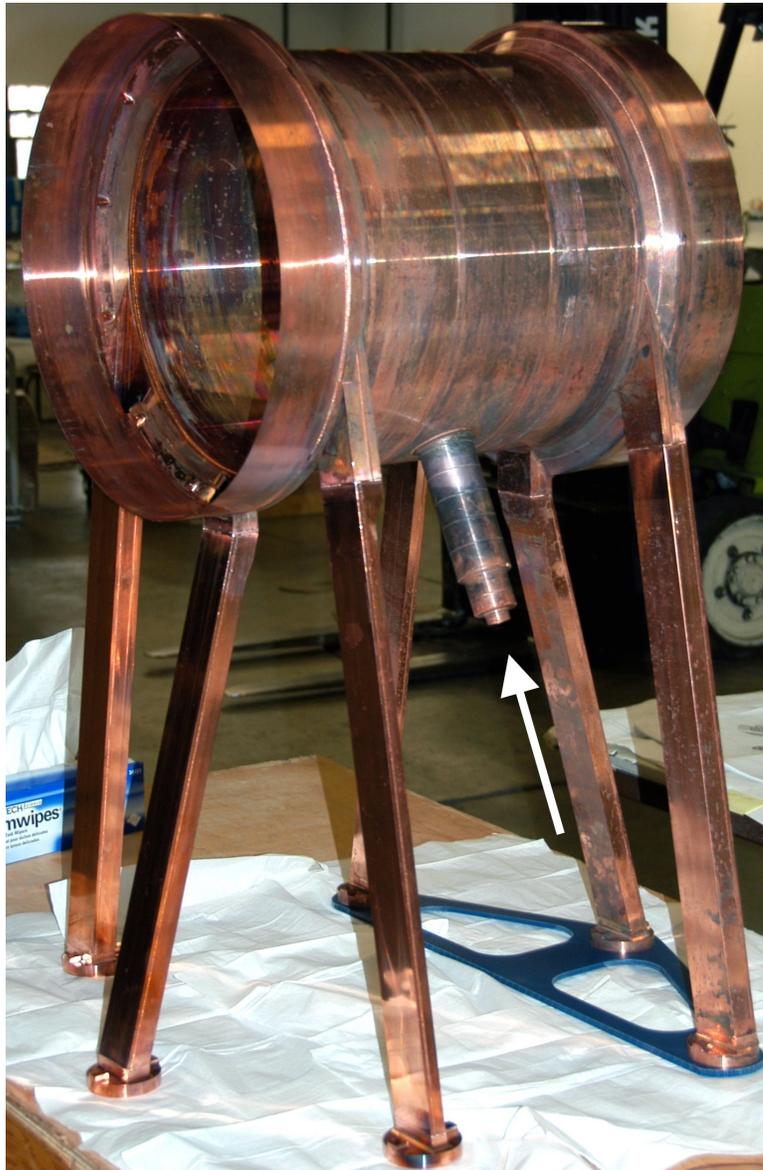


Figure 3.9: The vessel used to contain the TPCs, before installation of endcaps and welding to the cryostat door. The vessel is resting on six cable conduits, or “legs”. The white arrow indicates the penetration for the cathode high-voltage. From [47].

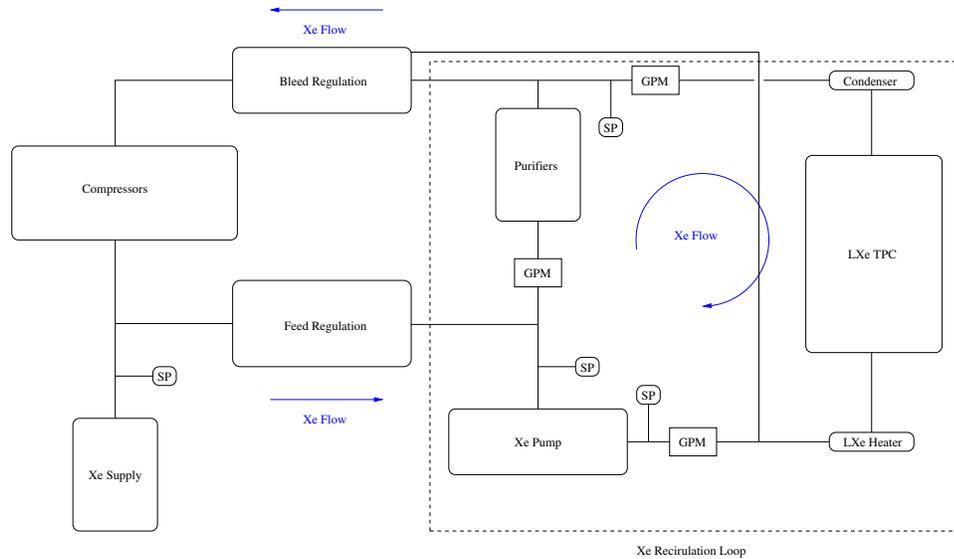


Figure 3.10: A schematic of the xenon handling system. Xenon pressure is controlled by feed and bleed systems, while xenon in the TPC is recirculated through the purifiers by the pump. Sampling Ports (SP) are locations from where xenon gas may be removed from the system for testing.

The xenon vessel is kept over-pressured with respect to the HFE, since simulations show that the vessel is under less stress when it is at positive pressure [54]. The xenon pressure is typically maintained at 1100 torr, and the HFE at 1040 torr, for a typical differential of 8 kPa outward pressure.

The slow control system sets a deadband, typically 5 kPa, around this setpoint pressure differential. When the deadband is exceeded, proportional valves in the feed or bleed regulation system are opened automatically to respectively add or remove xenon from the chamber. Xenon removed from the chamber is recovered from the recirculation loop (1-2 atm) to high pressure storage cylinders (100-800 atm) using two 2-stage gas compressors [55]. The xenon pressure is only regulated relative to the HFE; the HFE pressure thus determines the absolute pressure, as described in Section 3.5.

Boil-off xenon gas from the vessel is constantly recirculated through the external handling system, in which it is re-purified, and in which the purity can be monitored. To reduce the exposure of the xenon to impurities, a custom magnetic pump consisting of clean materials and all-metal seals was designed and built by the EXO-200 collaboration [56]. The pump provides steady flow at 14 slpm. At this rate, 200 kg of xenon can be entirely recirculated in about 40 hours.

Two SAES Model PF4-C3-R-1 noble gas purifiers, with zirconium getters as the active component, are installed in parallel in the recirculation loop. These are certified to purify rare gases to 1 ppb levels of O₂, H₂O, CO, CO₂, N₂, H₂, and CH₄. Large recirculation rates are important to process the source xenon in a practical amount of time, so filters supplied with the getters were removed to reduce impedance. Tests indicate that the getters operate more efficiently at higher temperatures [57], so additional PID-controlled heaters are installed on the getter and are used to raise the getter temperature from the default of 400 C to 550 C.

The electronegative impurity content of the xenon gas is monitored inline with three gas purity monitors (GPMs). These devices measure the thermionic emission current from a heated tungsten filament positioned in the center of the gas flow [58]. The current drops as the concentration of electronegative impurities increases. Devices are positioned at the input and output of the noble gas purifiers and at the output of the TPC, as seen in Figure 3.10. The GPMs are sensitive to about 1 part per billion of oxygen in the xenon.

Sampling ports are located at the xenon supply and before and after the purifiers. Xenon gas can be extracted from these ports and removed to a mass spec-

trometry system for later analysis; see section 3.10.1.

3.5 Cryostat

The LXe vessel is housed in the inner portion of a twelve sided, double walled cryostat, pictured in Figure 3.11. The schematic in Figure 3.12 shows the size of the cryostat and the associated lead shielding. The cryostat is made from 2.7 cm thick copper plates. Respectively, the inner and outer cryostat measure 149 and 159 cm in height and length. The inner region of the cryostat is filled with an organic cryofluid, HFE-7000 [59], providing thermal coupling to the vessel. The HFE also shields the detector from radiation; the 50 cm thickness corresponds to about 3.3 attenuation lengths for a 2.615 MeV γ . The inner cryostat is cooled by a Polycold PFC-672HC refrigerator [60] via a heat exchanger welded to the cryostat. The inner vessel is wrapped in five layers of aluminized polyester film to minimize conductive and radiative heat transfer [61]. Finally, the outer region is evacuated for further thermal insulation.

The pressure of the HFE must be regulated to be in balance with the pressure in the xenon vessel, as discussed in Section 3.4. At liquid xenon temperatures, the vapor pressure of HFE is ~ 10 Pa, too low to allow for control. Thus, the pressure of the HFE is regulated separately from the temperature using a ~ 100 L ballast tank directly coupled to the cryostat. The tank is kept about half full, and thermo-electric coolers [62] installed externally control the temperature (and vapor pressure) of the HFE.

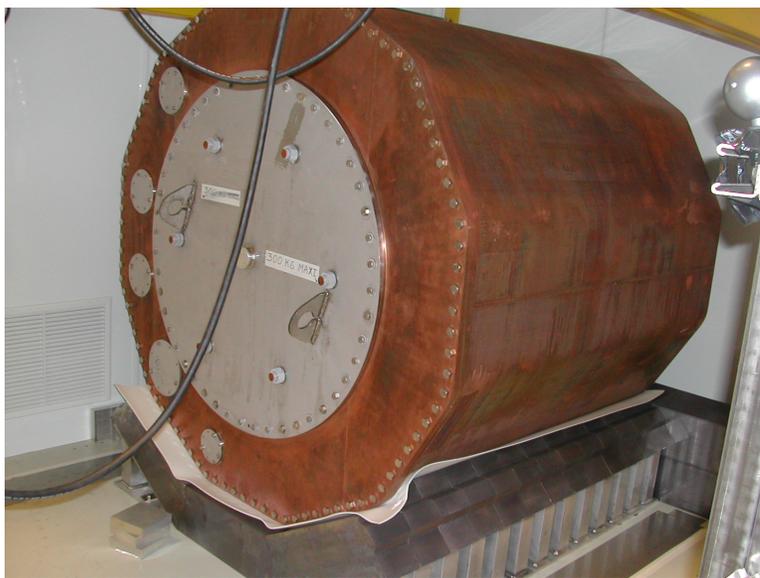


Figure 3.11: The copper cryostat. A temporary steel door is shown and is later replaced with a copper door.

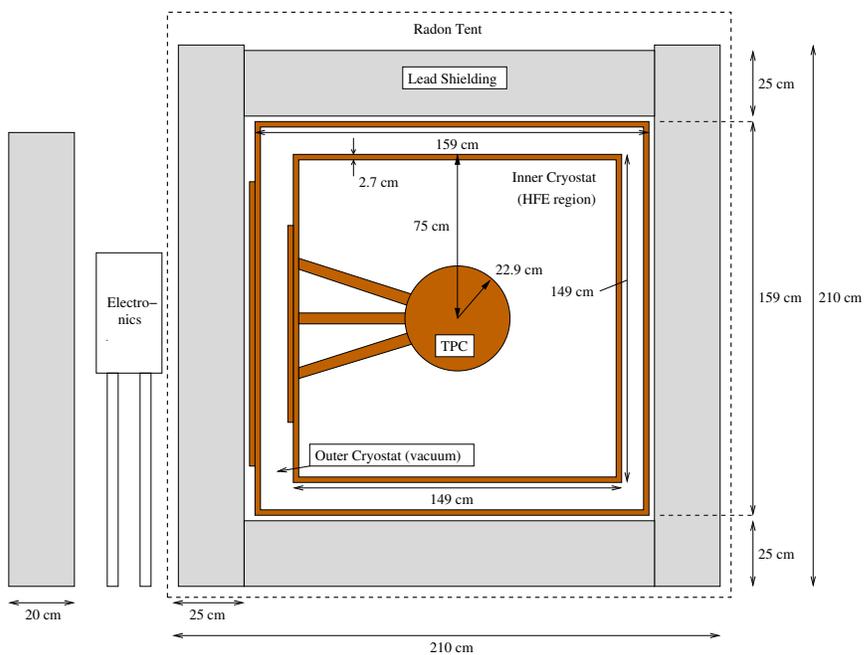


Figure 3.12: A cross-sectional view of the copper cryostat, with lead shielding and TPC in place. The heavy dashed line indicates the extent of an enclosure which can be purged to remove radon. Figure not to scale.

Each region of the cryostat has a circular opening in the face for insertion of the detector, as shown in Figure 3.13. The opening is 83.6 cm for the inner cryostat and 106.1 cm for the outer cryostat. Elastomer seals could not be used at the temperatures required for liquid xenon, so the inner cryostat was sealed with a custom “U-Mega” gasket made by Jetseal [63]. Due to the radiopurity required for parts near the detector, the gasket is made of phosphor-bronze. The gasket is coated with 125 μm of indium to make the seal. The indium itself is radioactive, but has only a β -decay at 0.496 MeV, below the threshold of interesting energies, and it is shielded from the detector by one layer of copper and by the HFE.

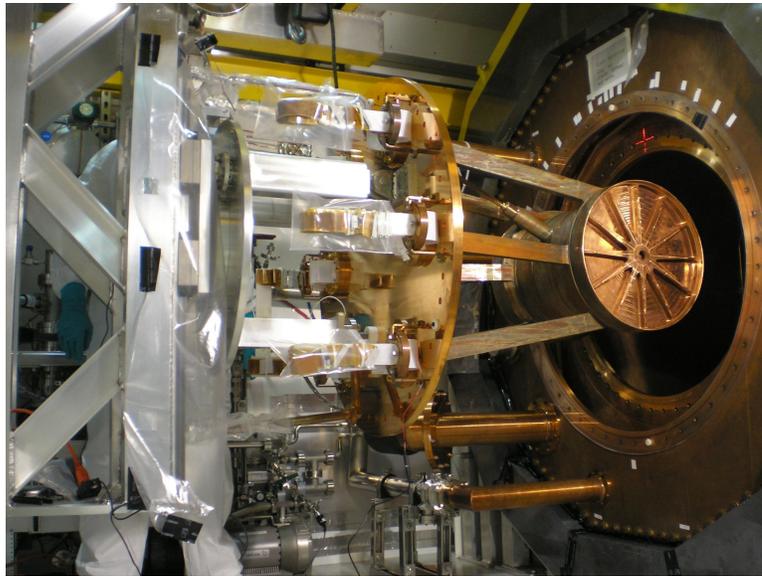


Figure 3.13: Insertion of the sealed vessel into the cryostat on January 13, 2010, showing both regions of the cryostat when open. The inner cryostat can be seen through the opening in the outer cryostat, and the lead shield can be seen surrounding the outer cryostat. From [47].

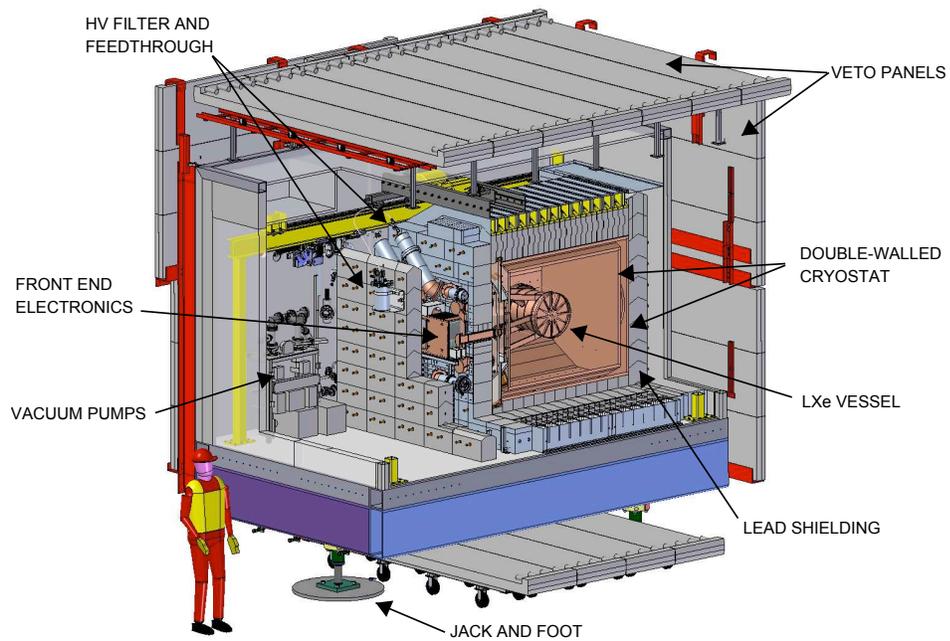


Figure 3.14: A 3-D model of the entire detector assembly. The detector is housed in a class 1000 cleanroom, which is surrounded on the outside by scintillator panels used for tagging muons. The detector is completely surrounded by a 25 cm thick lead shield, except for penetrations in the front with access to the detector. A secondary 20 cm thick lead shield provides additional shielding for these front penetrations. The vessel housing the detector can be seen inside the double-walled cryostat. From [47].

3.6 Shielding

The detector must be shielded from long-lived radioactive backgrounds in the environment. The major source comes from trace impurities in salt composing the walls of the experimental hall. These have been measured for the most prevalent isotopes, and average activities of 60 ± 47 Bq/kg, 0.54 ± 0.42 Bq/kg, and 0.76 ± 0.69 Bq/kg were found for ^{40}K , ^{232}Th , and ^{238}U , respectively.

For passive shielding, a 25 cm thick lead shield completely surrounds the cryostat on all sides. (See Figures 3.12 and 3.14.) The shield is composed mostly of 500 lb bricks. The bricks are chevroned so that no line of sight to the cryostat exists. The front shield has several openings to allow penetrations for admitting xenon to the TPC and for cabling. To mitigate this, a second 20 cm thick wall of similar bricks was installed in front of the cryostat. The HFE in the cryostat and the copper comprising the cryostat offer additional layers of shielding. The HFE is ~ 50 cm thick and the two walls of the cryostat are together 5.4 cm thick. In total, the shielding corresponds to about 29 attenuation lengths for a 2.615 MeV γ , 12 from lead, 4 from copper, and 3 from HFE.

Radon-222 in the *air gap* between the inner lead wall and the face of the cryostat presents a background unshielded by lead. The crucial background is the daughter nucleus ^{214}Bi , which has a 1.6% branching ratio to a γ at 2448 keV, quite close to the expected ^{136}Xe $\beta\beta$ Q-value of 2458 keV. To mitigate this, an “aged” air purge was implemented as follows. The unshielded air was isolated by enclosing the cryostat and inner lead wall with panels of 0.8 mm thick stainless steel sheet metal,

all sealed with copper tape. The inside of this enclosure, or “tent”, was purged with air that had been isolated for ~ 1 month, about 10 times the ^{222}Rn half-life of 3.8 d.

The purge air was introduced initially at 1 scfh, and later at 3 scfh, but measurements showed that the radon content inside and outside the tent were consistent with each other in either case. It was then determined that the tent was leaking to the HVAC system at greater than 72 scfh; a test where cleanroom air was introduced at this flow rate did not increase the pressure inside the tent. It was concluded that the leak was causing an under-pressure in the tent with respect to the cleanroom that was not compensated by the inflow of aged air, and that un-aged air was consequently being drawn from the cleanroom into the tent through leaks in the panels. The locations of the leaks from tent to cleanroom and tent to HVAC were not identified, and on 02/07/12 the purge was stopped for the remainder of the run. For this run, it will be seen that ^{214}Bi decays present a significant, but not overwhelming, source of background to $\beta\beta 0\nu$ decay.

3.7 Muon Veto

The vertical muon flux at WIPP has been measured to be $(3.10 +0.05 -0.07) \times 10^{-7} \text{ s}^{-1} \text{ cm}^{-2} \text{ sr}^{-1}$ [64]. Simulations show that this flux will lead to more background events in the region of interest for $\beta\beta 0\nu$ than is acceptable [65]. Thus, a muon veto detector is installed to tag coincidences with events in the TPC. Twenty nine panels of Bicron BC-412 plastic scintillator cover four out of six sides of the clean room module housing the detector. The panels measure 5 cm thick \times 65 cm tall, and

Table 3.1: Calibration γ sources. Activities listed as measured September 1, 2009, 12:00 PST. Note that there are two sources of each isotope, each with a different activity. These are labelled “strong” and “weak”.

| Isotope | Weak (kBq) | Strong (kBq) | γ energies (keV) | Half-life (yr) |
|-------------------|------------|--------------|-------------------------|----------------|
| ^{137}Cs | 2.820 | 13.14 | 662 | 30.2 |
| ^{60}Co | 0.5295 | 7.060 | 1173, 1333 | 5.3 |
| ^{228}Th | 1.417 | 34.04 | 2615 | 1.9 |

are variously either 315 cm or 375 cm wide. In each panel, eight 2” photomultiplier tubes, four at each end, read out the signal.

3.8 Calibration system

To calibrate the detector, sources of γ -rays are placed near the TPC, in the inner (cryofluid) region of the cryostat. To allow for insertion and removal, a 1/4” copper guide tube penetrates the cryostat and is bent around the TPC, as shown in Figure 3.15. Calibration sources are inserted into the tube via a wire and sprocket system. The system contains several marks to indicate the position of the source in the cryostat, allowing for repeatable calibration. The standard positions used are listed in Table 3.2; other positions are possible but were not used. Three sources are used, with two source strengths per source. The sources are listed with activities and γ -ray energies in table 3.1.

3.9 DAQ/Electronics

Figure 3.16 shows a schematic of the EXO-200 data acquisition electronics (DAQ) system. Two boxes made of low radioactivity copper are placed outside the

Table 3.2: Source calibration positions. Coordinates are aligned with the origin at the center of the TPC, on the cylinder axis. The x-axis points perpendicular to and away from the cryostat door, the y-axis points vertically, and z-axis points along the axis of the TPC to the right when viewed from outside the cryostat.

| Source position | Nominal coordinates (cm) |
|-----------------|--------------------------|
| S2 | (0, 0, -30.4) |
| S5 | (25.4, 0, 0) |
| S8 | (0, 0, 30.4) |
| S11 | (0, 30.4, 0) |
| S17 | (0, -30.4, 0) |

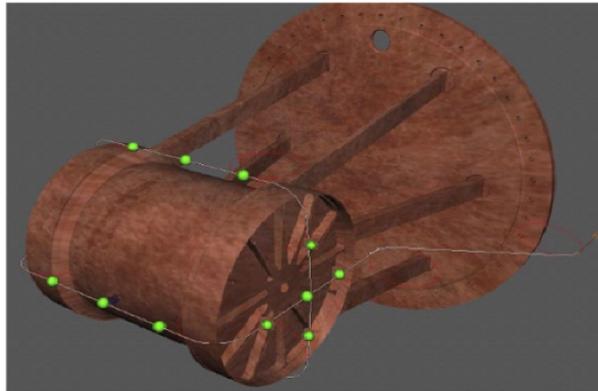


Figure 3.15: The calibration guide tube in position around the LXe vessel. Green dots indicate locations designated for calibration sources. From [47].

cryostat to hold eighteen Front End Cards (FECs). The FECs amplify, shape and digitize signals coming from the u-wires, v-wires and LAAPDs. Analog to Digital Converters digitize the signal with a 1 MHz sampling rate.

Each channel passes through a replaceable shaping circuit, consisting of two integrators, two differentiators, and one preamplifier with an intrinsic differentiation time. The shaping circuits act as a notch filter to remove noise having time scales different than the expected signal times of $\sim 1\text{-}100 \mu\text{s}$. The shaping constants used are shown in Table 3.3. During data runs before the one used in this work, the constants for the U-wires were more similar to those of the V-wires, but the integration obscured information about induction signals generated on $3 \mu\text{s}$ time scales when charge drifts between the V- and U-wires, while the differentiation cut into the signal peak, obscuring energy information. To compensate, the U-wire shaping constants were adjusted to allow a wider band pass. The differentiation time of each U-wire preamplifier is set by a 1 pF capacitor and a 60 M Ω resistor. These capacitors are rated to ± 0.5 pF, implying a 50% uncertainty on the differentiation time. Other components have uncertainties of 5-10%, as well. The effects of this channel-to-channel variation will be discussed in Chapter 4. Each channel also contains a capacitor that injects a fixed amount of charge, used during routine calibrations to track variations in gain.

Optical connections transmit the digitized signals to a Trigger Event Module (TEM). The TEM groups signals into events containing 2048 samples from all wire and LAAPD channels. The grouping is arranged such that the trigger is in the middle of the recorded waveform. Thus, each event corresponds to 1024 μs of data

from before the trigger and 1024 μs of data after the trigger. Assembled events are then sent to a control PC for storage on disk. Events can be triggered by a single U-wire, V-wire, or LAAPD channel crossing a threshold, or the sum signal of all U-wire, V-wire, or LAAPD signals crossing a threshold. Forced triggers can also be made at a pre-defined interval.

Table 3.3: Nominal values of the shaping constants for integrators and differentiators in the Front End Cards used to read out signals from the TPC. Note that there are two integrators and differentiators applied to each channel, each with the same time, while there is a single preamplifier. The U-wire differentiation time varies significantly due to component uncertainty; only the nominal value is provided.

| Signal | Integrators (μs) | Differentiators (μs) | Preamp. Differentiator (μs) |
|--------|-------------------------------|-----------------------------------|--|
| APD | 3 | 10 | 300 |
| U-Wire | 1.5 | 40 | 60 (nominal) |
| V-Wire | 3 | 10 | 60 |

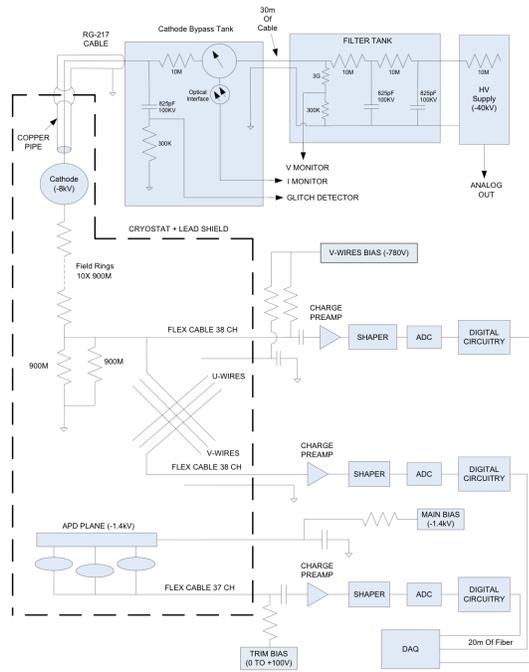


Figure 3.16: A schematic of the EXO-200 electronics system, from [47].

3.10 Material selection

All materials used in the EXO-200 detector, as well as some critical external components, have been gone through a careful screening program to select only materials with low concentration of radioactive isotopes [66]. Gamma radiation from these materials represents the primary source of background for EXO-200. Tolerances for the concentration are determined using MC simulations which include details about the composition and position of materials. Tolerances are tighter for more massive objects and for material near the active xenon. The primary isotopes of concern are ^{238}U , ^{232}Th , and ^{40}K , all high-energy γ emitters. These isotopes are long-lived, with half-lives $\sim 10^9$ yrs, and are present in all materials in trace amounts; thus they can never be eliminated entirely. Multiple methods were used to measure the radioactive content of the various source materials, including neutron activation analysis (NAA), mass spectrometry (MS), and direct counting. In NAA, a small sample (~ 1 g) is exposed to a high neutron flux from a reactor, and the nuclei are activated. The decay of activated nuclei is then counted in a germanium detector. Knowing the neutron flux and the capture cross section of the target nucleus allows one to infer the concentration of the non-activated target nucleus. This technique is slow to implement and is subject to contamination from activation of nuclei other than the one measured. Furthermore, materials that can be activated, such as copper, cannot be studied this way, since the entire sample becomes activated and hence unsafe to handle.

In MS, a small sample (~ 1 g) is ionized and an electric field separates the com-

ponent species. Two types of MS were used, differing in the method of ionization. In inductively-coupled plasma MS (ICP-MS), a sample is digested in acid, typically HNO_3 , and the solution is preconcentrated. The resulting solution is heated until a plasma forms. Glow-discharge MS (GD-MS), by contrast, uses a solid sample as an electrode to create a glow discharge in a buffer gas. ICP-MS requires a soluble sample, while GD-MS requires a conductive sample. MS is very sensitive; for EXO-200 measurements, these techniques can be sensitive to 10^{-12} g/g, with ICP-MS generally achieving higher sensitivity due to the preconcentration of the sample.

A detector component can also be placed in a germanium counter, which will directly measure the rate of the various γ decays. Counting this way is most relatable to EXO-200 measurements, but it is slow and requires a large sample, making it not useful for critical components that must remain intact. This method is not as sensitive as the others but it can be useful as a cross-check.

Lead used in shielding the detector was also screened for ^{210}Pb , an β emitter and parent of β emitters. Due to the mass of the lead and its proximity to the detector, bremsstrahlung for these β s becomes a potential background. The concentration of ^{210}Pb was measured by counting 5.3 MeV α decays of ^{210}Po , a daughter nucleus, in a silicon detector. The activity of ^{210}Pb was measured at or below 25 ± 4 Bq/kg for several samples of lead.

Table 3.4 shows a summary of U, Th, and K measurements made for several of the most important components: the copper comprising the vessel, the LAAPDs, the signal cabling, the reflective Teflon (TE-6472) tiles, acrylic supports inside the vessel, the HFE-700 cryofluid, and the lead shielding. Some of these materials were supplied

as lots or several samples were taken. In these cases, the values or limits reported in this table reflect are the most conservative of all lots or samples measured. The sample cables were etched before measurement to more accurately represent the EXO-200 cables. The contamination in the LAAPDs lies predominantly in the aluminum contacts.

This table shows that the signal cables are the major background concern, due to their high activity. These parts were photo-etched, which requires several chemicals that introduce radioactive impurities. This is somewhat mitigated by the small total mass of the cables, but the cables are immersed in the liquid xenon, which increases the impact. Simulations show that when all factors are combined, the cables are in fact the largest source of background. The next largest concern is the copper vessel, due to its proximity and mass. Only limits could be placed on the concentration of impurities, so the background impact is unclear. Note that the Teflon is not a significant background source, due to its extremely low impurity content.

3.10.1 Xenon source material

The xenon used for EXO-200 was enriched from natural xenon by ultracentrifugation to 80.6% in ^{136}Xe . Other natural isotopes remain, with ^{134}Xe constituting 19.1%, and other natural isotopes forming an unimportant remainder. The measured mass spectrum is shown in Figure 3.17. Prior to admitting it to the detector, the xenon was recirculated through the purifiers (see Section 3.4) until

Table 3.4: Summary of concentration measurements of K, Th, and U in important detector materials. Several measurement methods are indicated: Ge stands for germanium counting, ICP-MS for inductively-coupled plasma mass spectroscopy, and GD-MS for glow-discharge mass spectroscopy. For some materials, the value reported represents the least stringent limit over several lots. Signal cables are measured after etching.

| Material | Method | K (10^{-9} g/g) | Th (10^{-12} g/g) | U (10^{-12} g/g) |
|--|--------|--------------------|----------------------|---------------------|
| Copper | ICP-MS | < 55 | < 2.4 | < 2.9 |
| Copper | Ge | < 120 | < 35 | < 63 |
| LAAPDs (aluminum contacts) | ICP-MS | 190 ± 40 | 45 ± 2 | 76 ± 4 |
| LAAPDs (aluminum contacts) | Ge | 490 ± 160 | < 630 | < 360 |
| Signal cables - copper cladding | ICP-MS | - | < 130 | 463 ± 34 |
| Signal cables - polyimide substrate | ICP-MS | - | < 646 | 1320 ± 196 |
| Teflon (TE-6472) | NAA | 1.8 ± 0.2 | < 0.26 | < 0.78 |
| Teflon (TE-6472) | Ge | < 740 | < 112 | < 200 |
| Acrylic supports | NAA | < 2.3 | < 14 | < 24 |
| HFE-700 | NAA | < 1.78 | < 8.4 | < 7.3 |
| Lead | ICP-MS | - | < 1 | < 1 |
| Lead | GD-MS | < 15 | < 7 | < 8 |

inline measurements showed it was pure at ~ 1 ppb. For studying trace impurities, both electronegative and radioactive, the xenon can be directly sampled and removed from the system to a mass spectrometer/liquid nitrogen trap [67]. This technique allows for measurement of any species with a significant partial pressure at liquid nitrogen temperatures, including oxygen, nitrogen, and krypton. Oxygen and nitrogen content were measured before passing the enriched xenon through the purifiers. Nitrogen was found at $329 \pm 49 \times 10^{-9}$ g/g and oxygen was not detected, with a sensitivity of 0.4×10^{-9} g/g. This concentration of nitrogen corresponds to approximately 65 mg in the entire supply, an amount that can safely be removed by the purifiers. A common radioactive background for liquid noble gas experiments is ^{85}Kr , which undergoes β -decay with a Q-value of 687.1 keV and a half-life of 10.8 yr. The concentration of natural Kr was measured to be $25.5 \pm 3 \times 10^{-12}$ g/g ($^{nat}\text{Kr}/^{136}\text{Xe}$). This leads to a rate of 10^4 decays/yr, low enough to be sub-dominant to $\beta\beta 2\nu$ decay. The Q-value of the decay is too low to interfere with observing $\beta\beta 0\nu$.

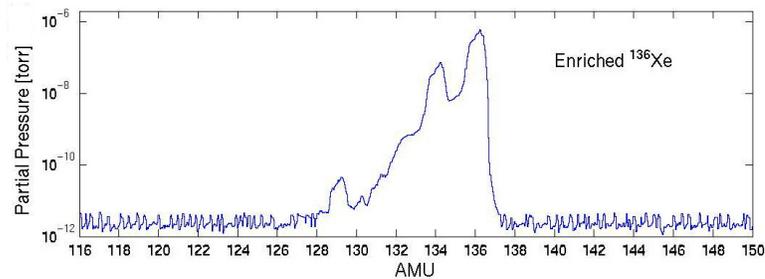


Figure 3.17: Isotopic spectrum of enriched Xe used in EXO-200. From [47].

Chapter 4

Event Reconstruction

Each DAQ trigger results in a recorded frame consisting of a 2.048 ms waveform (2048 samples at 1 μ s per sample) from each of 76 U-wire channels, 76 V-wire channels, and 74 LAAPD channels. A typical physics run lasts 24 hours and contains 10,000-20,000 frames, and daily two-hour calibration runs contain about 250,000 frames. The large number of events and channels necessitates a fast and automatic way of identifying signals and extracting information for analysis.

Reconstruction is accomplished in three stages: pattern recognition, parameter estimation, and clustering. First, candidate physical signals from U-wires, V-wires and LAAPDs are identified. Then position, timing, and energy information are extracted. Finally, charge signals nearby in time and space are associated into localized energy deposits, termed *clusters*, and *charge clusters* are associated with scintillation signals.

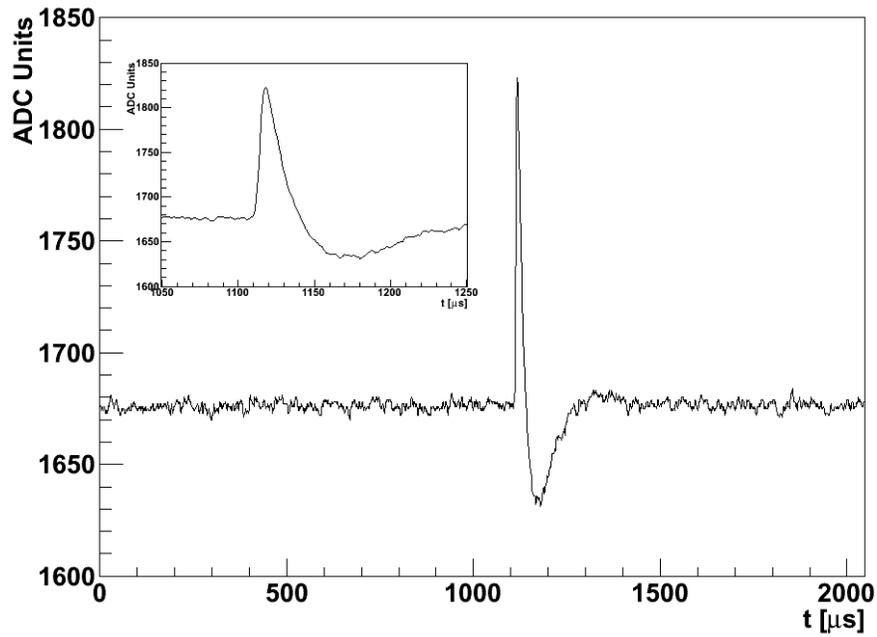
4.1 Event characteristics

Figure 4.1 shows typical U- and V-wire signals; Figure 4.2 shows two superimposed. The initial rise in each waveform is caused by induction from charge drifting towards the wires. The V-wire signal rises due to induction when charge starts drifting in the bulk. U-wires are mostly shielded from induction by the V-wires, so

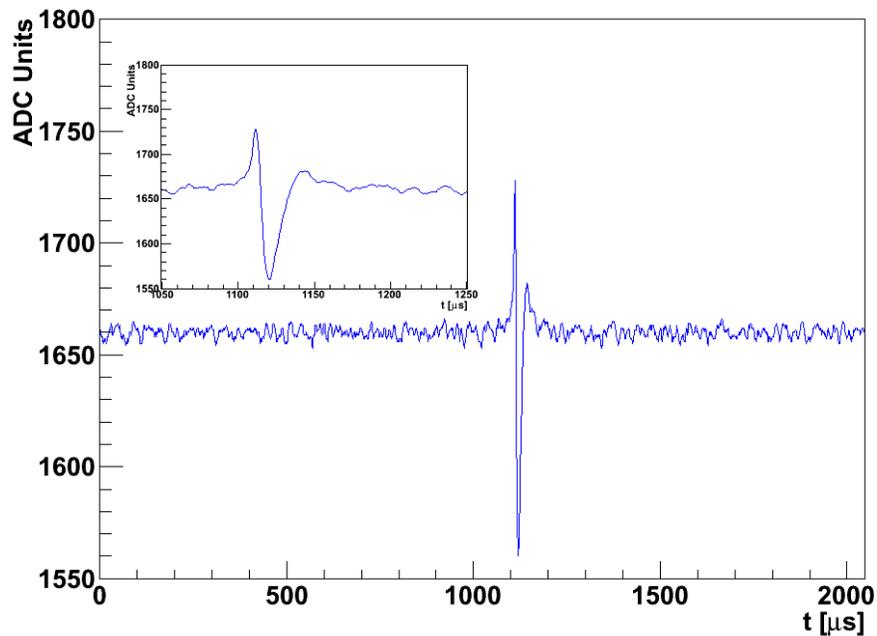
the signal does not begin until after charge passes the V-wires. At this time, the V-wire signal reverses polarity. The U-wire signal peaks when charge is collected on the U-wires. Both signals subsequently drop to baseline under the action of the shaping differentiators.

A typical APD waveform is shown in Figure 4.3. To improve signal to noise ratio, the signals from all APDs in each half-TPC are summed, as shown in Figure 4.4. The sum is used for triggering and analysis. Scintillation is produced on timescales of ~ 1 ns, much smaller than the sampling rate, so the shape of the signal is determined entirely by the shaping electronics.

Figure 4.5 shows the responses of all channels to a typical event. The heat plots show the responses at every sample time; schematics of the detector show the location of each channel and the response at the time the signal peaked. Several features of the TPC are apparent. Typical charge deposits have a spatial extent of ~ 3 mm, so only a small number of wire channels record a signal. More V-wires collect signals than U-wires, since only an induction signal is necessary. The electric field ensures that each charge signal is recorded in only one half-TPC; the event shown occurred in TPC 1. Scintillation light, conversely, yields signals on many APDs in both half-TPCs, since it is emitted isotropically and can be reflected several times by the teflon panels and aluminized APD platters.



(a) U-wire charge collection waveform



(b) V-wire charge induction waveform

Figure 4.1: U-wire (4.1a) and V-wire (4.1b) signals from a typical energy deposit. The channels with the largest amplitude signals are shown. The U-wire signal shows a short induction rise time from charge that has passed the V-wires. The V-wire signal is bipolar, with a rising induction signal as the charge approaches the wires, and a falling signal after the charge has passed. Both signals are restored to baseline by the electronic differentiators. The signals are taken from physics run 2637, event 750, channels 22 and 59.

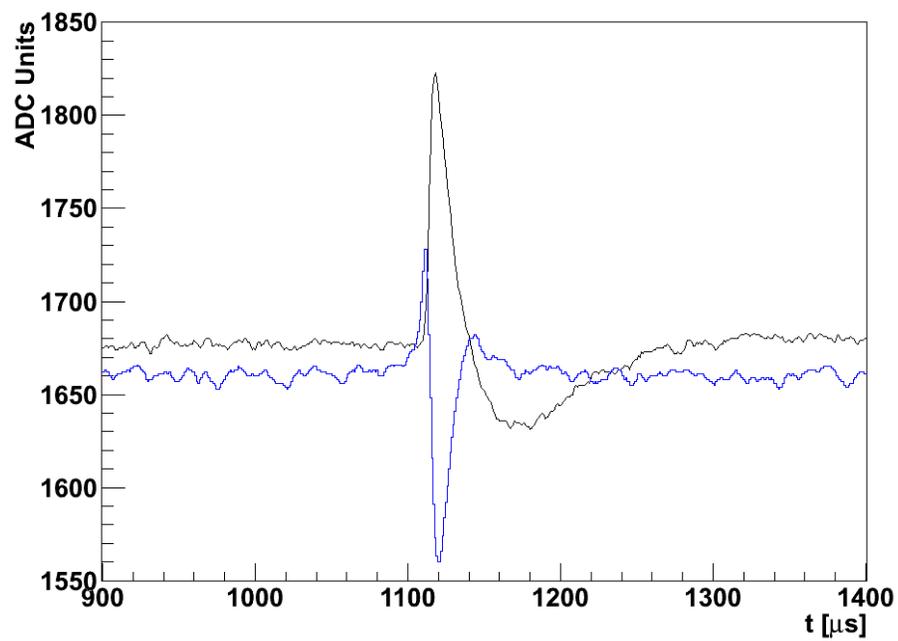


Figure 4.2: Comparison of the above U-wire (black) and V-wire (blue) signals. Notice the V-wire rise time begins earlier than the U-wires, since the V-wires are not shielded from induction in the bulk. The signals are taken from physics run 2637, event 750, channels 22 and 59.

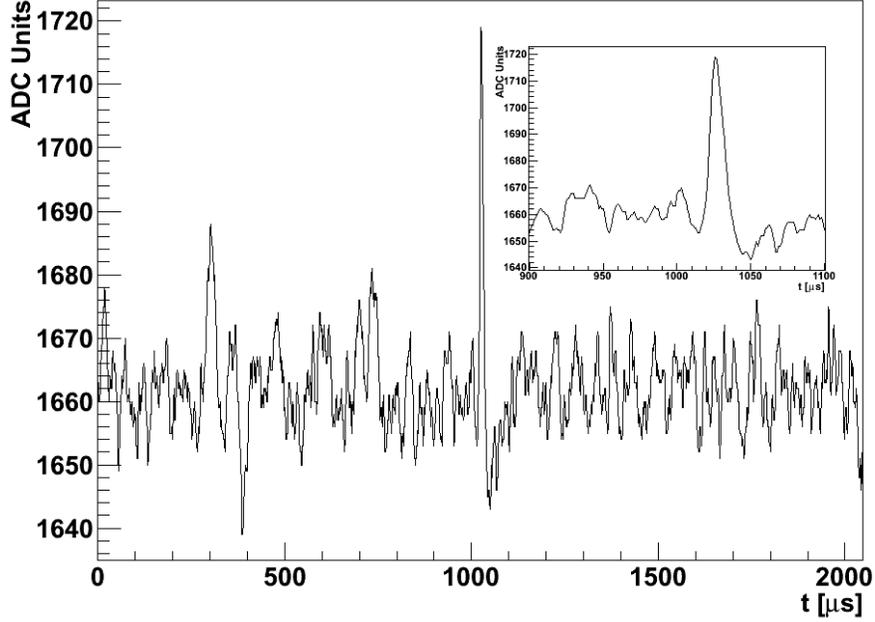


Figure 4.3: A single-channel APD waveform corresponding to the wire waveforms shown above. This channel had the signal with the largest amplitude of the APD channels. The shape is entirely determined by the shaping electronics. Note that the APD waveform begins slightly earlier than the charge signal, at $1024 \mu\text{s}$; the delay in the charge signal is due to the charge drift. The signal shown is taken from physics run 2637, event 750, channel 157.

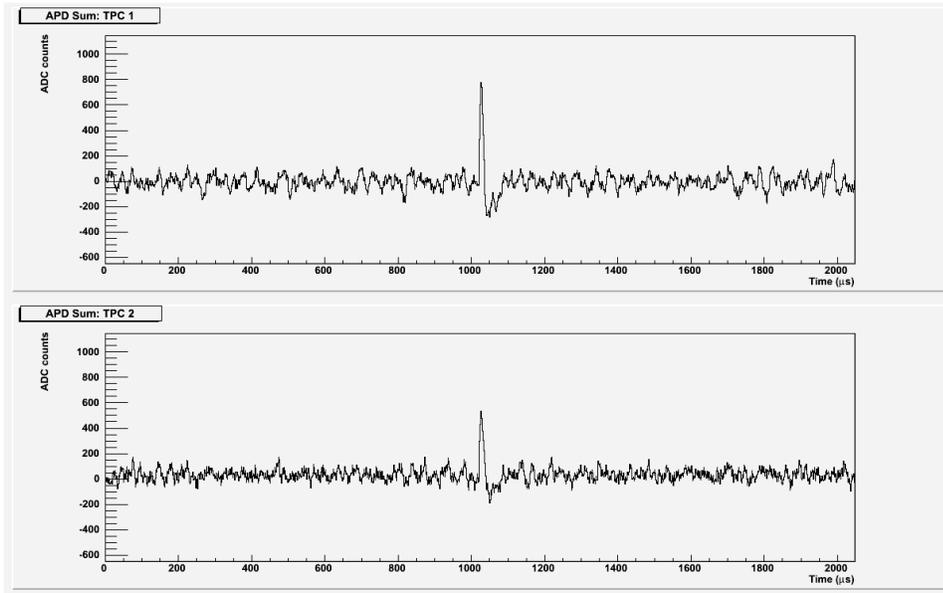
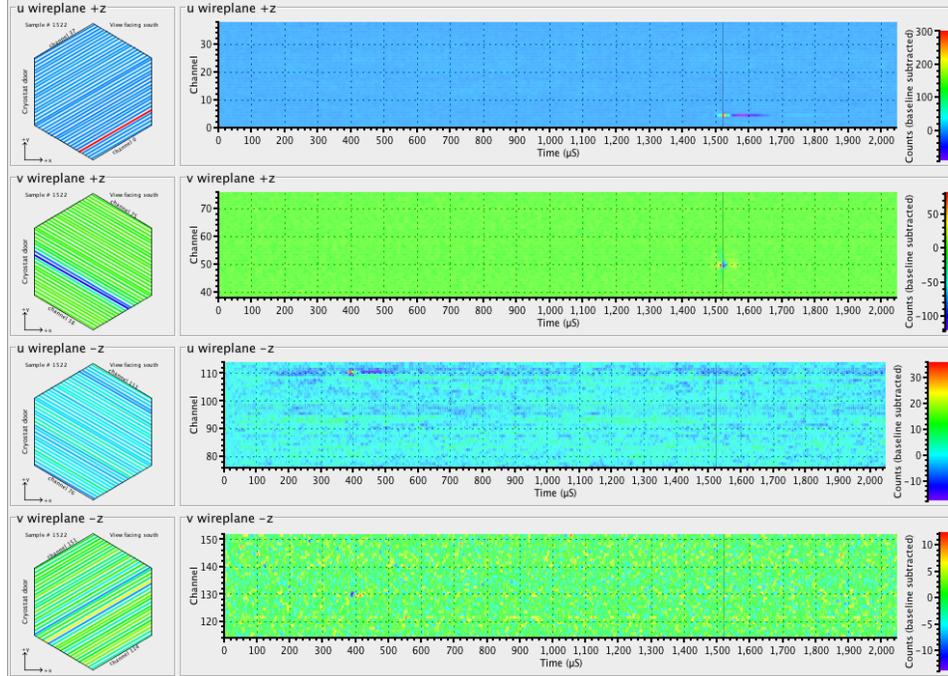
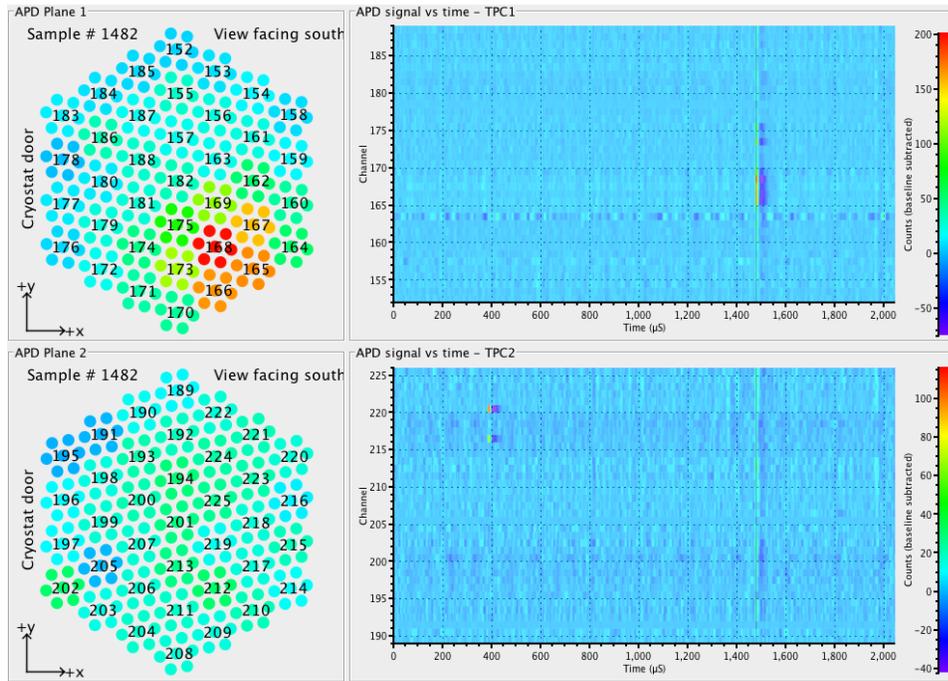


Figure 4.4: APD sum signals for run 2637, event 750. Note the improved signal-to-noise relative to the single-channel APD waveform in Figure 4.3. The signal is larger in TPC 1, the half-TPC in which charge was deposited.



(a)



(b)

Figure 4.5: An event display of a typical energy deposit in the TPC. The right side of each subfigure contains a heat plot of wire (a) or APD (b) channel responses, including a vertical grey bar indicating the time of peak U-Wire or APD signal. The left side contains a schematic of the detector showing the position of each channel and the signal strength at the time of the peak. (The APDs are indicated in the “gangs” of seven per channel.) The event shown occurs near sample time $1500 \mu\text{s}$ in TPC 1 (+z).

4.2 Pattern Recognition

Two types of signal identification are used. Matched filters are used to identify the presence of a signal on U-wire, V-wire and APD waveforms. Once a primary signal is found, the waveform is filtered again to find secondary signals, using a method termed *unshaping*. The matched filter is inefficient at identifying multiple U-wire signals on a single waveform, so this unshaping uses a filter with a narrower time response to achieve better resolution.

4.2.1 Matched filters

Matched filtering is a technique for identifying signals the presence of noise. A template, representing the signal of interest, is correlated with a noisy data signal. The magnitude of the matched filter output correlates to the likelihood of the presence of the template. We can write the data signal as $r(t) = s(t) + n(t)$, where $s(t)$ is the signal of interest and $n(t)$ is the noise. Then applying a generic filter $H(t)$ gives a filtered signal, $r_f(t)$:

$$\begin{aligned} r_f(t) &= \int_{-\infty}^{\infty} \tilde{H}(f) \tilde{r}(f) e^{i2\pi ft} \, df \\ &= \int_{-\infty}^{\infty} \tilde{H}(f) \tilde{s}(f) e^{i2\pi ft} \, df + \int_{-\infty}^{\infty} \tilde{H}(f) \tilde{n}(f) e^{i2\pi ft} \, df \\ &= s_f(t) + n_f(t), \end{aligned} \tag{4.1}$$

where tildes denote Fourier transforms, and $s_f(t)$ and $n_f(t)$ are the filtered signal and noise, respectively. It can be shown that under the assumption of white Gaus-

sian noise, the signal-to-noise ratio, $|r_f(t)|^2/|n_f(t)|^2$, is maximized when the filter is related to the conjugate of the signal of interest as

$$\tilde{H}_{mf}(f) = k\tilde{s}^*(f)e^{-i2\pi ft_d}, \quad (4.2)$$

where k is an arbitrary constant and t_d is a reference time. This filter $\tilde{H}_{mf}(f)$ is then the matched filter. In the case of EXO-200, the noise is non-Gaussian due to pickup. To account for this, the signals and filter are “whitened” by dividing by the magnitude of the noise, $\tilde{n}(f)$. Thus we use the whitened signal

$$\tilde{r}'(f) = \frac{\tilde{r}(f)}{|\tilde{n}(f)|} \quad (4.3)$$

and whitened filter

$$\tilde{H}'_{mf}(f) = \frac{\tilde{H}_{mf}}{|\tilde{n}(f)|}. \quad (4.4)$$

For EXO-200, a matched filter template, $s(t)$, is chosen for each type of wire signal and for APD signals. The templates are model signals that are modified by a transfer function that accounts for the electronic shaping discussed in Section 3.9. The U-wire and V-wire model signals are generated using a simulation of a charge drifting in the electric field of the TPC. These models may vary slightly from day to day due to fluctuations in drift velocity, but are otherwise static, and are always constant during a daily run. For the APDs, a step function is used as the model.

The shaping times of the V-wire and APD transfer functions are taken to be the nominal values listed in Table 3.3. Thus, the V-wire template and APD template

can be used for all signals of those types. The U-wire transfer function (and hence matched filter template) is different for each channel to account for the hardware variations; see Section 4.2.2. An example U-wire waveform fitted to a template function is shown in Figure 4.6.

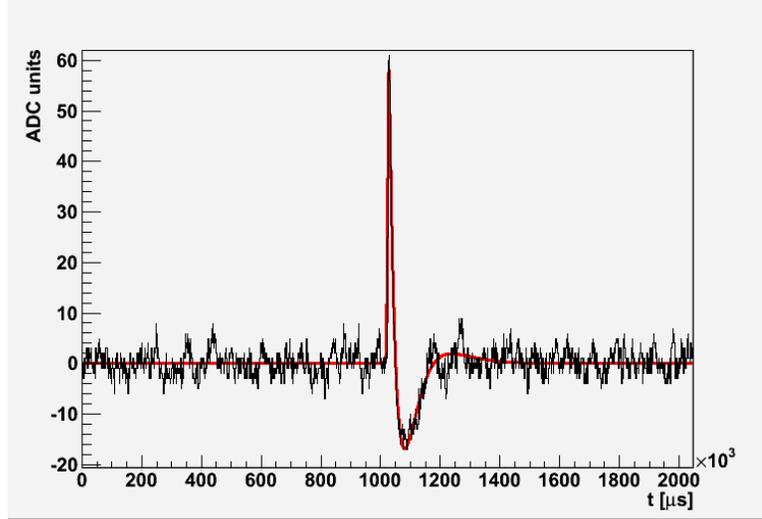


Figure 4.6: Example U-wire signal (black) fitted to a template signal (red).

For U- and V-wires, the matched filter is applied to every channel of every recorded event. To improve the signal-to-noise ratio for scintillation signals, all the APD signals from each half-TPC are summed. The filter is applied to both sums individually. For both U-wire and APD signals, the matched filter is applied at every time sample, returning a response function like the one in Figure 4.7. Signals that above a threshold are accepted as candidate events. Thresholds of 5 (4) times the mean absolute deviation of the baseline are used for the U-wires (APDs). The sample time of peak response is taken as the approximate signal collection time.

The V-wire matched filter is only applied at the approximate U-wire signal time, resulting in a single number, a Figure of Merit (FOM). The distribution of the

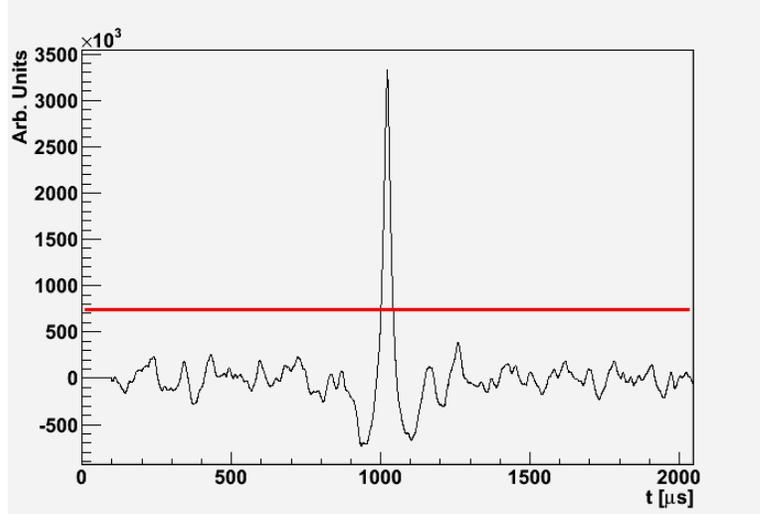


Figure 4.7: Match filter response to a U-wire signal, shown as a function of the sample time at which the match filter is evaluated. The red line indicates the threshold for signal detection.

FOM in response to noise input is shown in Figure 4.8. It is fitted to a Gaussian distribution with RMS $\sigma_V = 3.858 \pm 0.010$ (arb. units).

The efficiency of the V-wire matched filter was studied as a function of the FOM threshold. It is critical to have a low threshold, since it is difficult to apply fiducial cuts to events where some V-wire signals are missed. This problem is especially acute for multi-site events which may consist of several low-energy clusters. Calibration run 2448, a strong ^{228}Th at position S5, was studied, and the ratio of the number of charge clusters reconstructed with V-wire coordinates to the number with U-wire coordinates was taken as the filter efficiency. The chosen threshold of $3.6 \sigma_V$ leads to an efficiency of 98% or higher above 273.7 keV, and a false positive rate of 0.9%. The efficiency curve for this threshold is shown in red in figure 4.9.

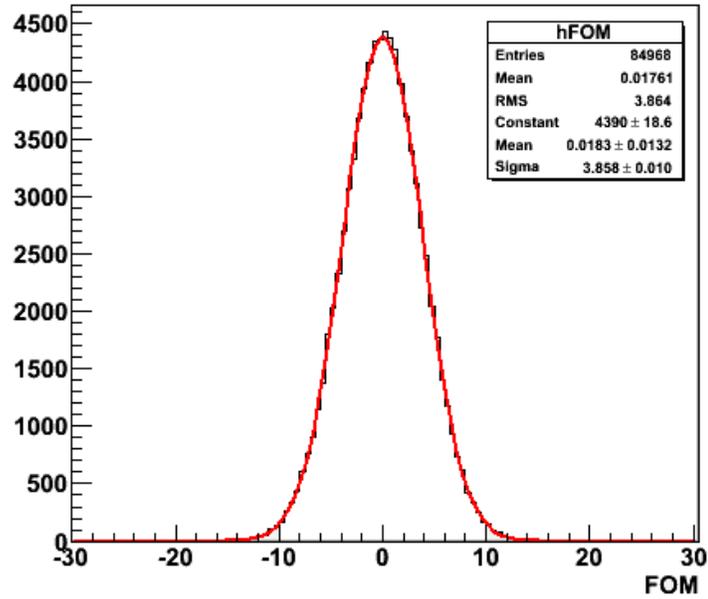


Figure 4.8: Distribution of the v-wire match filter response (FOM, or Figure of Merit) to noise signals. The distribution is fit to a Gaussian with width $\sigma_V = 3.858$ (arb. units). The signal detection threshold is $3.6\sigma_V$, or 13.9 (arb. units), leading to a detection efficiency of at least 98%. See also Figure 4.9. Courtesy of G. Giroux.

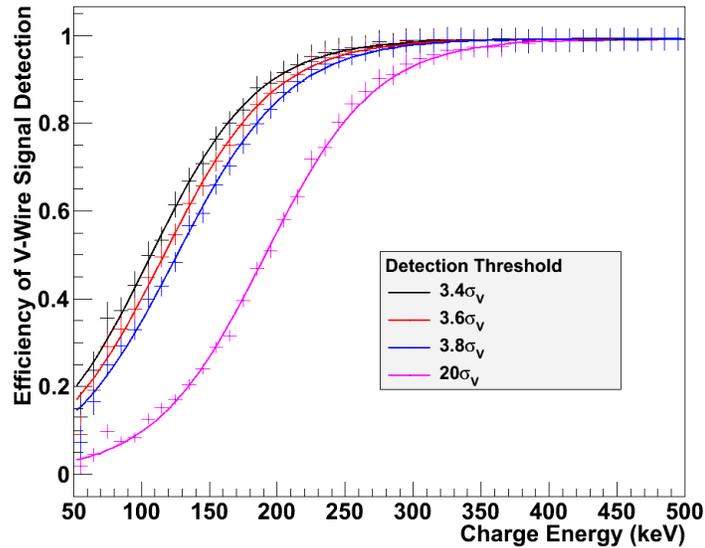


Figure 4.9: V-wire signal detection efficiency as a function reconstructed charge energy. The points shown correspond to matched filter thresholds of $3.4\sigma_V$, $3.6\sigma_V$, $3.8\sigma_V$, and $20\sigma_V$. A threshold of $3.6\sigma_V$ (red) was chosen for this analysis. The continuous curves are best fit sigmoids. Figure courtesy of G. Giroux.

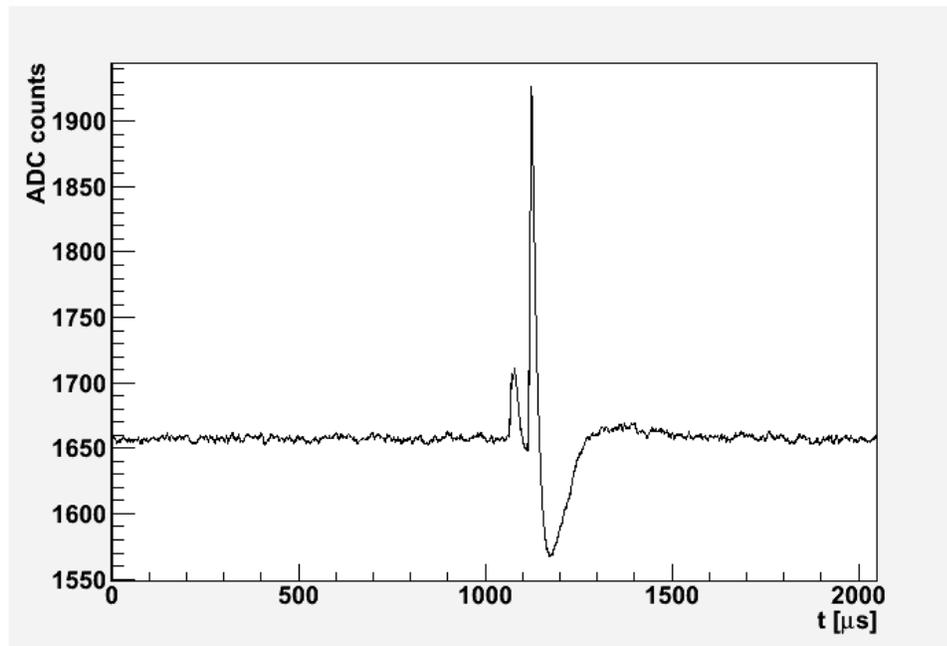
4.2.2 U-Wire shaping time variance

Due to hardware variation, template signals generated using the nominal U-wire differentiation time of $60 \mu\text{s}$ are poor fits to actual signals. To compensate for this, the preamplifier differentiation times for each U-wire channel is measured. A calibrated amount of charge is injected, and the channel response is fit to a template signal in which the differentiation time is allowed to float. The best fit time is not a true measurement of the preamplifier differentiation time alone, since other component uncertainties are folded into the response.

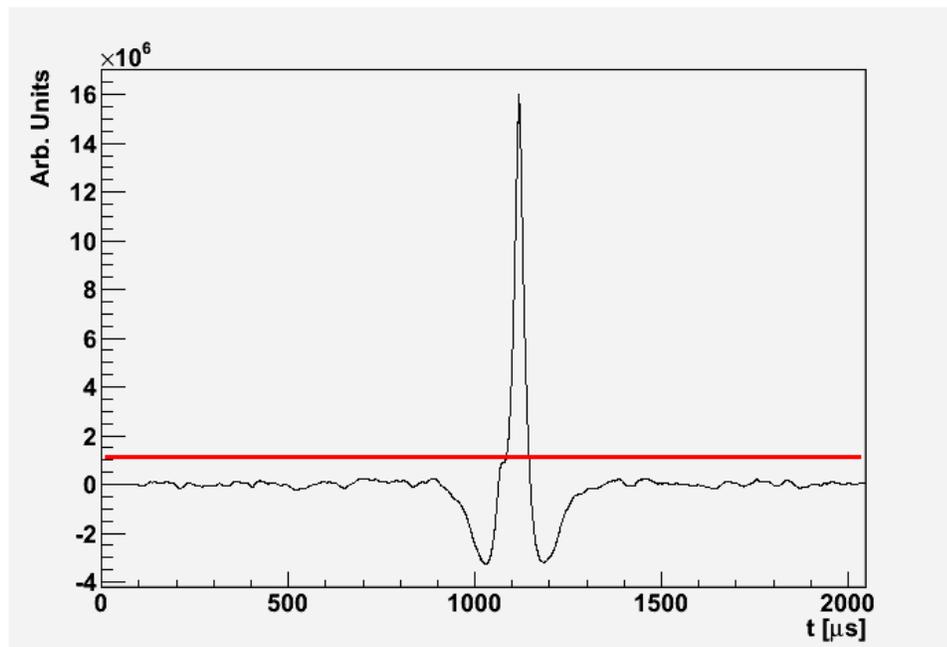
4.2.3 Unshaping

Analysis shows that the matched filter frequently does not find small signals that are near in time to a larger signal in the same waveform. Figure 4.10 shows a waveform 4.10a and the match filter response 4.10b for a frame with two clearly distinct charge deposits. The second, larger signal arrives before the first signal returns to baseline. The response only crosses threshold once, so only one signal is identified. The width of the match filter response is determined by the width of the signal, indicating that the response is too broad to identify signals that arrive within the time-width of the template function.

To achieve better separation of multiple signals, we apply a filter with a narrower time response. Signals are lengthened in time by the electronic shaping, so the effects of this shaping must be removed before a new filter can be applied; hence this method is termed “unshaping”. Unshaping is performed on a $256 \mu\text{s}$ section of



(a)



(b)

Figure 4.10: A waveform (a) with two clear signals that are identified as one signal by the matched filter. The signal threshold is only exceeded once, since the matched filter time response is too broad (b). Unshaping is thus required to properly detect both signals.

an identified signal waveform. The inverse of each differentiator and integrator is sequentially applied to remove the effects of the shaping electronics. A $2 \mu\text{s}$ triangular filter is then applied. Given an unshaped signal $s(t)$, the output $r(t)$ of the triangle filter can be found recursively at each sample time:

$$r(i) = r(i - 1) + s(i) - 2s(i - k) + s(i - 2k), \quad (4.5)$$

where $0 < i < 256 \mu\text{s}$ represents the integer time of the sample and $k = 2\mu\text{s}$ is the filter width. [68] The responses of signal peaks to this filter are much narrower in time, and signals 3 to 5 μs apart can be discriminated. Signals above 5 times the RMS of the baseline of the re-filtered signal are kept as candidates to be treated separately during signal clustering.

4.3 Parameter Estimation

Energy and timing information are extracted from identified U-wire and APD signals by fitting to the matched filter template signal. The amplitude and the time of peak signal are allowed to float. The fitted values are taken as the *reconstructed energy* and collection time of the signal. We use the notation Q_{rec}^{wire} and S_{rec}^{APD} for the charge and light signal reconstructed energy, respectively. The timing of charge signals is determined entirely this way, and V-wire signal times are taken from the associated U-wire collection time. V-wire signals are not fitted, since the induction signals are much weaker than the collection signals.

Position information is determined from the signal timing and from the loca-

tion of the wires with identified signals. Coordinates in the the plane of the wires are defined, with the U- and V-axes defined parallel to the respective wires. The U- and V-coordinates of a signals are found from the coordinates of the nearest wire channels with signals above threshold. In the case of signals on multiple wires, each signal is weighted according to the signal strength. The U-position is weighted by the amplitude of the signal, while the V-position weighted by the match filter response (FOM). The U- and V-wires are crossed at 60 degrees, so for convenience a coordinate transformation into rectilinear coordinates (X, Y) is made:

$$\begin{aligned} x &= \text{sgn}(z)(v - u) \\ y &= (u + v)/\sqrt{3}, \end{aligned} \tag{4.6}$$

where $\text{sgn}(z) = +1$ for TPC 1 and -1 for TPC 2. With these definitions, the X-axis points horizontally into the TPC, and the Y-axis points vertically up (see also Figure 3.3). The axial position, or Z-position, is determined using the time difference between the charge collection and the associated scintillation signal. The scintillation signal is prompt, so this difference is the drift time of the charge. The drift velocity is measured during calibration runs to be $1.71 \text{ mm}/\mu\text{s}$, so the drift distance can be found.

4.4 Clustering

β and $\beta\beta$ -decays overwhelmingly deposit energy in a localized region, about 3 mm in range. By contrast, γ s from radioactive backgrounds may Compton scatter

and deposit energy in multiple, well-distinguished locations. We will call a localized energy deposit a *charge cluster*; a single ionizing particle may result in one or more clusters. By separating *single-site* events with a single charge cluster from *multi-site* events with more than one cluster, we can identify a large fraction of background γ -decays.

By analogy, we will also define a *scintillation cluster* to consist of all scintillation signals deposited at a given time. These signals are prompt, so all charge clusters in an event will contribute to a single scintillation cluster. It is useful to associate each charge cluster to a scintillation cluster and identify the scintillation cluster with the ionizing event. A single DAQ frame may contain multiple scintillation clusters. Reconstruction identifies clusters and associates charge with scintillation clusters.

Charge cluster identification is driven by the 3 mm cluster size of a β deposit. Since this is about one third of the collection wire pitch, most charge collection signals will occur on a single wire or two adjacent wires. Charge clusters are hence restricted to only consist of signals from a single wire or two adjacent wires. All signals in a charge cluster must also be collected within a clustering time of $t_{cl} = 3.5 \mu\text{s}$; this time is chosen to exclude signals from depositions in near time coincidence while ensuring all signal due to a given deposit is included. V-wire signals coincident with the U-wire signals are associated into the charge cluster, as well.

Scintillation clusters consists of light signals occurring within $6 \mu\text{s}$ of one another. This clustering time is the sum of the electronic integrator times. A charge cluster is associated with a scintillation cluster if the charge clusters occurs up to

$t_{d_{max}} = 117 \mu s$ after the scintillation clusters. The time $t_{d_{max}}$ is the maximum drift time of the charge in a half-TPC.

The raw *reconstructed energy* associated with a cluster is taken as the sum of the reconstructed energies of each associated signal. Thus

$$E_Q^{recon} = \sum Q_{rec}^{wire} \text{ and} \quad (4.7)$$

$$E_S^{recon} = \sum S_{rec}^{APD}, \quad (4.8)$$

where E_Q^{recon} and E_S^{recon} are the raw charge and scintillation cluster reconstructed energies, respectively.

Chapter 5

Monte Carlo

Accurate simulations of interactions in the detector are required for generating energy-based probability distribution functions (PDFs) of physics and background processes. These PDFs will be used to make binned maximum likelihood fits to extract process rates, as described in Chapter 7. The simulations will also be used for extracting detection efficiencies for signal and background processes.

In this chapter we describe the generation of simulated data, and some strengths and weaknesses of the method. A detailed detector model is implemented in GEANT4 [69], and decays can be accurately simulated in major components. Energy deposits are digitized to simulate voltage waveforms on the wires and APDs. Only charge spectra (and not scintillation) are accurately generated, and the detector model contains known omissions, such as a model of recombination. The charge spectra will be used to create the required PDFs.

5.1 EXOsim

EXOsim, software based on the GEANT4 Radioactive Decay Module, was developed for simulating nuclear decays within the EXO-200 detector and its environment. The module allows the specification of parent nuclei by atomic number and mass numbers and generates the appropriate α , β , and γ particles. For decaying

nuclei with radioactive daughter particles, particles from the entire decay chain are generated.

Background assays showing the presence of ^{238}U and ^{232}Th make it necessary to simulate these entire chains, which are shown in Figure 5.1. The presence of ^{222}Rn in air makes it an important sub-chain of ^{238}U , and its decay is simulated separately in some cases. The critical backgrounds in these chains are two de-excitation gammas, one 2.448 MeV gamma (1.57% branching ratio) from ^{214}Bi and one 2.615 MeV gamma (36% branching ratio) from ^{208}Tl , respectively in the $^{238}\text{U}/^{222}\text{Rn}$ chain and the ^{232}Th chain. The 2.448 MeV gamma is rare, but extremely close to the 2.458 MeV Q-value for $\beta\beta$.

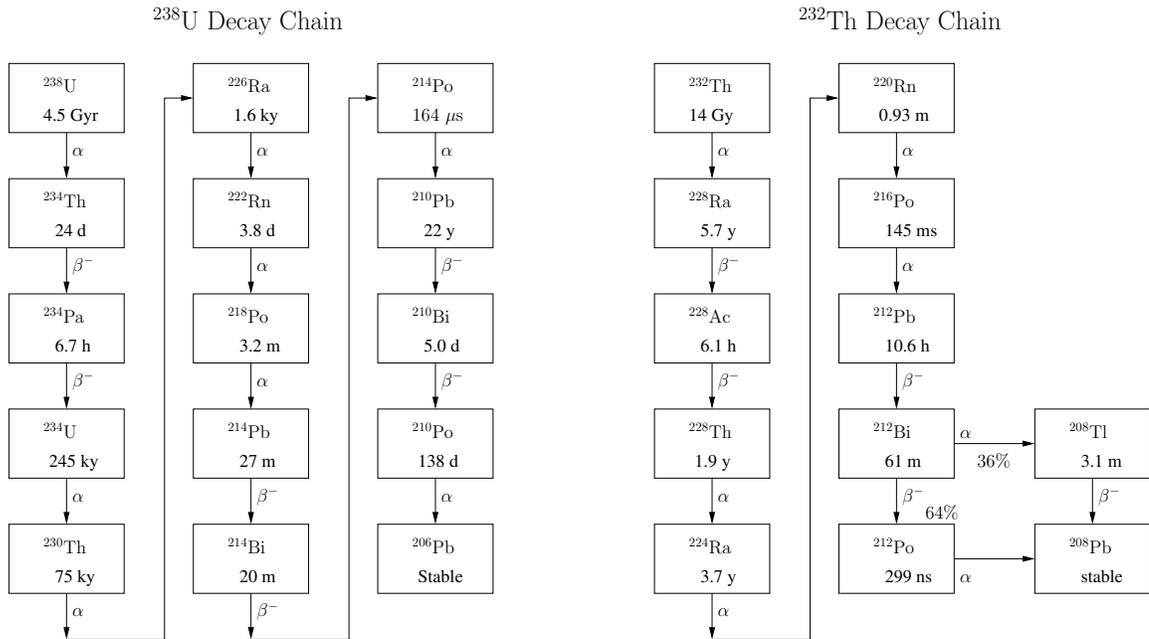


Figure 5.1: Radioactive decay chains for ^{238}U and ^{232}Th . The daughters ^{214}Bi and ^{208}Tl emit de-excitation gammas with energies near the $\beta\beta$ Q-value.

Generators for $\beta\beta$ are not included in GEANT4, so custom generators were added to EXOsim. These include $\beta\beta 0\nu$ and $\beta\beta 2\nu$, and $\beta\beta 0\nu$ with Majoron emission

[70] of ^{136}Xe and ^{134}Xe as well as the analogous double positron decays of ^{124}Xe . Decays to the first excited daughter state [71] of ^{136}Xe and ^{134}Xe can also be generated. Only the generators for $\beta\beta 0\nu$ and $\beta\beta 2\nu$ decays of ^{136}Xe were used in this analysis; these generators are described in Section 5.1.2.

Once the energies of the decay particles have been identified, GEANT4 tracks the particles until energy is deposited in the liquid xenon, accounting for the interaction cross-sections with the various parts of the detector. The position and energy of the deposits are recorded. Total gamma attenuation constants and lengths of the major detector components are shown in Figure 5.2 [72]. Near 2.5 MeV, the $\beta\beta 0\nu$ region of interest, the interaction cross section of most materials is similar and dominated by Compton scattering. For example, Figure 5.3 shows the attenuations lengths for xenon due to Compton scattering, the photoelectric effect, and pair-production; Compton scattering accounts for 85% of the total. Identifying Compton scatter events is thus a crucial aspect of discriminating background gamma radiation from physics signals. It is also important that the material attenuation lengths are generally within an order of magnitude of the others. This low variation reduces the model sensitivity to inaccuracies or missing features. .

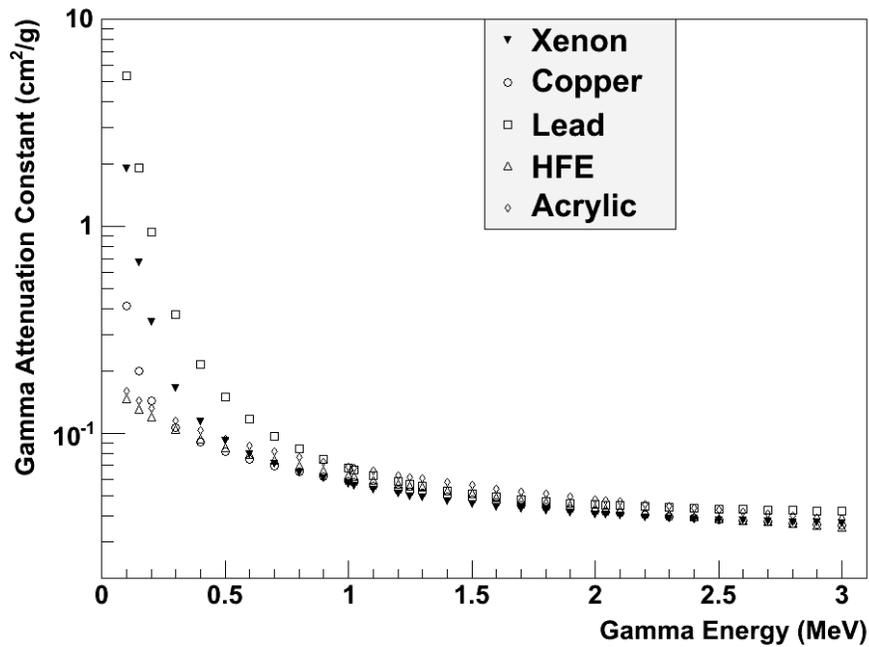
5.1.1 Pixelation

Energy signals will be used to calculate the shape of signals seen by the wires and APDs. To facilitate this, energy deposits are associated into *pixelated charge deposits* (PCDs). Each PCD consists of energy deposits that occur within a cube

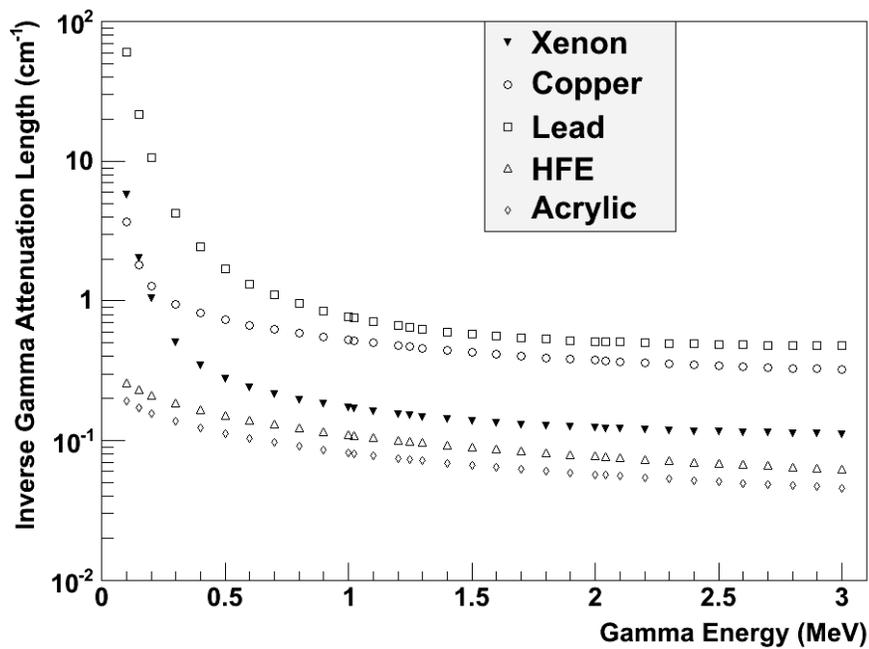
of 0.15 mm and within 0.5 μs of one another. Each PCD is much smaller than a charge cluster, and a single ionizing event typically creates many PCDs. The sum energy of all PCD's associated with an event give the energy, E_{PCD} , which will be used to form the MC energy spectra. Note that since there is no recombination model, E_{PCD} is simply the total deposited energy, containing no information about the division of energy into ionization and scintillation.

5.1.2 $\beta\beta$ generators

Each $\beta\beta$ generator simulates two β particles with final state kinetic energies T_1 and T_2 , chosen at random from the appropriate energy spectrum. The spectrum is included as a differential decay rate, $d\Gamma$, a function of $K = T_1 + T_2$, $D = T_1 - T_2$, and the Q-value of the decay, T_0 . Angular correlations are integrated out for the purposes of selecting final state energies. For each simulated decay, a value for K and D is chosen. For 0ν decays, K is fixed at T_0 . Values for D are generated by applying the Acceptance-Rejection Monte-Carlo technique to the energy spectrum. That is, a random value for \tilde{D} is generated in the entire range of D , from $-K$ to K , and a random value $\tilde{\Gamma}$ for $d\Gamma$ is generated below the maximum value of the function. If the random value of $\tilde{\Gamma}$ is less than $d\Gamma(\tilde{D})$, then \tilde{D} is simulated; otherwise, the value is rejected. This procedure ensures that the distribution of simulated values of D approximates the analytic spectrum. For 2ν decays, the sum energy is not fixed and must also be chosen by Acceptance-Rejection Monte Carlo. A sum spectrum, $\frac{dN}{dK}$, is formed by numerically integrating $d\Gamma$ over the entire range of D . Then random



(a) Gamma attenuation constants in major detector components, as a function of gamma energy. The differences between the constants are small in the $\beta\beta\nu$ region of interest near 2.5 MeV.



(b) Gamma attenuation lengths for major detector materials. Near 2 MeV, most materials have an attenuation length between 1 and 10 cm.

Figure 5.2: Gamma attenuation properties of EXO-200 materials.

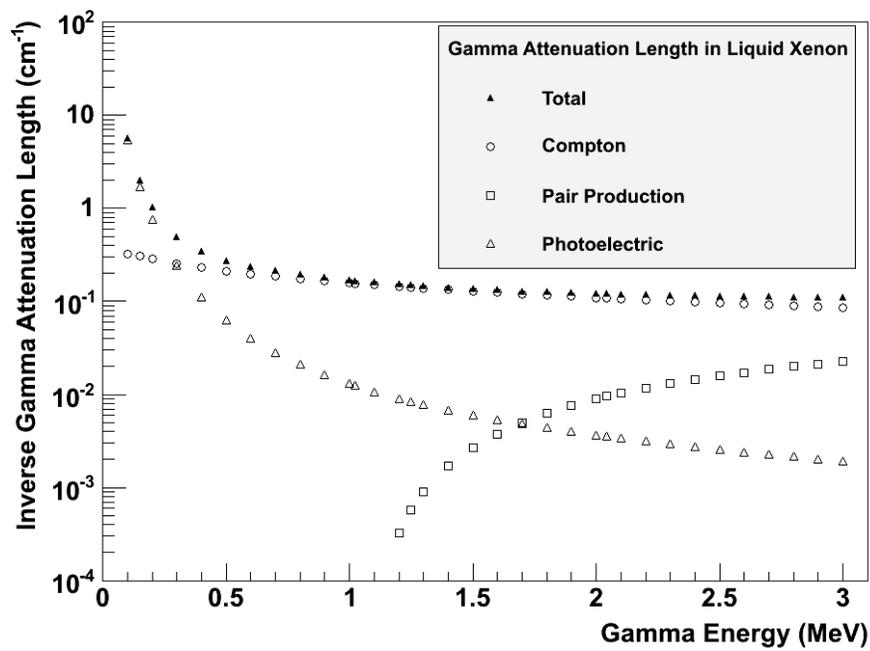


Figure 5.3: Gamma attenuation lengths in xenon for energy-depositing processes. At the $\beta\beta\nu$ Q-value of 2458 MeV, Compton scattering is the dominant process, contributing 85% to the total attenuation length; pair production and photoelectric interactions account for 13% and 2%, respectively.

values for K and $\frac{dN}{dK}$, \tilde{K} and \tilde{N} , are chosen and kept if $\frac{dN}{dK}(K)$ is less than \tilde{N} .

The spectral functions $d\Gamma$ for $\beta\beta 0\nu$ and $\beta\beta 2\nu$ are

$$d\Gamma = \begin{cases} F(E_1, Z)F(E_2, Z)p_1p_2E_1E_2 & (\beta\beta 0\nu) \\ F(E_1, Z)F(E_2, Z)p_1p_2E_1E_2(T_0 - K)^5 & (\beta\beta 2\nu), \end{cases} \quad (5.1)$$

where p_i and E_i are the final state momenta and total energy. $F(E, Z)$ is the Fermi function accounting for Coulomb interaction between a nuclear decay daughter ion and the final state β . It depends on the β energy, E , and momentum, p , as well as the daughter atomic number, Z . The exact form of the Fermi function is given by

$$F(E, Z) = 2(1 + \gamma)(2pR/\hbar)^{-2(1-\gamma)}e^{\pi\nu}\frac{|\Gamma(\gamma + i\nu)|^2}{\Gamma(2\gamma + 1)^2}, \quad (5.2)$$

with

$$\nu = Ze^2/\hbar v, \gamma = [1 - (Z\alpha)^2]^{1/2}, \quad (5.3)$$

where α is the fine structure constant, Γ is now the gamma function, and R is the final state nuclear radius. In practice, calculating this expression is computationally expensive. EXOsim uses an approximation from Schenter and Vogel to find the Fermi function [73], with m the electron mass:

$$F(E, Z) = \frac{E}{p}\exp[\alpha(Z) + \beta(Z)(E/m - 1)^{1/2}], \quad (5.4)$$

where for $E/m < 2.2$,

$$\begin{aligned}\alpha(Z) &= -0.811 + 4.46 \times 10^{-2}Z + 1.08 \times 10^{-4}Z^2 \\ \beta(Z) &= 0.673 - 1.82 \times 10^{-2}Z + 6.38 \times 10^{-5}Z^2\end{aligned}\tag{5.5}$$

and for $E/m > 2.2$

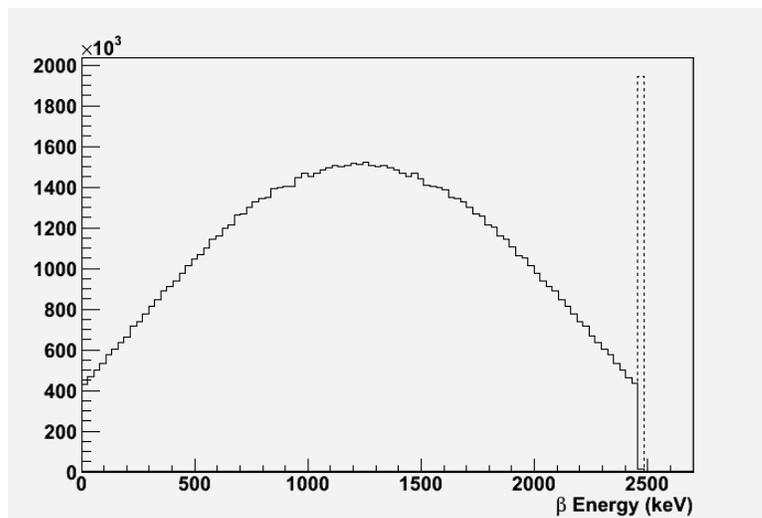
$$\begin{aligned}\alpha(Z) &= -8.46 \times 10^{-2} + 2.48 \times 10^{-2}Z + 2.37 \times 10^{-4}Z^2 \\ \beta(Z) &= 1.15 \times 10^2 + 3.58 \times 10^{-4}Z - 6.17 \times 10^{-5}Z^2.\end{aligned}\tag{5.6}$$

The sum and individual β spectra generated for $\beta\beta 0\nu$ and $\beta\beta 2\nu$ are shown in Figure 5.4.

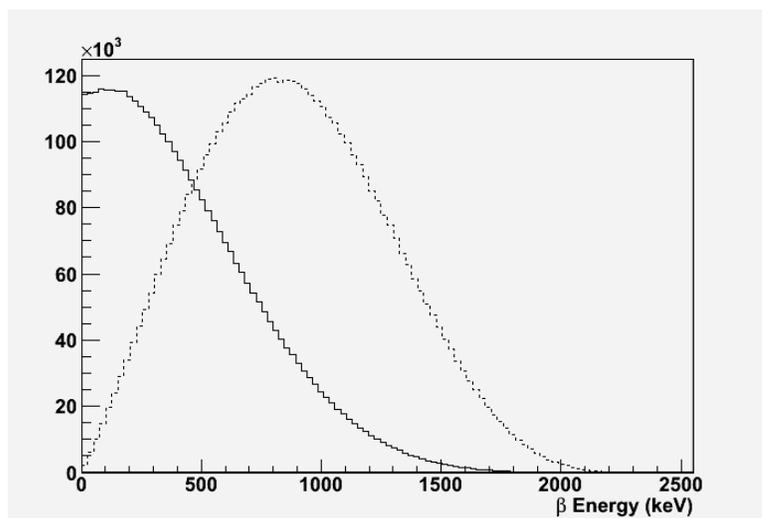
5.1.3 Detector model

The GEANT4 detector model includes most detector features. A list of the parts in the model is shown in Table 5.1. It is critical that components comprising or near the TPC are modeled accurately, as well as the massive components such as the HFE, lead, and cryostat. The air surrounding the detector is also modeled so we can account for Rn accumulation between the lead wall and the cryostat. However, several smaller components have some inaccuracies:

- The Teflon reflectors are modeled as a single cylinder, as opposed to flat, slightly overlapping tiles. A gap in the tiles near the cathode is filled with Teflon in the simulation.



(a) $\beta\beta 0\nu$



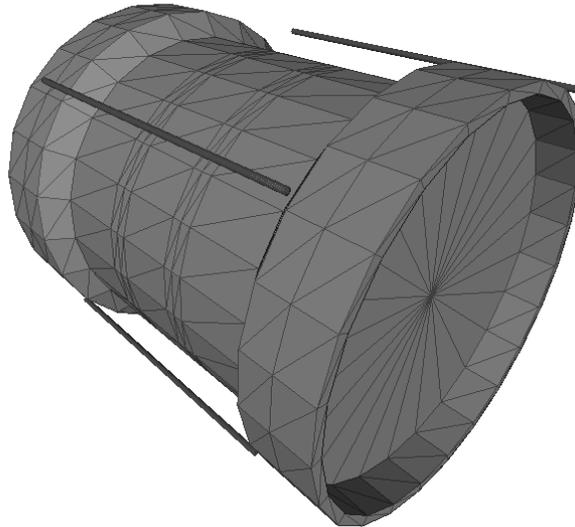
(b) $\beta\beta 2\nu$

Figure 5.4: Simulated single (solid) and sum (dashed) β spectra for $\beta\beta 0\nu$ and $\beta\beta 2\nu$. The single β spectra are arbitrarily scaled for viewing.

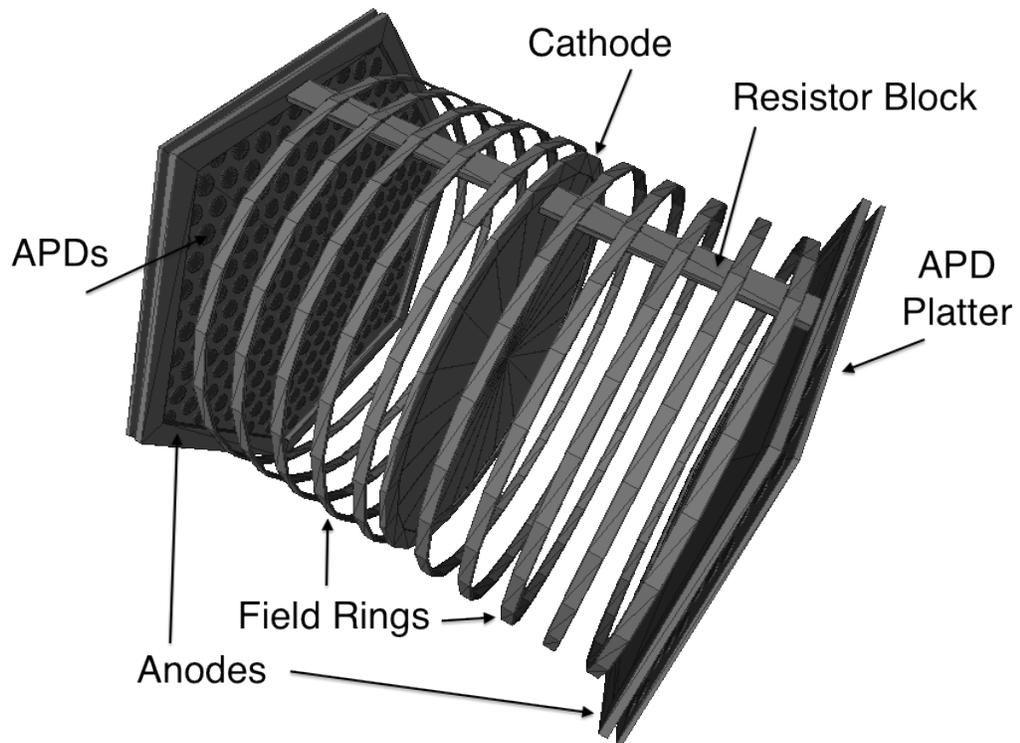
- Some features on the liquid xenon vessel are missing: the four external legs, four external ribs, and the star pattern of ribs on the endcaps. See Figure 5.5a.
- Parts of the calibration tubing that curve around the anodes are omitted, see again Figure 5.5a.

Additionally, some small parts are not included in the simulation. These include twelve acrylic holders for the field shaping rings, 1.27 cm x 0.75 cm in cross section, the thin Kapton cables connecting to the cathode ring, and the high voltage cable external to the vessel.

None of these inaccuracies should significantly affect the simulated energy deposited, nor should they contribute significant errors in the simulated EXO-200 background rates. The lack of Teflon overlap minimally affects the effective thickness. The calibration tubing is external to the TPC and it is quite small compared to the mass of the vessel. The vessel features omitted comprise a small fraction of the detector mass, and are on the outermost portions. Accumulated, the acrylic holders and the Kapton cables each have a mass of about 100 g, and the mass of the high voltage cable is about 200 g. Compared to 200 kg of the xenon, these parts will not contribute significantly to the charge or light attenuation. Also, these parts are among the cleanest components in the detector (see Table 3.4), so they will contribute negligibly to the background spectrum. It will be seen below that good agreement between simulation and data can be achieved for calibration sources, validating these claims.



(a) The exterior of the copper TPC as rendered in the GEANT4 simulation. Comparison to 3.4 shows that the endcap “spoked” features and the input/output “legs” are not reproduced in the detector model. The external tubes are the model of the calibration tubing, which omits portions near the endcaps.



(b) The interior of the TPC as rendered in the GEANT4 simulation. The teflon tiles and half of the (non-adjacent) field shaping rings are suppressed to show the interior.

Figure 5.5: GEANT4 models of the detector.

Table 5.1: Detector components included in GEANT4 model, with material and properties used for simulation. The operating temperature of the detector is taken to be 173.15 K. Acrylic refers to PMMA $((C_5O_2H_8)_n)$, and air is taken to be 70% N, 30% O.

| Component | Material | Density (g/cm ³) |
|----------------------------------|---------------|------------------------------|
| Surrounding Air | Air | 1.29×10^{-3} |
| Lead shield | Pb | 11.35 |
| Inner and outer cryostat | Cu | 8.96 |
| Cryostat insulating vacuum | H | 10^{-25} |
| HFE-7000 | $C_3F_7OCH_3$ | 1.77 |
| Calibration source tube body | Cu | 8.96 |
| Calibration source tube interior | Air | 1.29×10^{-3} |
| Liquid xenon vessel | Cu | 8.96 |
| Liquid xenon | Xe | 3.023 |
| APD platters | Cu | 8.96 |
| APDs | Si | 2.33 |
| Teflon reflector tiles | PTFE | 2.15 |
| Resistor blocks | Acrylic | 1.19 |
| Anode wire support beams | Acrylic | 1.19 |
| Cathode ring | Cu | 8.96 |
| Field shaping rings | Cu | 8.96 |

5.2 Digitization

Digitized waveforms for wire and APD signals are created from the charge pixels deposited in the detector by GEANT4. The procedure is similar in either case. First raw waveforms are generated with a fine sampling time of $0.05 \mu\text{s}$. For wire signals, these are calculated using Ramo's Theorem; the raw APD signal is taken as a step function. The signals are then resampled with a $1 \mu\text{s}$ sampling time to match the output of the physical DAQ. Noise sampled from calibration data is added to the waveform so that the reconstruction matched filter will yield realistic efficiencies. Lastly, the waveforms are scaled from energy units into ADC counts using fixed factors.

5.2.1 APD digitization

For each PCD, the total energy, E_{PCD} , is used to determine the number of photons reaching each APD channel. The nominal scintillation yield is taken to be $Y = 14,000$ photons/MeV, representing the approximate average number of photons generated for a given energy of deposit. The photons are not tracked in each simulation, for reasons discussed in Section 5.2.3. Instead, a single simulation of 10^6 5 MeV alpha decays generated uniformly in the detector was made with the photons being tracked until absorption [74]. The number of photons striking each APD plane was recorded as a function of deposit position, and the ratio of photons collected to photons deposited was used to define an acceptance function, $A(r, z)$. The acceptance varies between 1% and 14%.

The function, $A(r, z)$ is piecewise defined in six ranges of radius; in each range, the dependence on z is approximated by two polynomials, one for each half of the TPC. The acceptance is shown for each radial range in Figure 5.6. These curves show the acceptance on the APDs in TPC 1 ($+z$); for deposits in TPC 2 ($-z$) the mirror of the curve is used. The acceptance generally has a step near $z = 0$ due to shadowing effects from the cathode wires and support ring. At large radius, the anode also has shadowing effects near $|z| = 190$.

Optical properties of the detector were implemented in the detector model to ensure the accuracy of the acceptance function. Acrylic and copper components were modeled as totally absorbing; the anodes and cathode were taken to have 8.3% and 10% opacity, respectively, due to the fractional area filled by phosphor-bronze wires. Properties for other materials are listed in Table 5.2.

Finally, fluctuations in the ratio of charge to light are simulated with a random factor, F , taken from a gaussian with width $\sqrt{h_i}$; $h_i = Y A(r, z) E$ is the number of photons hitting an APD plane without the fluctuation factor. These fluctuations are known to be non-Poissonian, but the energy spectrum of the simulated scintillation is not used in any analysis. All together, the number of photon hits on the plane, H , is given by

$$H = F \times Y \times A(r, z) \times E. \quad (5.7)$$

The photons are assumed to be distributed evenly over an entire plane, so the number of hits on a particular APD channel is taken to be $H/37$, where 37 is the number of channels per plane.

For each APD channel with a hit, a raw waveform is created. The photon signal is a step function at the hit time, rising from 0 to the number of hits. Multiple signals can be superimposed if multiple signals occur in the same frame. The signal is then shaped with an electronics transfer function, determined by the shaping times of the DAQ. Real noise is added to the signal. The signal is resampled and scaled into ADC counts, and a baseline of 1664 ADC counts is added.

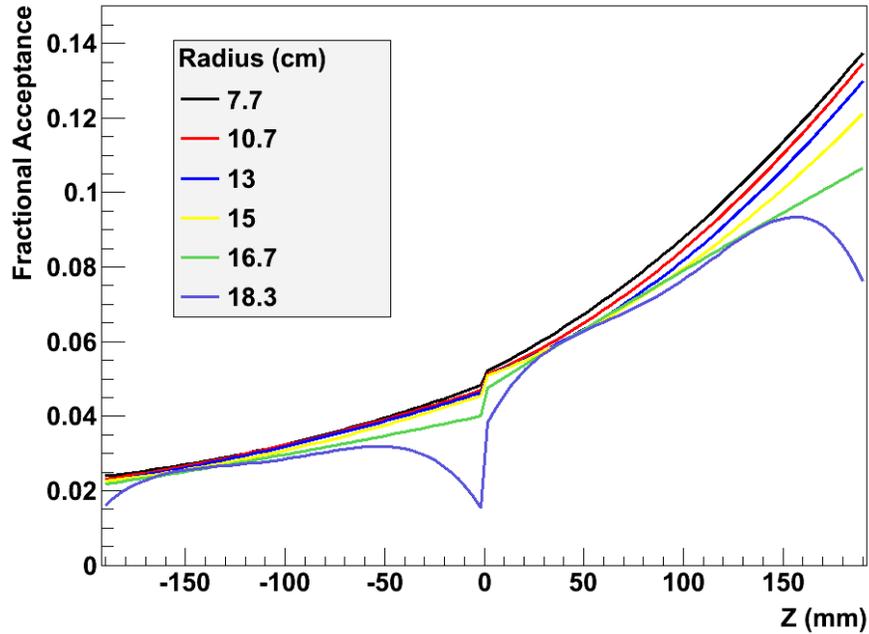


Figure 5.6: Simulated APD collection efficiency $A(r, z)$ as a function of position. Each curve corresponds to a range of radii, with the largest radius given by the legend, and the smallest radius given by the boundary of the next smaller range (or 0). The acceptance for APDs in TPC 1 (+z) is shown. The step near $z = 0$ is due to shadowing effects from the cathode.

5.2.2 Wire digitization

Waveforms on wires are generated by calculating the current induced using Ramo's Theorem [75]. A moving charge of magnitude q_m with velocity \vec{v} induces a

Table 5.2: Optical properties of detector components used for photon propagation in simulation. Values are quoted for 177 nm (7 ev) photons. *Complex refractive index is used to model reflectivity.

| Component | Material | Refractive Index | Total Reflectivity |
|--------------|---------------------------------|-----------------------|--------------------|
| Liquid Xenon | Xe | 1.6 | N/A |
| Reflectors | PTFE | 1.4 | 0.74 |
| APD platter | Al and MgF ₂ coating | N/A | 0.85 |
| APDs | Si | 0.68 + <i>i</i> 1.68* | N/A |

current $i(t)$ on an electrode of interest:

$$i(t) = -q_m \times \vec{E}_w \cdot \vec{v}. \quad (5.8)$$

The *weighting field*, \vec{E}_w , is an electric field that represents the response of the electrode to the presence of charge in the detector; it is the field that results from setting the potential on the electrode of interest to 1 and the potential on all other electrodes to 0. Then the charge $Q(t)$ induced as the charge travels from \vec{X} to $\vec{X} + d\vec{X}$ is then the integral of $i(t)$,

$$Q(t) = -q_m \times (V_w(\vec{X} + d\vec{X}) - V_w(\vec{X})), \quad (5.9)$$

where V_w is the electric potential corresponding to \vec{E}_w . In general, the charge signal is bipolar, as the weight field points in different directions during the charge trajectory. This is consistent with the fact that the integrated current must be 0 for a pure induction signal. For signals collected on a wire, the charge signal is determined by Ramo's Theorem while the charge is drifting, rises to the total deposited charge at the time of collection, and stays constant. Collection signals are

therefore monopolar and integrate to the total collected charge.

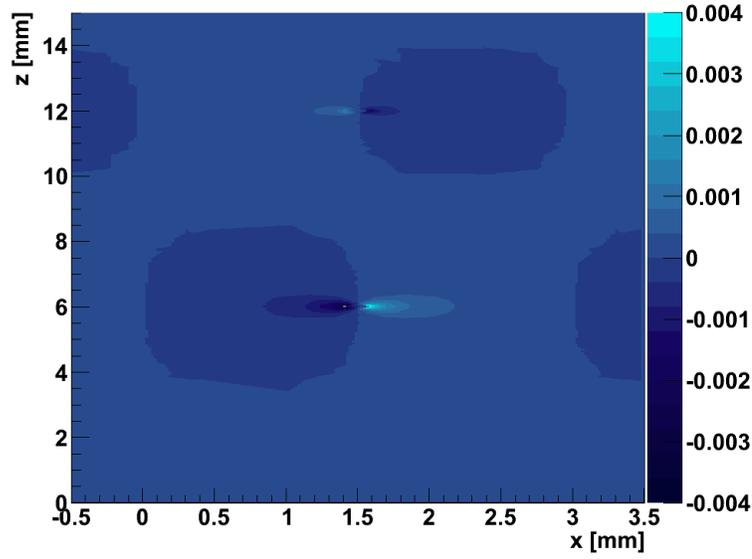
For EXO-200, induced waveforms are calculated for each PCD, and all the waveforms in a frame are then summed together. As the charge deposited is proportional to the event energy, we use E_{PCD} rather than q_m to find the signal. To calculate the waveform, first the nearest U-wire and V-wire to the PCD are identified. Signals are found for these channels and the nearest adjacent channels on either side. A two-dimensional electric field has been calculated in MAXWELL for finding the charge trajectory. The calculation assumes cylindrical symmetry, and the results are shown in Figure 5.7. This field is then used to calculate the charge trajectory in intervals of the stepping time, $dt = 0.05 \mu\text{s}$:

$$d\vec{X} = -\hat{E}(z, x) \times V(z) \times dt, \quad (5.10)$$

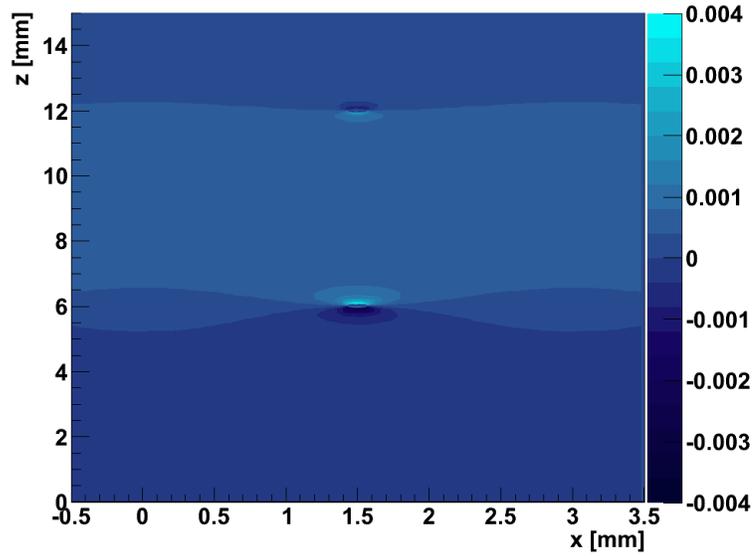
where $d\vec{X}$ is the change in the charge position and \hat{E} is the normalized electric field. The drift velocity $V(z)$ is different in the collection region between the wires, due to the potential on the V-wires; it is taken to be

$$V(z) = \begin{cases} 1.71 \text{ mm}/\mu\text{s} & 0 < |z| < 192 \text{ mm}; \text{ Bulk Drift} \\ 2.25 \text{ mm}/\mu\text{s} & 192 < |z| < 198 \text{ mm}; \text{ Collection} \end{cases} \quad (5.11)$$

The electric field was calculated using voltages different than those used during physics runs: the V-wires and cathode are set to -4.35 kV and -62.5 kV as opposed to -745 V, -8 kV, respectively. This potentially affects the charge drift velocity, and



(a) Drift field x-component



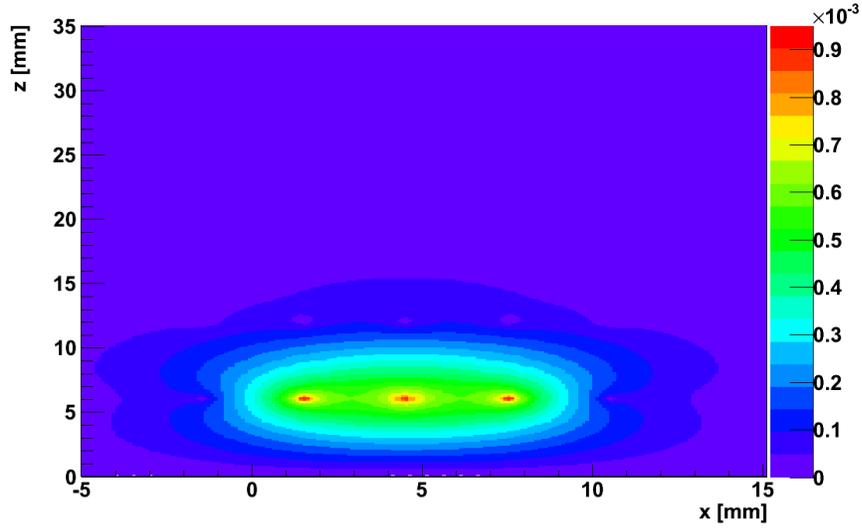
(b) Drift field z-component

Figure 5.7: Components of the TPC electric field in the vicinity of the anode. The anode is viewed edge-on, so that the axis of the TPC (z -direction) is vertical. Components of the field in the x and z directions are shown. Positive field component values indicating fields in the direction of negative x and negative z on the respective figures. In general, the field is directed away from the V-wires and towards the U-wires.

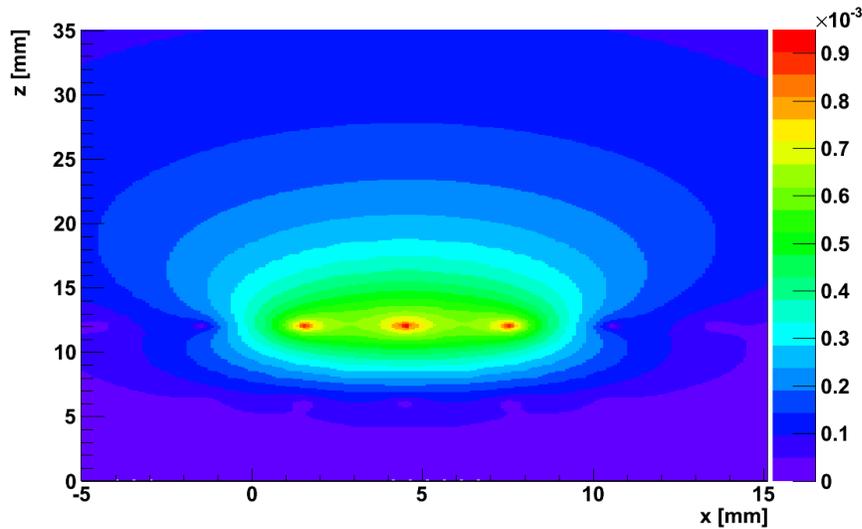
the amplitude of collection signals could be affected if the amount of charge collected on the V-wires to U-wires increases. However, the drift velocity is measured and applied separately, and is not taken from the field calculation. Furthermore, from reference [76], Eqn. (20), the V-wires are fully transparent to charge if the ratio of collection electric field to bulk electric field is larger than $(1 + \rho)/(1 - \rho)$, where $\rho = 2\pi r/d$, with r the wire radius and d the wire pitch. For the EXO-200 TPC, this threshold is 1.31, while calculations show that the field ratio for both the simulated and implemented voltages are both ~ 2 , well into the fully transparent regime.

The bulk drift velocity is measured during calibration runs, but the collection drift velocity is taken from calculations, so induction effects near the wires may not be properly simulated. At each sample time t_{samp} (i.e., charge position \vec{X}), the charge signal $Q(t_{samp})$ is calculated from Equation 5.9, taking q_m to be the charge deposited from GEANT4. The weight field of a wire gang is calculated assuming symmetry along the axis of the wires. A 2-dimensional potential is calculated using MAXWELL for event adjacent U-wires and the nearest eleven V-wires. Thus, for both types three gangs of 3 wires plus one additional wire on either side are simulated. The central U- or V-gang is set to potential 1 and the other gangs set to 0. The weighting potential, V_w , corresponding to the weight field is shown for U- and V-wires in Figure 5.8, viewing the wires edge-on.

For use in the digitizer, this potential is translated so that the central gang aligns with the electrode of interest. This method relies on the approximation that non-adjacent gangs have little effect, which is valid since the nearest neglected gangs are at least four wire pitches away from the central gang.



(a) U-wires.



(b) V-wires

Figure 5.8: Weighting field potential used to calculate charge waveforms from simulated energy deposits. A 2-dimensional potential is calculated for U-wires 5.8a and V-wires 5.8b; the wires are viewed edge-on in this plot. Eleven U-wires and eleven V-wires are simulated. The potential on the central gang is set to a unit value, and the color plot indicates potential relative to that value. The central gang is seen in red on each plot, and the nearest gang above or below can be seen at 0 potential. Wires are spaced at 3 mm intervals from $x = -11.5$ mm to $x = 19.5$ mm, with the U-wires at $z = 6$ mm and the V-wires $z = 12$ mm.

The signal generated is added to the waveform of all signals on that channel from a given event. The raw charge waveforms are then shaped with the electronics shaping transfer function for the wires. The shaped signals are resampled, and noise is added. Finally the signals are scaled into ADC counts by a factor of $\frac{1}{300} \times \frac{1}{18.7}$, where 300 is the number of electrons read at full scale and 18.7 eV is the nominal energy required to liberate a single electron in liquid xenon.

5.2.3 Difficulties using scintillation information

A simulated photon collection efficiency is used in APD digitization, since accurately simulating the propagation of scintillation in the TPC raises several difficulties. Primarily, photon tracking is computationally expensive, and we require the efficient production of several dozen MC spectra. Furthermore, as described in Section 3.3.2, the reflectance of the interior detector materials, primarily the Teflon tiles surrounding the field cage, is not precisely known. Finally, no model of charge recombination is available that correctly models the enhanced anticorrelation between charge and scintillation. Thus any relative scaling between the charge and light energies will be suspect.

5.3 Unused Monte Carlo capabilities

Several features are implemented in the simulation code, but are not used in this analysis. First, the digitizer can reproduce the effects of charge attenuation during drift due to capture on electronegative impurities in the liquid xenon. The

correction is not made, since in the analysis below, the purity of the xenon is not known before data taking and varies from run to run. It is more convenient to instead correct for the attenuation in data and compare to simulated data that has not been attenuated. Second, the digitizer can simulate the effects of a non-zero charge energy resolution by scaling the charge energy with a random gaussian-distributed factor. However, it will be seen that the detector energy resolution is dependent on the deposit energy, and measurement uncertainty in its parameterization contributes to the systematic uncertainty of the lifetime measurement. Thus, the parameterization must be included as fitted parameters, and the simulated data must not be modified for detector resolution effect until fits are made.

5.4 Notes on PDF generation

The simulations generate spectra of the total energy, E_{PCD} , deposited in the liquid xenon. These spectra assume perfect energy resolution, so before being used as PDFs, they must be convolved with a detector energy resolution function. This function is energy dependent, and must be determined during calibration. Furthermore, the region of the TPC in which the electric field is uniform is not modeled in the simulation. To account for this, the uniform region must be measured, and PCDs outside must be excluded from the total energy. For these reasons, we delay discussion of PDF generation until Section 7.4.

Chapter 6

Calibration

The low-background data will ultimately be fit to energy-based PDFs generated using MC data sets. Since the MC data is generated using the true energy deposited in the detector, a precise calibration of the TPC energy response is required, described below. Fits will be discussed in Chapter 7.

The gamma sources in listed in Table 3.1 are used to characterize the detector. First, several position-based or electronics-based variation in the energy scale are identified, and corrections are developed to homogenize the detector response. These corrections are as follows:

- Electron lifetime (z-dependent attenuation of charge)
- Shielding grid inefficiency (z-dependent reduction in charge signal)
- Relative U-wire response
- Three-dimensional scintillation response function

Next, we determine the optimal way to combine the charge and light spectra, leading to a *rotated energy* spectrum. We characterize the *rotated energy* scale and energy resolution as a function of energy. This is essential to ensure that the PDFs have the correct spectral shape. To separate signal β -like events from background γ s, two spectra are made, one for single-site and one for multiple-site events. Each

is calibrated separately to account for possible effects of the event topology on the energy scale.

6.1 Calibration Runs

Determine energy corrections and data selection criteria, large number of events are required to reduce statistical uncertainties. To achieve this, dedicated runs were made prior to the low-background campaign. At least one run with each of the strong calibration sources was taken at each of the five source locations (see Table 3.2). This ensures sensitivity to position dependent effects. Strong sources runs use only a 50 Hz random DAQ trigger to avoid dead time associated with large numbers of triggers in a single frame.

In addition, daily runs with the weak sources are used to check detector stability, obtaining about 250,000 DAQ triggers per run. The position used is varied to monitor position dependent time variations. Weak source runs match the physics trigger, which triggers on one of the following: APD sum signal (threshold 800 ADC counts, baseline subtracted); APD individual signal (threshold 1750 with no baseline subtraction), a single U-wire signal (1617 ADC counts with no baseline subtraction), or a 0.1 Hz random trigger.

6.2 U-Wire response

The U-wires exhibit non-uniform responses due to variances in the components of the DAQ electronics. For consistent interpretations of the energy signal, these

gains must be normalized. It is convenient to use the pair production double escape peak of ^{228}Th . In pair production, an incoming gamma with energy greater than $2m_e = 1.02$ MeV is converted into an electron-positron pair, with the excess energy going into kinetic energy of the pair. The positron comes to rest and annihilates with an atomic electron to create two photons of energy $m_e = 0.511$ MeV. These photons typically deposit energy elsewhere in the detector, since the attenuation length of a 511 keV γ in LXe is only 4 cm. The remaining energy of 1593 keV is deposited as a single cluster in the detector by the created electron (which is collected) and the loss of kinetic energy of the positron. This should lead to a peak in the multi-site spectrum; we call this the *double-escape* peak, since even though the photons may not truly “escape” the detector, the ones that do not can be distinguished as belonging to a different charge cluster. The photons may escape the detector, leaving only a single-site event, but due to the size of the detector this is rare. Consequently, the double-escape peak is much less significant above the background in the single-site spectrum than in the multi-site. Crucially, all of the energy in the double escape peak is deposited by β s, making it an ideal calibration peak for $\beta\beta$ decay. A similar process in which only one photon escapes leads to a single escape peak at 2104 keV, but this peak is not used for calibration since it mixes β and γ deposits. For the relative U-wire response correction, a single calibration point suffices in any case. Both of these features appear in the multi-site charge cluster spectrum of strong ^{228}Th calibration data, shown in Figure 6.1.

Using the strong Th calibration data, a spectrum of the reconstructed charge

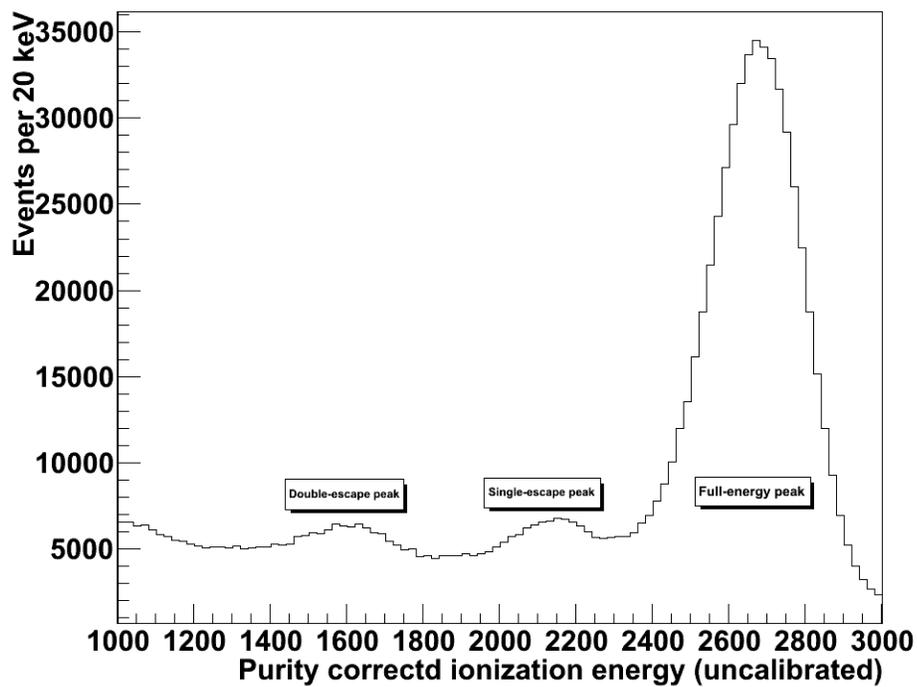


Figure 6.1: ^{228}Th calibration source spectrum of reconstructed charge cluster energy. Only multi-site and single-wire events are included, to ease identification of the double and single escape peaks, indicated at 1593 keV and 2104 keV. The double escape peak is entirely composed of β deposits, making it useful for wire gain calibrations.

cluster energies, E_Q^{recon} , was formed for each U-wire.¹ To make a cleaner escape peak spectrum, only clusters from multi-site events were included; further, the clusters were required to deposit on a single U-wire. The position of the ^{228}Th double escape peak in each spectrum was measured by fitting the peak to the sum of a gaussian and a linear background and extracting the best-fit gaussian mean. Figure 6.2 shows the fit for two adjacent U-Wire channels, 15 and 16, which have a clear discrepancy in peak position. The ratio of the fitted peak energy to the mean is taken as the response correction factor, ξ . Thus the corrected U-wire signal energy, E_{uc}^{wire} , given by

$$Q_{uc}^{wire} = \xi \cdot Q_{rec}^{wire}, \quad (6.1)$$

recalling the reconstructed wire signal energy Q_{rec}^{wire} from Chapter 4. We must re-sum the wire signal energies to find the corrected charge cluster energy, E_{uc} :

$$E_{uc} = \sum \xi \cdot Q_{rec}^{wire}. \quad (6.2)$$

For all channels, ξ is less than 10%, and it is less than 5% for the large majority. As a stability check, the responses were re-measured using strong calibration runs 3416-3428. ξ did not drift by more than 1% from the initial values during the physics run.

¹To achieve sufficient statistics, strong ^{228}Th calibration runs 2417, 2418, 2421, 2422, 2423, 2424, 2426, 2431, 2432, 2433, 2434, 2447, and 2448 were all combined.

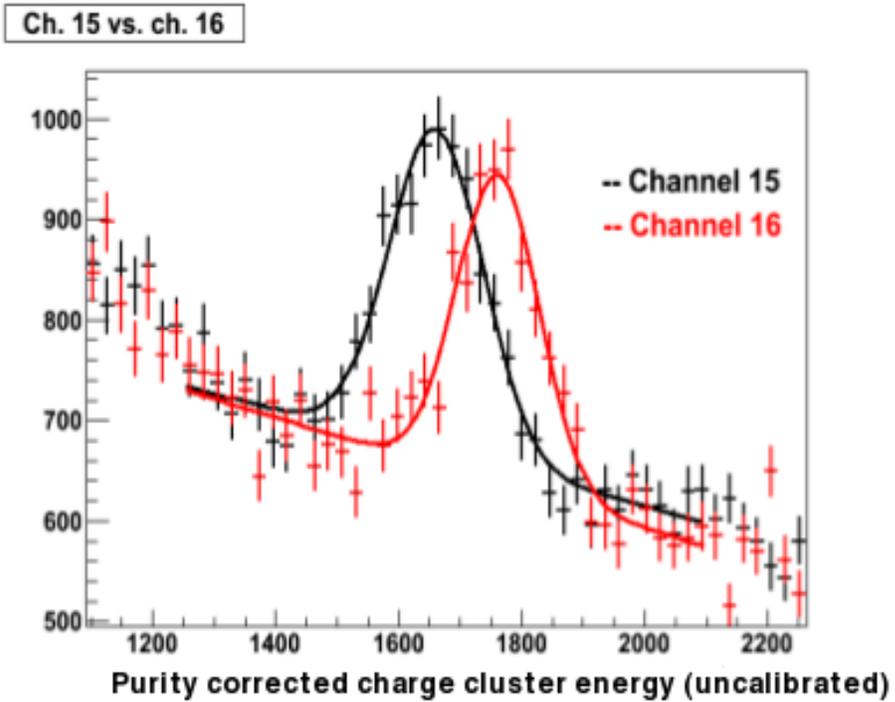


Figure 6.2: Double escape peak in the ^{228}Th multi-site cluster energy spectrum, as measured on two adjacent U-Wire channels, fit to a gaussian with linear background. The observed discrepancy in the peak energy is due to gain differences on the two wires, which are corrected using the results of the fit. Figure courtesy of G. Giroux.

6.3 Electron lifetime

Electronegative impurities in the LXe capture drifting charge, decreasing the charge collection signal. A *free electron lifetime*, τ , can be defined to describe the characteristic drift time until capture. Charge deposits further from the anode drift for longer times, leading to a non-uniform attenuation along the axis of the detector (z-direction). The electric field ensures that the charge motion is predominantly in the z-direction, so we do not consider the effects of charge attenuation due to radial motion. Concentrations as low as 0.1 ppb O₂ can affect the lifetime, so the ability to purify xenon is extremely important for reducing charge loss from this effect.

The amount of charge captured can be described by an attachment rate constant, k . The constant depends on the impurity species and concentration, as well as the electric field strength applied. The largest attachment constant for typical impurities is associated with O₂; for the EXO-200 electric field, $k \sim 10^{11} M^{-1} s^{-1}$ for O₂ [77]. The constant for N₂ is several orders of magnitude lower. The charge captured per unit time is given by

$$dQ(t)/dt = \sum -k_i \times [X_i] \times Q(t), \quad (6.3)$$

where $Q(t)$ is the charge of drifting electrons at time t , and $[X_i]$ and k_i are the concentration and attachment constant of species X_i , respectively. Integrating this, we see that the effect of the electron lifetime is an exponential attenuation of the

charge collection signal

$$Q_{coll} = Q_{dep}e^{-t_d/\tau}, \quad (6.4)$$

where t_d is the drift time, $Q_{dep} = Q(0)$ is the initial charge deposited, and $Q_{coll} = Q(t_d)$ is the charge collected at the anode after the charge has drifted. The observed electron lifetime is then given by

$$\tau_X = \frac{1}{\sum k_i \times [X_i]}, \quad (6.5)$$

so ultimately the attenuation measured may arise from a combination of different impurities. We define a correction factor, κ ,

$$\kappa(t_d) = e^{+t_d/\tau}. \quad (6.6)$$

Applying this correction factor to the charge energy yields an estimate of the original charge deposited:

$$E_{pc} = \kappa E_{uc}, \quad (6.7)$$

where E_{pc} is purity corrected charge energy, and E_{uc} is the output of the previous U-wire response correction. For multi-site events, this correction can be applied to each charge cluster individually:

$$E_{pc} = \sum \kappa(t_d^{(i)}) \cdot E_{uc}^{(i)}, \quad (6.8)$$

with the sum over all charge clusters, indexed by i . The electron lifetime in the

EXO-200 TPC is measured using the full-absorption peak from radioactive source calibration spectra. The full-absorption peak is mono-energetic, but the amount of charge collected will decrease with as the distance of the interaction site from the anode increases, according to Equation 6.4. Two methods are considered, a *peak method* which measures the change in full-absorption peak position with drift time, and a *resolution method* which optimizes the energy resolution of the full-absorption peak by varying the lifetime. Both methods require each event to have a well-defined drift time. Multi-site data do not have a single well-defined drift time, so only single-site data are used.

The impurity concentration can vary slowly in time; measurements of the lifetime during each calibration run are found to characterize the variance with sufficient resolution. The physics run data is then corrected using a polynomial fit to the time variation. For the best purities obtained with EXO-200, this effect is 3% at maximum drift time, corresponding to an lifetime of $\sim 3000 \mu\text{s}$.

6.3.1 Peak method

In the peak method, calibration source events are divided by drift time into 24 even bins of $6.875 \mu\text{s}$, with 12 bins in each TPC half. Events in each bin have a similar drift time and hence a similar attenuation. By finding the mean charge collected in each bin, the attenuation as a function of drift time can be measured. Events are selected that have a radial coordinate $r < 160 \text{ mm}$ and pass a minimum scintillation threshold (2000, 1000, or 400 ADC counts for ^{228}Th , ^{60}Co , or ^{137}Cs ,

respectively). The energies of the events in each bin are histogrammed to form a spectrum, and the full-absorption peak is fit to extract the mean charge signal amplitude in that bin.

The fit function, $F(E)$, shown in Equation 6.11, is the sum of a gaussian and a smoothed step function. These represent the peak and the Compton shelf, respectively.² The latter must be included since the detector resolution is not sufficient to fully separate the peak; it is expressed as a complementary error function (erfc), i.e. the convolution of a step function with a gaussian. The peak energy, μ and the width, σ are the same for both functions, and are allowed to float in the fit. Also floating are the overall amplitude A and the relative scales of each function, g , f , and h . A factor of $\sqrt{2}$ in the argument of the erfc is required to account for the definition of the gaussian width. In the case of ^{60}Co , both full-absorption peaks are fit simultaneously. The ratio of the peak energies, μ_1 and μ_2 , is fixed to be the true ratio of the line energies,

$$\mu_1/\mu_2 = 1333 \text{ keV}/1173 \text{ keV}. \quad (6.9)$$

Similarly the widths, σ_1 and σ_2 , are fixed to follow the expected variation of resolution with energy as $\sqrt{\mu}$,

$$\sigma_1/\sigma_2 = \sqrt{\mu_1/\mu_2}. \quad (6.10)$$

²This is the functional form used in GF2, a least-squares peak-fitting program used for fitting γ peaks in Germanium detectors. See also <http://www.phy.anl.gov/gammasphere/doc/gf2.hlp>.

$$F(E) = \begin{cases} A\{ge^{-(E-\mu)^2}/2\sigma^2 + (1-g)\text{erfc}[(E-\mu)/\sqrt{2}\sigma]\} & {}^{228}\text{Th}/{}^{137}\text{Cs} \\ Af\{ge^{-(E-\mu_1)^2}/2\sigma_1^2 + (1-g)\text{erfc}[(E-\mu_1)/\sqrt{2}\sigma_1] \\ +(1-f)(he^{-(E-\mu_2)^2}/2\sigma_2^2 + (1-h)\text{erfc}[(E-\mu_2)/\sqrt{2}\sigma_2])\} & {}^{60}\text{Co} \end{cases} \quad (6.11)$$

Fits from a typical run are shown in Figures 6.3 and 6.4. The fitted peak positions of each bin are then graphed; Figure 6.5 shows an example of the charge collected as a function of drift time, both for all data, and for each TPC half separately. The graph is fit to an exponential in each half, with the fitted constant being taken as the electron lifetime in that half.

6.3.2 Resolution method

When left uncorrected, the effect of the electron lifetime tends to broaden the full-absorption peak, as the amount of charge collected varies for events of the same energy. This broadening should be eliminated when the correct electron lifetime is used, since all the corrected charge signals should be restored to the full-absorption peak energy (subject to detector resolution).

The resolution method uses this observation by scanning over many different electron lifetime values and identifying the one that minimizes the resolution of the full-absorption peak. This is then selected as the measured electron lifetime. At

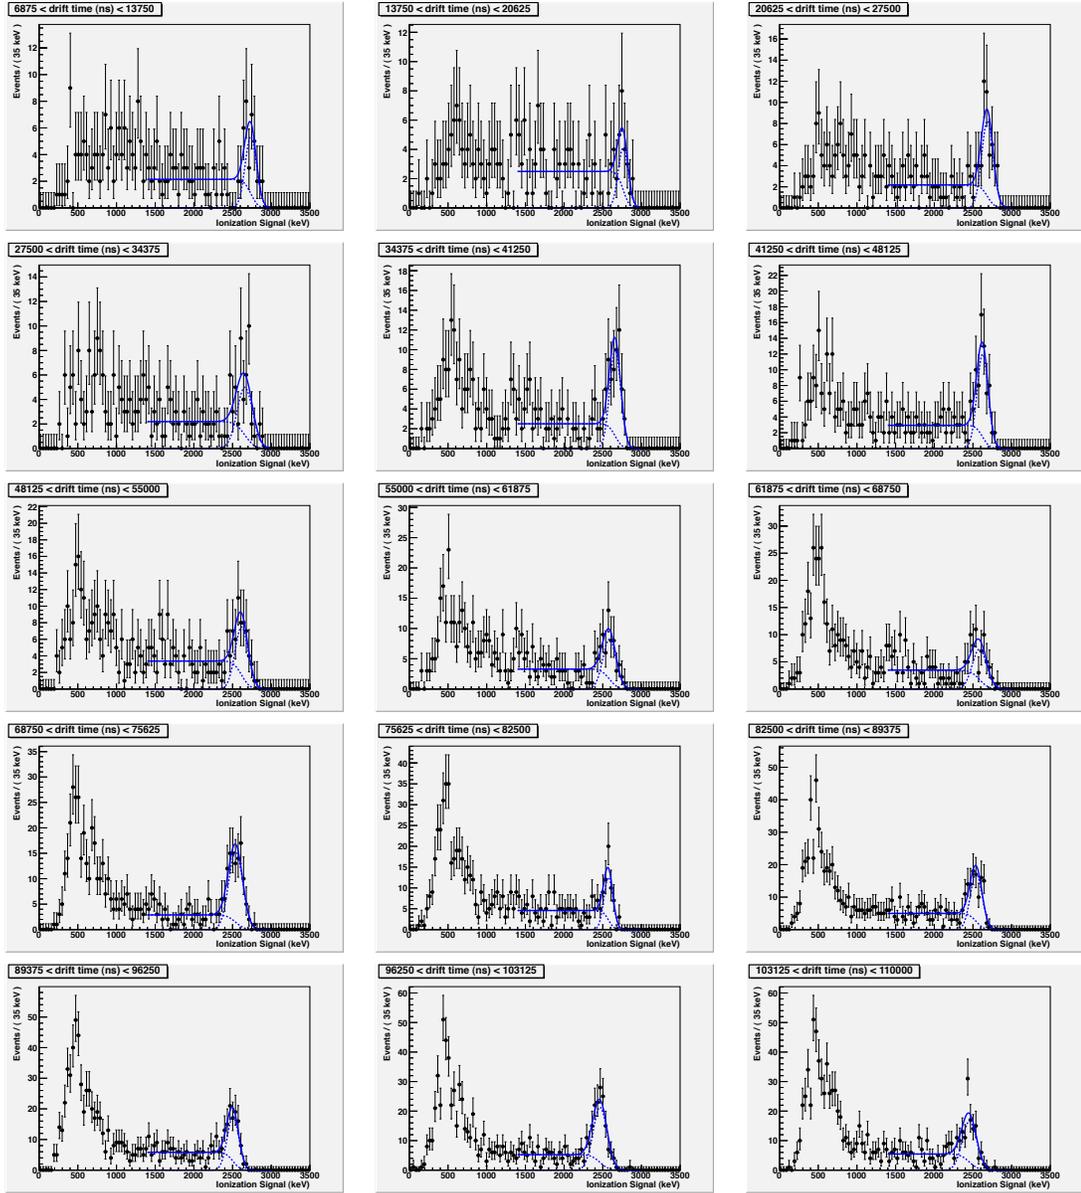


Figure 6.3: Example of fits used in the peak method to find the full-absorption peak position. Each plot corresponds to TPC 1 events with a different range of drift times. Data are taken from run 3124, a calibration run with ^{228}Th at the cathode (S5). The dashed curves show the gaussian and erfc components of the fitted function, while the solid curve indicates the total. The full-absorption peak moves to lower energy as the drift time increases, as seen clearly in Figure 6.5.

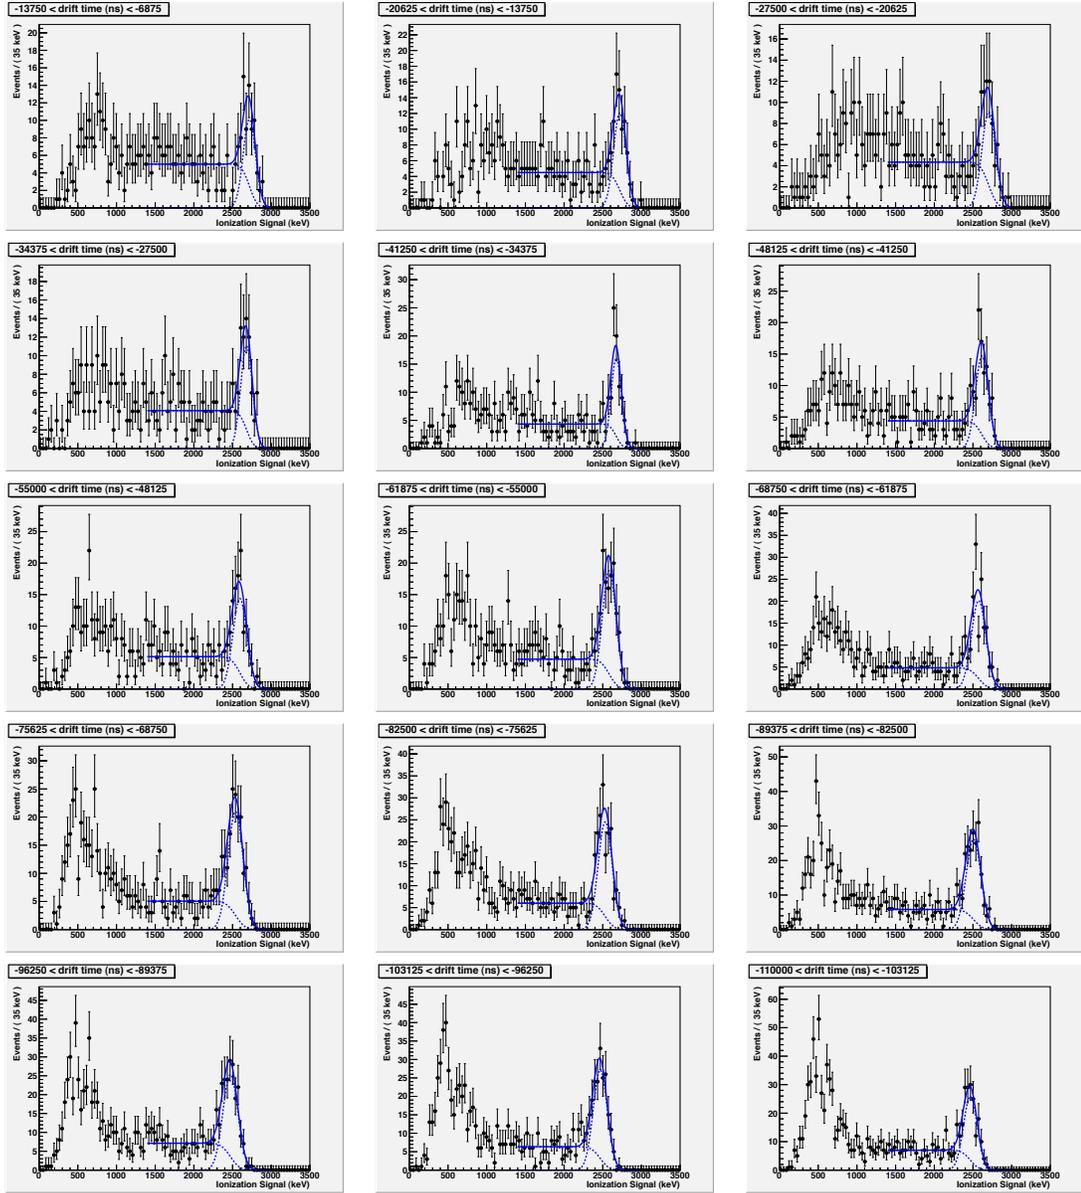


Figure 6.4: Example of fits used in the peak method to find the full-absorption peak position. Each plot corresponds to TPC 2 events with a different range of drift times. Data are taken from run 3124, a calibration run with ^{228}Th at the cathode (S5). The dashed curves show the gaussian and erfc components of the fitted function, while the solid curve indicates the total. The full-absorption peak moves to lower energy as the drift time increases, as seen clearly in Figure 6.5.

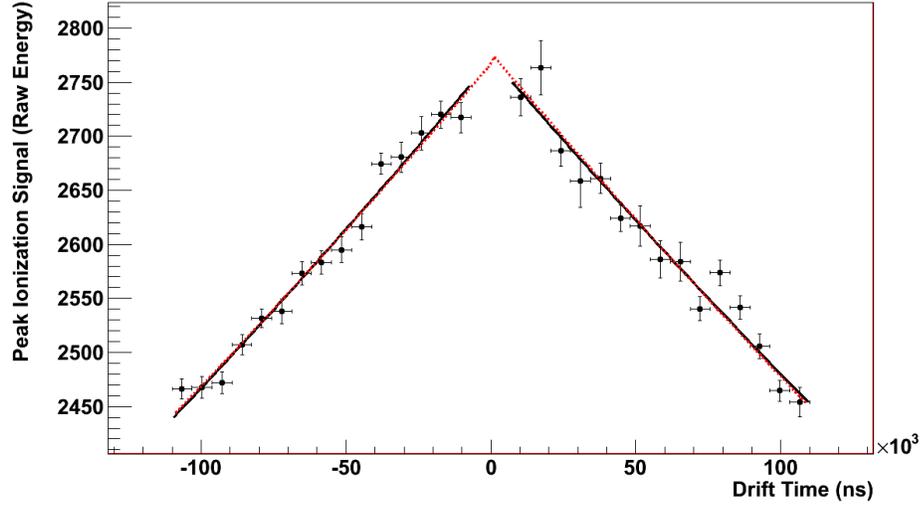


Figure 6.5: Average full-absorption peak charge energy, grouped by drift time, for a ^{228}Th calibration source (2615 keV γ). Least-squares fitting to an exponential results in the curves shown for TPC 1, TPC2, and both TPCs combined. The best fit lifetimes are $896.4 + 62.7 - 26.9 \mu\text{s}$ and $862.5 + 38.8 - 21.6 \mu\text{s}$ for TPC 1 and TPC 2 separately, and $876.3 + 39.4 - 17.5 \mu\text{s}$ combining both TPCs. Deposits in TPC 2 are shown with negative drift time solely for the purposes of display. Data are taken from run 3124, with the source at the top of the cathode.

each step in the scan, a different electron lifetime correction is applied to the charge energy spectrum, and the full-absorption peak is fit using a fit function similar to that used in the peak method:

$$G(E) = \begin{cases} A\{e^{-(E-\mu)^2/2\sigma^2} + g \times \text{erfc}[(E - \mu)/\sqrt{2}\sigma]\} & ^{228}\text{Th} \\ A\{e^{-(E-\mu)^2/2\sigma^2} + g \times \text{erfc}[(E - \mu)/\sqrt{2}\sigma] \\ + h \times (e^{-(E-\mu r)^2/2\sigma^2} + g \times \text{erfc}[(E - \mu r)/\sqrt{2}\sigma])\} & ^{60}\text{Co} \end{cases} \quad (6.12)$$

In this case, the ratio of ^{60}Co peak energies is still fixed at $r = (1333 \text{ keV})/(1173 \text{ keV})$, and the ratio of Compton shelf and full-absorption peak amplitudes, g , is fixed

to be the same for both peaks. For fit stability, σ is fixed to be the same for both peaks.

The fitted value of σ is plotted versus inverse lifetime, and the resulting graph is fitted to a parabola; an example is shown in Figure 6.6. The lifetime corresponding to the minimum of the parabola (the minimum resolution) is taken as the measurement, with uncertainty given by the statistical uncertainty from the parabolic fit. The uncertainty is asymmetric since longer lifetimes produce smaller corrections; a given fractional increase in lifetime produces a smaller change in resolution than the corresponding fractional decrease.

The range of lifetime values used extends into negative lifetimes, corresponding to a correction factor $\kappa < 1$ (see Equation 6.6). This improves determination of a minimum when the lifetime is very large (correction factor near 1). Figure 6.7 illustrates the improvement.

6.3.3 Electron lifetime during the physics run

The electron lifetime during the course of the physics run, as measured by the peak method, is plotted in Figure 6.8. The purity is quite high at the beginning of the run due to months of continuous recirculation during a previous run. In general, the electron lifetime remained between 2 and 4 ms during the campaign. This corresponds to charge attenuation lengths of 4-7 m, taking the attenuation length l_{att} to be the lifetime multiplied by the measured drift velocity of $v_d = 1.71$

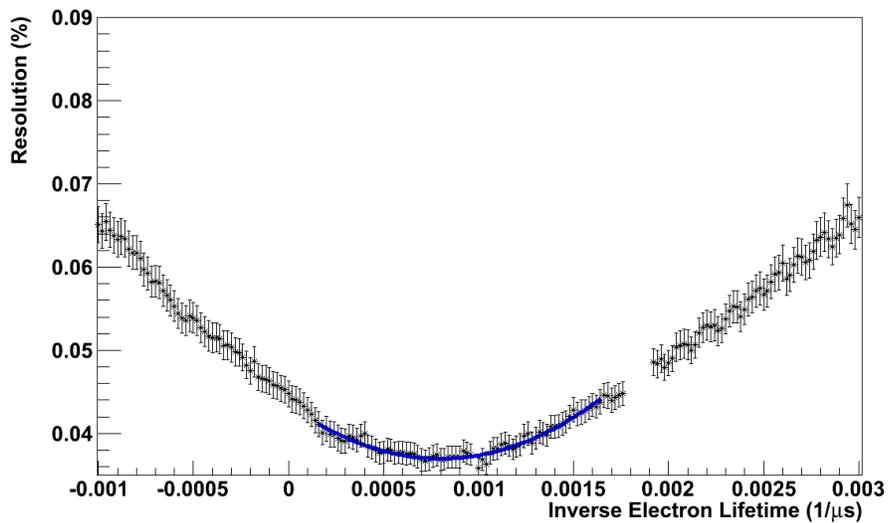
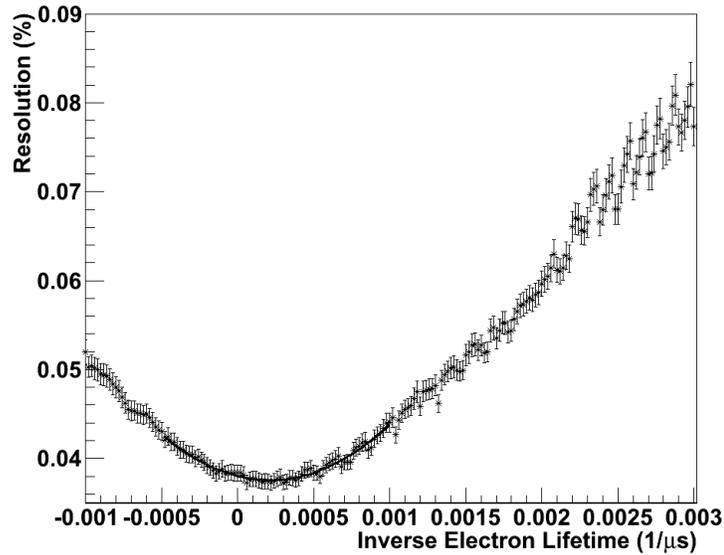
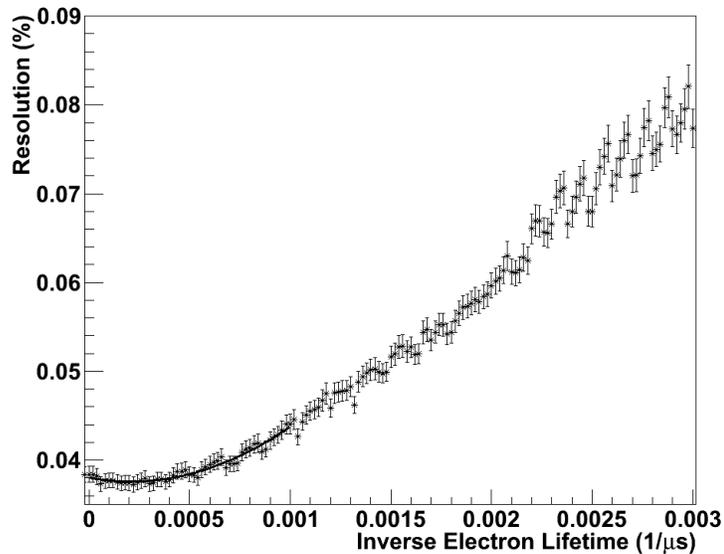


Figure 6.6: Electron lifetime measurement using the resolution method. The resolution of the full-absorption peak is measured as a function of the lifetime correction and fit to a parabola to extract the minimum. The fit shown uses data from calibration run 3124, taken 01/20/2012 with the weak ^{228}Th source. The electron lifetime was measured at $1249 + 34.14 - 32.37 \mu\text{s}$. The errors on the lifetime are taken from the fit, using the error bars from the fitted resolutions of the individual spectra. Thus, these errors are highly correlated.



(a) Resolution method applied to a run with a large electron lifetime. The data is from calibration run 3142, taken 01/23/2012 with the weak ^{228}Th source. The electron lifetime was measured at $4677 \pm 293 - 260 \mu\text{s}$.



(b) The resolution method applied to the same run without using negative lifetimes. The minimum is much less clear, and the electron lifetime is substantially more uncertain, with a measurement of $4780 \pm 1573 - 949 \mu\text{s}$.

Figure 6.7: Electron lifetime measured for a run with high purity, illustrating the effects of using negative lifetimes to improve the measurement

mm/ μ s:

$$l_{att} = \tau \cdot v_d \quad (6.13)$$

The maximum correction factor for these lifetimes is 3-6% of the charge collected, applied to events with the longest drift time. The brief drops to small lifetime near 1/18/12 and 4/1/12 are due to stoppages in recirculation, when the xenon is not being constantly repurified. Once recirculation is restarted, the purity recovers quickly, seen as a sharp rise in lifetime.

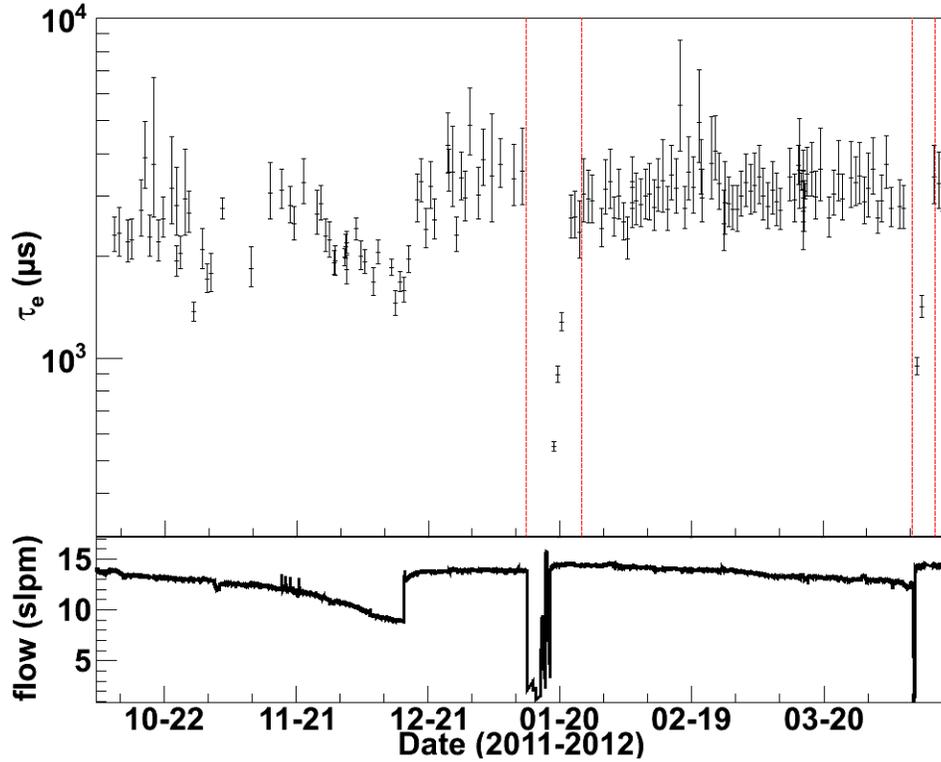


Figure 6.8: Electron lifetime measured using peak method, with both TPCs combined. Data in between the vertical bands are not included in the physics data set due to low purity levels. The lower plot shows the xenon recirculation rate during the physics run; there is a strong correlation between purity and flow rate. From [78].

The lifetime is measured during calibration runs. During the much longer physics runs, it slowly varies. To allow for interpolation between calibration runs, the

variation of lifetime with time is fit piecewise to polynomials of no higher than fourth degree. The fitted functions are shown in Figure 6.9, and the explicit polynomials are given in Table 6.1. Fitting to a range of time also avoids variance in the lifetime correction from run to run. Each TPC is corrected separately, to account for possible differences in electric field or other z-dependent effects. However, the error bands in Figure 6.9 show that the purities in each TPC are generally consistent.

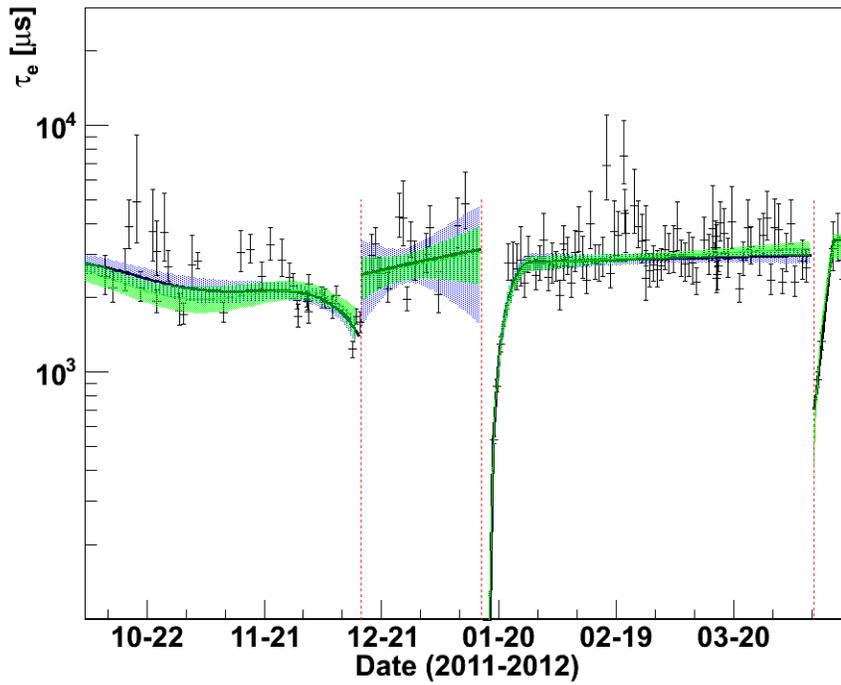
Table 6.1: Coefficients of polynomial fit to electron lifetime as a function of time. The fit is piecewise to incorporate several abrupt changes. The functional form is $\sum p_k(t - t_0)^k$, where t is in days, and t_0 is a reference time depending on the piecewise fit.

| Begin date | End date | p_0 | p_1 | p_2 |
|------------|------------|-----------|----------|------------|
| 2011-09-16 | 2011-12-16 | 2475.838 | 61.66379 | -4.091837 |
| 2011-12-16 | 2012-01-16 | 2433.325 | 20.49931 | 0.0 |
| 2012-01-16 | 2012-01-28 | -1006.269 | 594.3029 | -23.51381 |
| 2012-01-28 | 2012-04-10 | 2745.730 | 5.922156 | 0.0 |
| 2012-04-10 | 2012-04-15 | 366.3246 | 543.6926 | -0.5489047 |

| Begin date | End date | p_3 | p_4 |
|------------|------------|------------|---------------------------|
| 2011-09-16 | 2011-12-16 | 0.07377156 | $-4.128929 \cdot 10^{-4}$ |
| 2011-12-16 | 2012-01-16 | 0.0 | 0.0 |
| 2012-01-16 | 2012-01-28 | 0.0 | 0.0 |
| 2012-01-28 | 2012-04-10 | 0.0 | 0.0 |
| 2012-04-10 | 2012-04-15 | 0.0 | 0.0 |

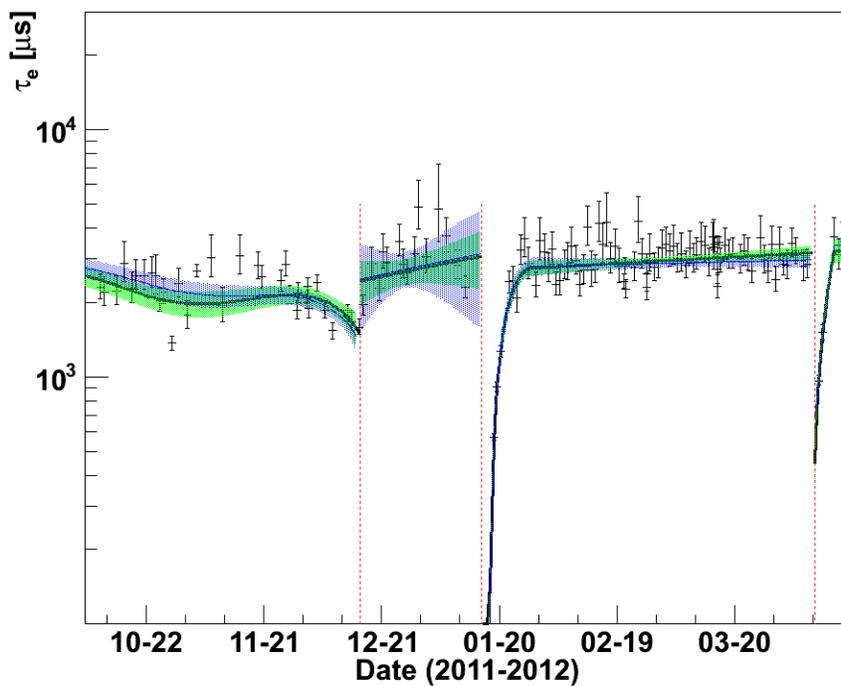
The electron lifetime measured by the resolution method can be seen in Figure 6.10. The values largely fall in the same range as the peak method, although they are systematically slightly higher using the resolution method, possibly due to a loss of sensitivity of one method at large lifetimes. One can also see that the resolution method produces unusually high lifetimes when the source is at an anode and the lifetime is largest. This indicates a loss of sensitivity of the method at the largest lifetimes, which becomes most pronounced when the source is at an anode, due to

TPC 1



(a) TPC 1

TPC 2



(b) TPC 2

Figure 6.9: Electron lifetime as measured by each TPC using the peak method. The curves and error bands are found from piecewise polynomial fits, blue to TPC 1 and green to TPC 2. Dashed vertical lines separate different segments of the polynomial. From [78].

the small drift time.

The two methods are compared in Figure 6.11. For the final analysis, the peak method was selected as having more consistent values between different calibration sources and source positions, as well as better understood uncertainties.

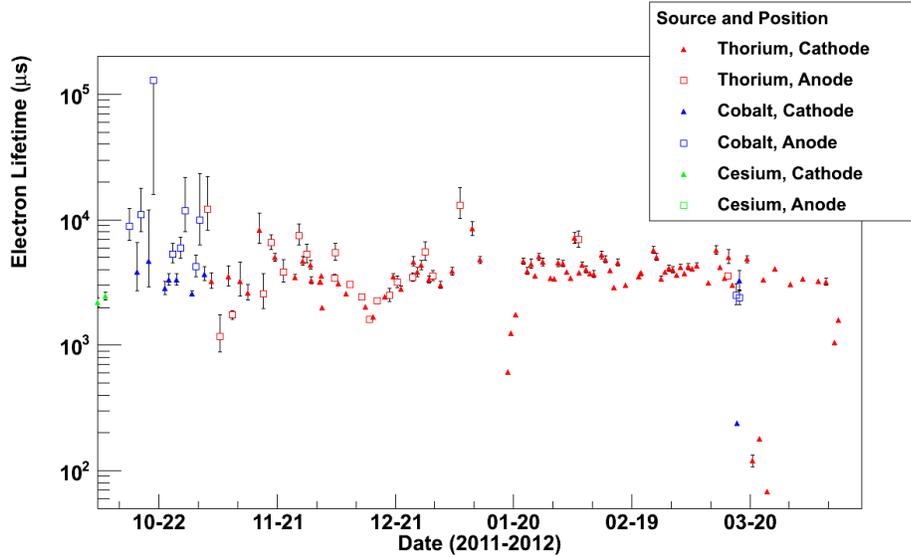


Figure 6.10: Electron lifetime measured by the resolution method. The source type and position are indicated. Note that the size of the error bars increases substantially at large lifetimes, and that some separation is seen between lifetimes measured at anode and cathode, especially for ⁶⁰Co.

6.4 Shielding grid inefficiency

Charge liberated from the liquid xenon results in positively charged xenon ions in the bulk of the TPC. Ultimately, these ions will drift to the cathode and be collected, but due to their mass, the drift velocity is low. At 170 K, the ion mobility is $3.7 \times 10^{-3} \text{ cm}^2/\text{Vs}$ [79], corresponding to a drift velocity of $\sim 0.01 \text{ mm/ms}$. This motion can be neglected during a 2.048 ms frame. These ions induce a constant charge on the collection wires (U-wires) with the opposite sign of the signal charge.

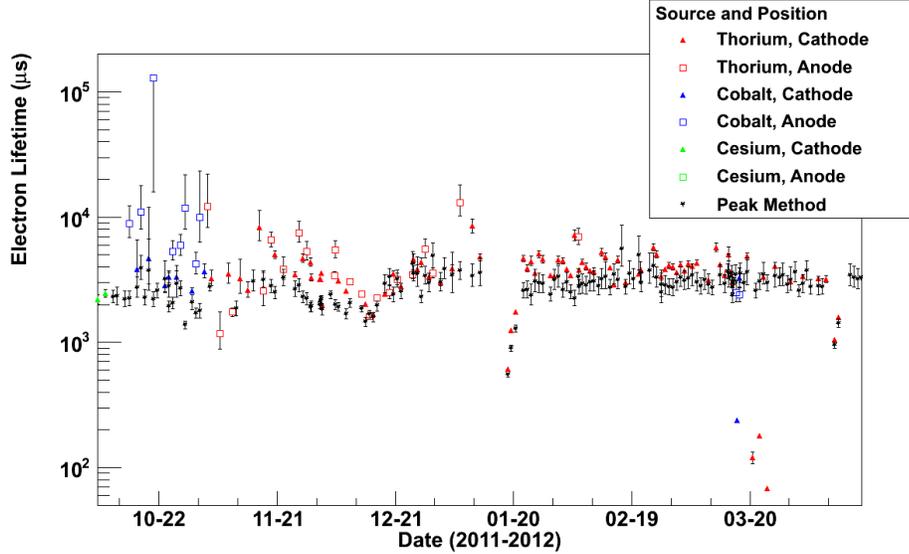


Figure 6.11: Comparison of the two electron lifetime methods. The source type and position are indicated for the resolution method. The peak method is seen to be more stable, with fluctuations consistent with the uncertainties.

The overall collection signal amplitude is reduced as a result. The shielding wires (V-wires) largely compensate this by screening the U-wires, but the screening is inefficient since the wires do not completely cover the plane. The ratio of charge induced, Q_{ind} , to charge deposited, Q_{dep} , depends only on the axial position and a geometrical factor, σ [76]:

$$Q_{ind}/Q_{dep} = \sigma \frac{|z|}{z_0}, \quad (6.14)$$

where $z_0 = 19.22$ cm is the distance from the cathode to the V-wires and z is the distance of the charge deposit from the cathode. For circular wires, the geometrical factor σ depends on the wire radius, r , pitch p , and the distance, l between the shielding wires and the anode. The EXO-200 TPC wires have a square profile, but since the wire pitch is much larger than the wire size, we can use the functional form for circular wires to good approximation. We take $r = 0.00635$ cm, $p = 0.3$ cm, and

$l = 0.62$ cm, yielding

$$\sigma = \frac{p}{2\pi l} \log \frac{p}{2\pi r} = 0.16 \quad (6.15)$$

for the TPC. The induced charge is most significant when charge is deposited near the V-wires.

This inefficiency is observed in calibration data. A study was made using the strong ^{228}Th source runs 2424, 2426, 2431, 2432, 2433, 2434, 2447, 2448 taken prior to the low-background data campaign. Events were divided into 75 evenly-spaced bins, each containing events in a 2.5 mm range of $|z|$. The purity-corrected ionization energies, E_{pc} were made into 75 spectra; to increase statistics while retaining the small bin size, events with the same drift time from both TPCs were mixed in each bin. (No significant variation was seen between the two TPCs.) The ^{208}Tl full-absorption peak at 2615 keV was fit for each spectrum using the fit function in Equation 6.12. Figure 6.12 shows clearly that the peak position clearly decreases at large $|z|$, near the anodes, as expected.

The effect was reproduced in MC ^{228}Th data by a similar procedure, and the MC data was then used to generate a correction. Events with charge collected on a single U-wire were studied separately from events where two U-wires collected charge, since the effect of the inefficiency should be increased for two U-wire signals.³

³To see this, consider that an amount of charge, Q_i , is induced by positive ions on two neighboring wires separately. If a signal deposits charge Q_{sig} all on one wire, the observed charge Q_{obs} is

$$Q_{obs} = Q_{sig} - Q_i \quad (\text{single wire deposit}). \quad (6.16)$$

However, if Q_{sig} is spread across two wires, depositing charges Q_{sig}^1 and Q_{sig}^2 , then the observed signals are

$$Q_{obs}^{(1)} = Q_{sig}^{(1)} - Q_i \quad \text{and} \quad (6.17)$$

$$Q_{obs}^{(2)} = Q_{sig}^{(2)} - Q_i. \quad (6.18)$$

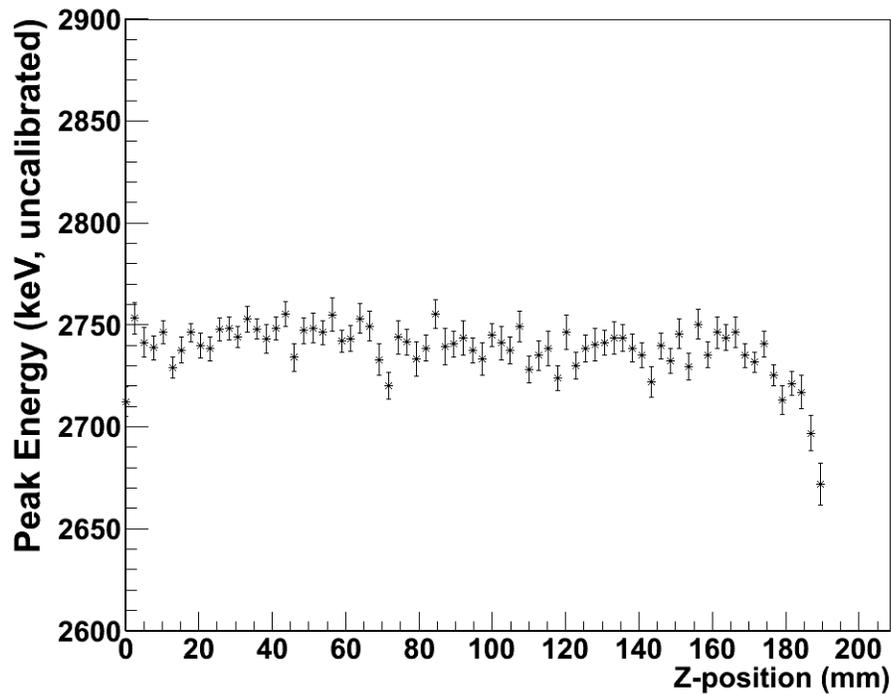


Figure 6.12: ^{208}Tl full-energy peak position in various z-bins, measured using dedicated strong calibration source data. Purity corrections have already been applied. Induction effects lead to a residual dependence of peak position with z-position, most noticeably at large values of z (far from the cathode). Signals striking any number of U-wires are included.

The resulting graphs of MC peak position, shown in Figure 6.13a, were then fit to a correction function, $f_{eff}(|z|)$, to extract correction parameters:

$$f_{eff}(|z|) = E_0 \left(1 - \frac{B}{z_0 - |z|} \right), \quad (6.20)$$

where we now work in terms of reconstructed energy, rather than strictly charge. This functional form follows from Equation 6.14 by solving for the measured charge signal $Q - Q_{ind}$. E_0 , B and z_0 are allowed to float; z_0 represents the anode-cathode distance, B represents the magnitude of the inefficiency, and E_0 is the observed ionization energy from a charge deposited infinitely far from the anode. The fit value

$$z_0 = 192.1 \pm 0.378 \text{ mm} \quad (6.21)$$

is consistent with the known anode position at $z = 192.2$ mm. Also, the fitted value of the inefficiency factor, B , is approximately twice as large for two U-wire signals as for single U-wire signals, consistent with the claim that the inefficiency will reduce the sum energy as many times as there are collection wires. To check that the MC represents the induction signals adequately, the data were also fit, fixing z_0 and B . The far-away energy E_0 is not calibrated in the data, so it was allowed to float, although its value is not significant. The quality of the fits was acceptable, as shown in Figure 6.13b and Table 6.2. The MC is able to model this induction effect

Summing to get the total charge observed, the combined two U-wire signal is then

$$Q_{obs} = Q_{obs}^{(1)} + Q_{obs}^{(2)} = Q_{sig} - 2Q_i \quad (\text{two wire deposit}) \quad (6.19)$$

well because the of the detailed simulation of the wire signals in the digitizer. In particular, the weight field used to calculate the induced signals is verified.

Table 6.2: Best fit parameters for the shielding grid inefficiency curve shown in 6.20. Errors are statistical only. Only the E_0 term floats in the fits to data, indicated by the vanishing errors on the other terms. Fit quality is indicated by the χ^2 per degree of freedom.

| Data type | U-wires | E_0 | B | z_0 (mm) | χ^2 /(d.o.f.) |
|-----------|---------|------------------|---------------------|------------------|--------------------|
| MC | Any | 2583 ± 0.886 | 0.0937 ± 0.0106 | 192.1 ± 0.38 | 105.3/72 |
| MC | 1 | 2604 ± 1.347 | 0.072 ± 0.017 | 192.4 ± 0.88 | 167.7/72 |
| MC | 2 | 2559 ± 1.24 | 0.135 ± 0.016 | 192.5 ± 0.46 | 241.5/72 |
| Data | Any | 2744 ± 0.75 | 0.0937 ± 0 | 192.1 ± 0 | 118.7/74 |
| Data | 1 | 2583 ± 0.886 | 0.072 ± 0 | 192.4 ± 0 | 79.33/74 |
| Data | 2 | 2583 ± 0.886 | 0.135 ± 0 | 192.5 ± 0 | 143.7/74 |

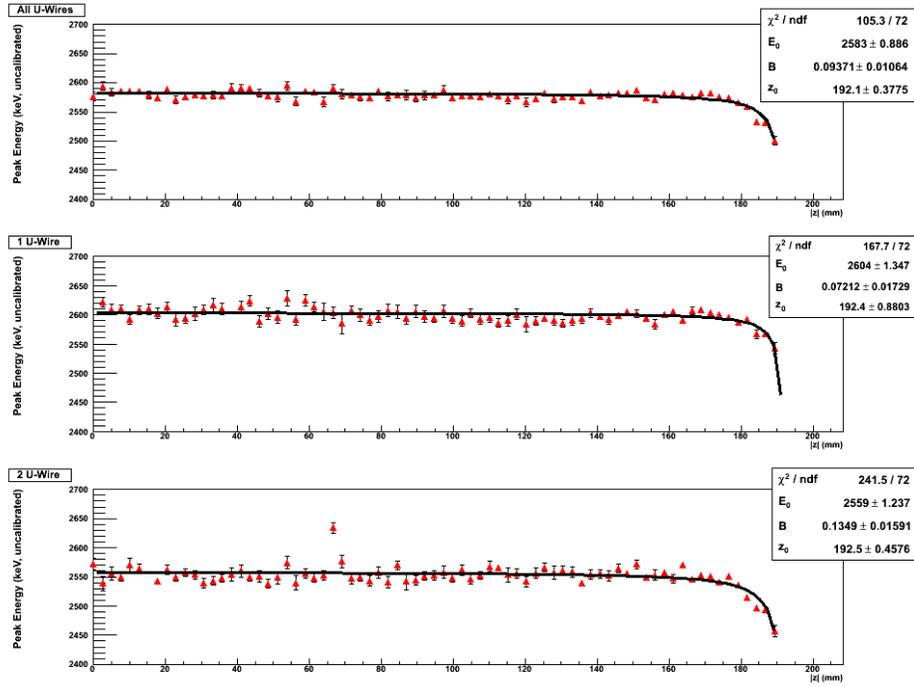
Near $|z| = 180$ mm, the loss of signal is $\sim 1\%$, significant enough to warrant applying a correction. The correction, λ , is obtained by from inverting Equation 6.20

$$\lambda = \frac{1}{1 - \frac{B}{z_0 - z}}. \quad (6.22)$$

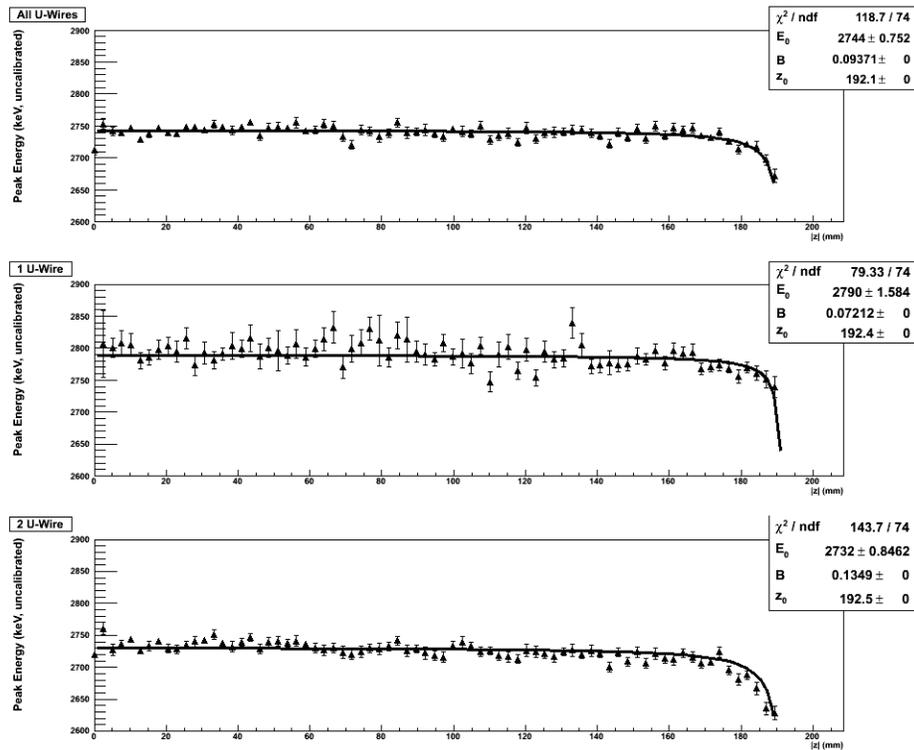
This is applied to the purity-corrected charge cluster energy, E_{pc} , to find the grid-corrected cluster energy, E_{gc} :

$$E_{gc} = \lambda \cdot E_{pc} \quad (6.23)$$

The appropriate parameters are used for events with single and two U-wire collection signals. Single-site events with more than two U-wires collecting charge are rare, and the two U-wire correction is applied in those cases without significant loss of energy resolution. Figure 6.14 shows that after the correction, charge deposits near the anodes no longer suffer loss of amplitude.



(a) MC



(b) Calibration data

Figure 6.13: The reduction in collected energy with z is reproduced in MC ^{228}Th data. Events with charge collection on 1 or 2 U-wires are studied separately and fit to extract correction parameters. As a check, the correction functions are fit to the data fixing those parameters, showing a good fit.

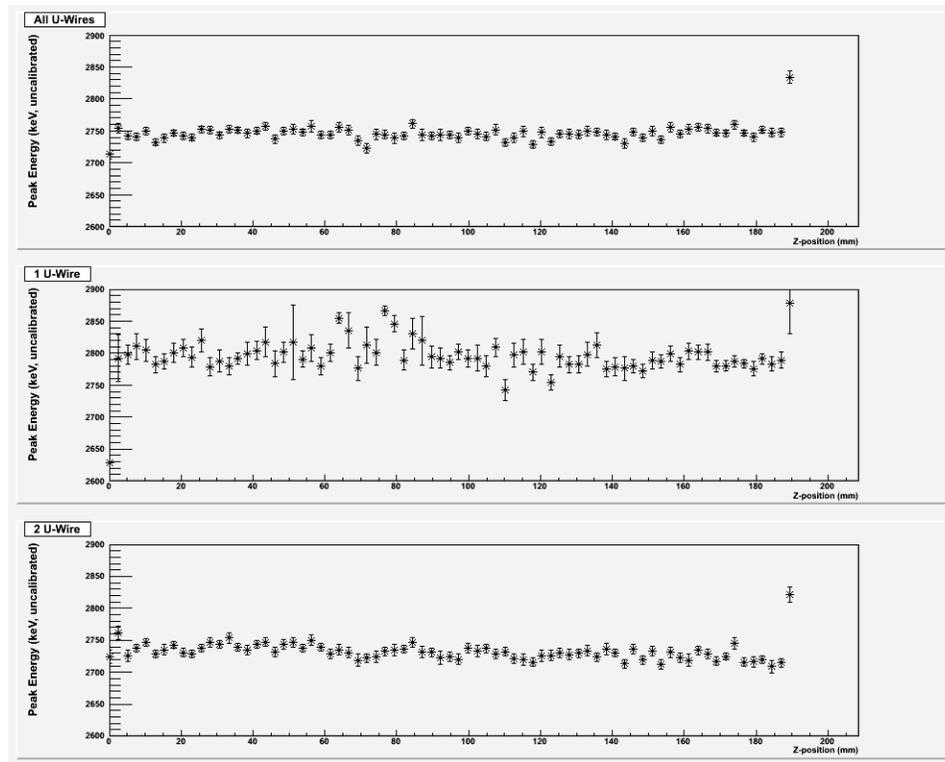


Figure 6.14: After the shielding grid correction has been applied, the full-absorption peak energy, E_{gc} , is a constant function of $|z|$.

6.5 Charge energy scale

After the corrections have been applied, the calibration sources can be used to calibrate the charge-only energy scale. The charge scale is not used in extracting $\beta\beta 0\nu$ rates, but it allows us to discern a possible difference between the energy scales of β and γ decays.

In addition to the four standard calibration source peaks, another peak in the ^{228}Th spectrum may be used. This peak occurs at 511 keV, and consists of “escaping” photons from the single and double-escape peaks that are identified in a separate charge cluster. Spectra of charge cluster energies are formed using a fiducial cut described in Section 7.2.3 below. The peak energies are fit, using the gaussian + erfc fit function given in Equation 6.11. A linear fit to the five calibration peak energies, shown in Figure 6.15, yields a parameterization for the energy scale for single-site and multi-site events:

$$E_Q^{cal} = \begin{cases} 72.22 + 0.9238 \cdot E_{gc} & \text{single-site} \\ 113.35 + 0.9287 \cdot E_{gc} & \text{multi-site.} \end{cases} \quad (6.24)$$

Here, E_Q^{cal} is the calibrated charge energy, and E_{gc} is the reconstructed charge energy, subject to the corrections above.

6.5.1 Potential β - γ energy scale discrepancy

The double-escape peak in the ^{228}Th spectrum, mentioned in Section 6.2, presents another possible calibration line. However, as shown in red in the single-

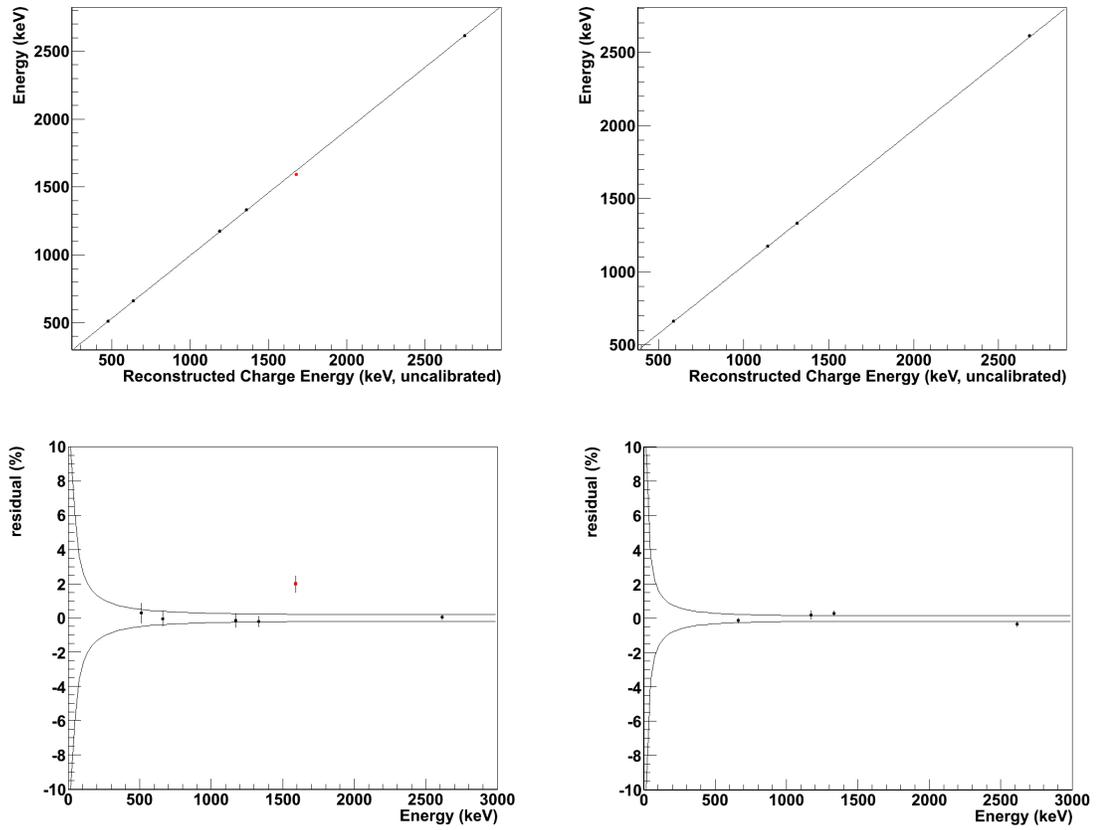


Figure 6.15: Linear calibration curves for the charge energy, created using the four calibration source gammas and the 511 keV peak. Single-site is on the left and multi-site on the right. Residuals are shown below with error bands. The best fit parameters are given in Equation 6.24. Red points in the single-site spectrum indicate the ^{228}Th double-escape peak, which does not align with other sources and is not used in the calibration or analysis.

site spectrum of Figure 6.15, this peak does not align with the γ source calibration, by $1.9\% \pm 0.45\%$. Since the double-escape peak is a purely β -like energy deposit, this leads to concern that the energy scale for β and for γ deposits may not be the same, implying that using γ calibration sources might lead to a systematic error in the search for $\beta\beta 0\nu$. No other purely β -like calibration source could be found to verify the discrepancy or create a separate, β -like calibration, so the double-escape peak was not used for source calibration. Further evidence for an anomaly in the double-escape peak can be seen in the calibration of the MC charge energy scale. Figure 6.16 shows that the residual for this peak is 0.6% offset from the best fit calibration. A possible explanation is that the energy scale is affected by the typical spatial extent of a charge cluster, which may differ for β and γ deposits. Attempts to quantify any differences in the β and γ energy scales are discussed in 7, particularly in Sections 7.7.1 and 7.9.2. A possible explanation of the discrepancy, along with a partial solution, is described in Section 6.8.4.

6.6 Scintillation spatial dependence

For a deposit of a given energy, the photon collection efficiency of the TPC is dependent on the location of the deposit. This arises from variations in APD gains, as well as differences in the solid angle subtended by the APD planes. A minor correction comes from shadowing effects from detector edges or internal components. A uniform spatial response is required, so a three-dimensional position-dependent correction function, or *light map*, is determined. Radioactive source calibration data

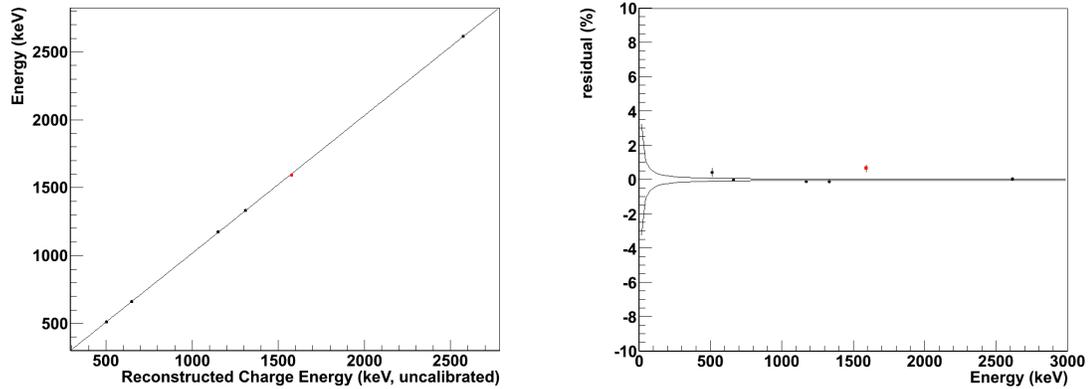


Figure 6.16: Linear calibration curves for the MC charge energy, created using the four calibration source gammas and the 511 keV peak. The best fit single-site energy scale is on the left, and the residuals are on the right. The red points indicate the ^{228}Th double-escape peak, which does not align with other sources, and is not used in the calibration.

was used to measure the LAAPD response to events in various parts of the detector from a source of fixed energy. This response was then used to generate the *light map*, which can be applied to low-background data.

Data from many ^{228}Th calibration runs (between 2424-3298 and 3342-3617) were used to generate the light map. These runs included all five source positions, to ensure as uniform a distribution as possible of events in the TPC. Position information is taken from reconstructed charge clusters associated with the scintillation clusters; only single-site events were used, since the position of energy deposition for multi-site events is ambiguous. The detector was divided non-uniformly into 1352 voxels, or regions, formed by 13 radial divisions, 13 division along the z axis, and 8 azimuthal divisions. The voxels were chosen to include sufficient statistics in every bin. The azimuthal divisions are evenly spaced. In z , 11 evenly spaced divisions are made. The central division is divided evenly three further times to improve the

resolution near the cathode, where the efficiency is rapidly changing. The radial bins vary in size to account for changing volume, and are listed in Table 6.3.

Table 6.3: Distribution of radial bins for the light map.

| Range of radial coordinate (mm) | Bin width (mm) | Number of bins |
|---------------------------------|----------------|----------------|
| $0 < r < 30$ | 30 | 1 |
| $30 < r < 90$ | 20 | 3 |
| $90 < r < 120$ | 10 | 3 |
| $120 < r < 168$ | 8 | 6 |

To ensure only true full-absorption peak events were used, the full TPC scintillation spectrum was made, and the 2615 keV full-absorption peak was fit to the usual gaussian + erfc shown in Equation 6.11. Only events with energy between 0.5 and 4.0σ were accepted for further analysis, where σ was the fitted width of the gaussian. The scintillation spectrum was then made for each voxel, and the full-absorption peak was refit. The mean peak energy, normalized to 2615 keV, was taken as the light map value for that voxel. A continuous correction function, $f_{light}(r, \phi, z)$, was determined by making a trilinear interpolation of the points, with ϕ being the azimuthal angle. This correction can then be applied to the sum of the counts on both APD planes by division:

$$E_S = E_S^{recon} / f_{light}(r, \phi, z), \quad (6.25)$$

with E_S^{recon} the raw reconstructed scintillation energy, and E_S the corrected energy. For multiple cluster events, an energy weighted sum of the correction for each cluster is used.

The light map, averaged over all azimuthal bins, is shown in Figure 6.17. The APD plane has the largest solid angle of acceptance in these regions, so it is expected that the efficiency is high. Except for one small region, the efficiency does not vary by more than 10% from the average. The exception occurs at large r and small $|z|$; here, a support ring for the cathode obstructs light collection in the secondary half-TPC. Both the solid angle and the cathode ring effects were anticipated by the MC studies of light collection described in Section 5.2.1, although the MC could not accurately predict their magnitudes. Finally, the light map appears slightly larger at positive z than negative z . This can be attributed to gain offsets between the two APD platters, with TPC 1 having the greater response.

Figures 6.18 and 6.19 show the two-dimensional projections of the light map onto the 8 azimuthal divisions and 13 divisions in z , respectively.

6.7 Summary of corrections

The corrections described are applied to the reconstructed signals in the same order as in which they are presented. The U-wire response correction, ξ , is applied to individual wire signals, which are re-summed to find the corrected charge cluster energy, E_{uc} . The electron lifetime and shielding grid corrections are applied to that energy, yielding a final corrected energy E_{gc} . Finally, the light map correction is applied to the scintillation cluster energy, yielding the corrected energy E_S . At this stage, we recall that each scintillation cluster can be associated with more than one charge cluster. We identify the *total ionization energy* associated with a scintillation

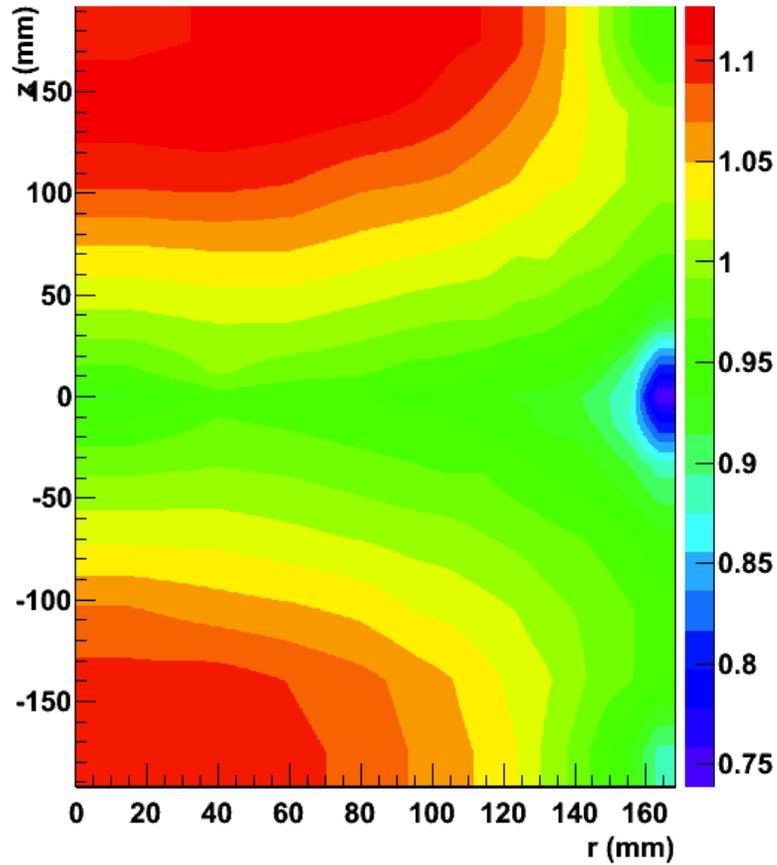


Figure 6.17: Light map correction function, averaged over all azimuthal bins to show gross features. Heat intensity shows the relative scintillation response to events in that voxel. Photon collection efficiency is largest at large $|z|$ and small r . A cathode support ring causes the largest loss in a small region at large r and small $|z|$. TPC 1 ($+z$) has a slightly larger response than TPC 2 ($-z$) due to gain variations. The light map is measured using the full-absorption peak in ^{228}Th calibration data. This figure and the following two figures are courtesy of S. Herrin.

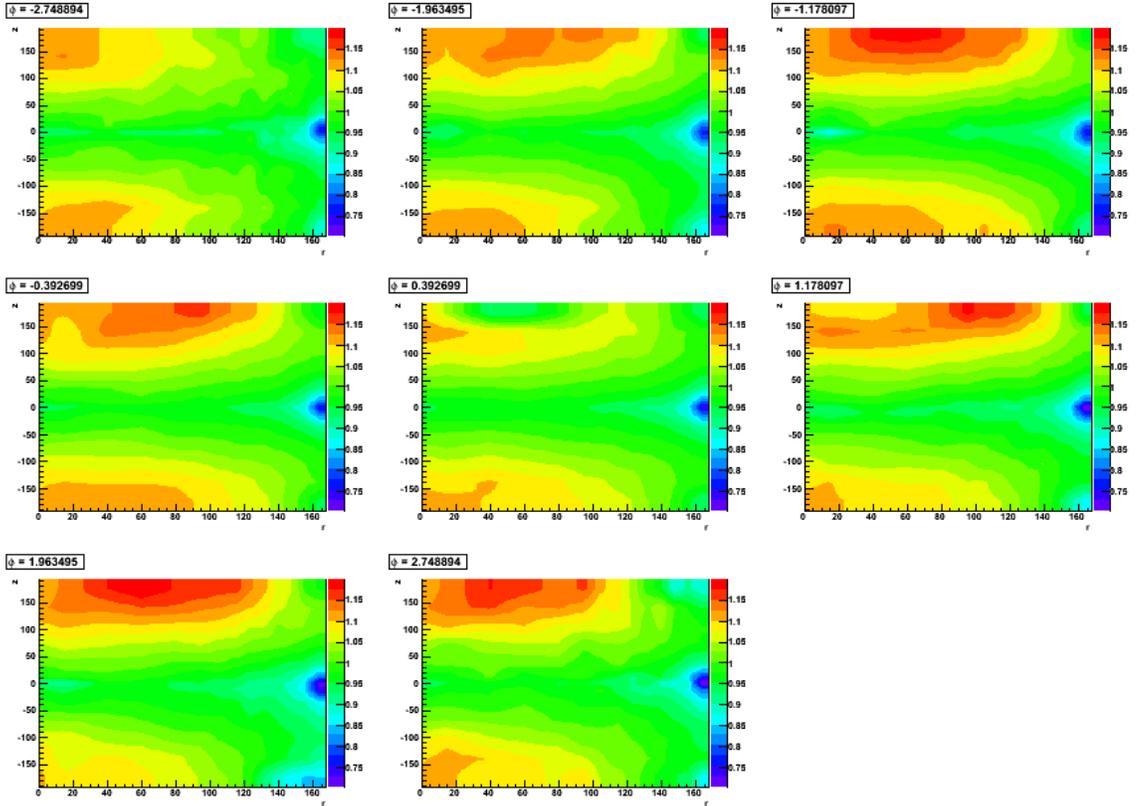


Figure 6.18: Azimuthal projections of the light map, showing the radial and z dependence in each division. Heat intensity shows the relative scintillation response to events in that voxel. The light map is measured using the full-absorption peak in ^{228}Th calibration data. There is some azimuthal dependence due to dead APD gangs on one plane; this is shown more clearly in Figure 6.19.

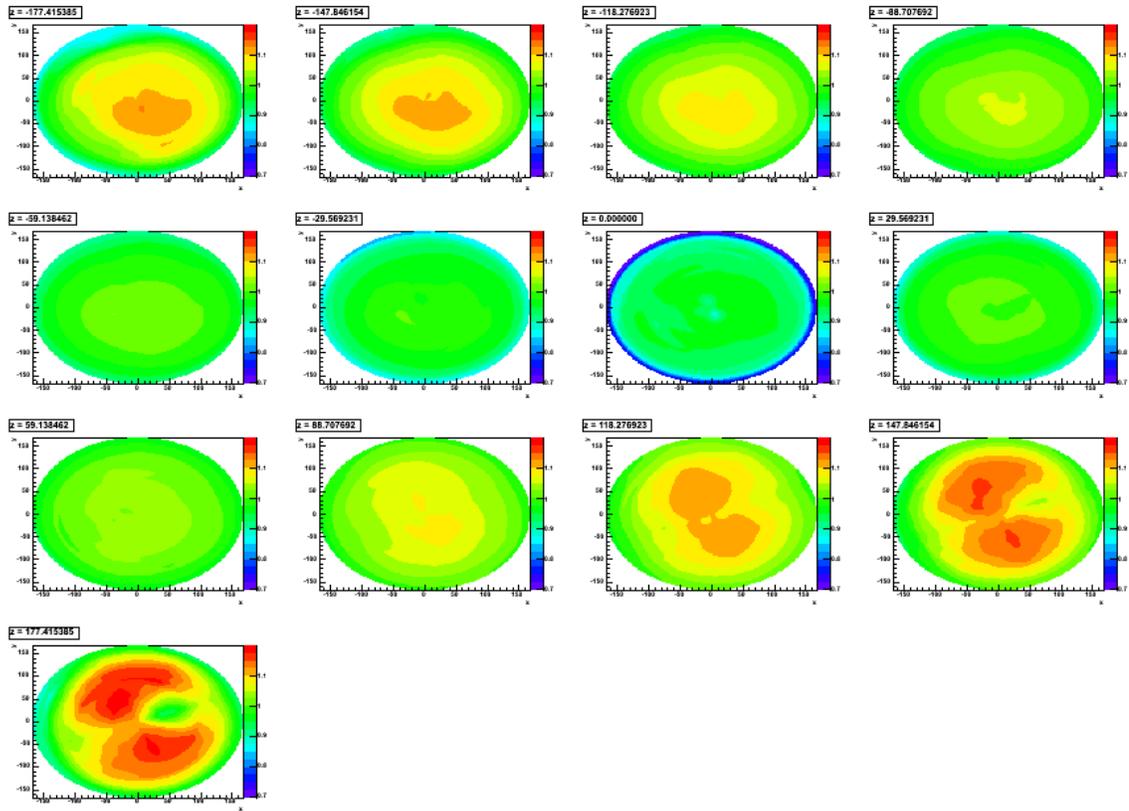


Figure 6.19: Axial (z) projections of the light map, showing the radial and azimuthal dependence in each division. Heat intensity shows the relative scintillation response to events in that voxel. A feature where the function becomes small is seen centered near $x = 50$ mm, $y = 25$ mm. This is a dead spot where APD gangs were disconnected due to noise. The light map is measured using the full-absorption peak in ^{228}Th calibration data.

cluster as the sum of charge cluster energies E_{gc} . This will be necessary for forming the combined energy spectrum.

$$E_Q = \sum E_{gc} \quad (\text{sum over all clusters}) \quad (6.26)$$

Table 6.4: Summary of corrections made to the reconstructed charge and scintillation energies. Corrections are applied initially to the wire signal energies, Q_{rec}^{wire} , and the scintillation clusters, E_S^{recon} .

| Correction | Correction Factor | Notation |
|-----------------------------|-------------------------|---|
| U-wire response | ξ | $E_{uc} = \sum \xi \cdot Q_{rec}^{wire}$ |
| Electron lifetime | κ | $E_{pc} = \kappa \cdot E_{uc}$ |
| Shielding grid inefficiency | λ | $E_{gc} = \lambda \cdot E_{pc} = \lambda \cdot \kappa \cdot E_{uc}$ |
| Total charge cluster energy | N/A | $E_Q = \sum E_{gc}$ |
| Light map | $f_{light}(r, \phi, z)$ | $E_S = E_S^{recon} / f_{light}(r, \phi, z)$ |

6.8 Combined charge and scintillation calibration

Once all the corrections are applied, we are ready to combine the charge and light energies, as discussed in Section 3.1. The calibration source data appear as shown in the 2-dimensional ionization-scintillation spectra in Figures 6.20 and 6.21. Here we plot E_S and E_Q , the scintillation and total ionization energies. Ellipses can be seen at the peak energies, exhibiting the anti-correction between charge and light. As discussed in Section 3.1, we can improve the energy resolution by choosing an energy variable along the short axis of the ellipse. This corresponds to a coordinate

rotation of this 2-dimensional spectrum:

$$\begin{cases} E_{rotated} = E_Q \cos \theta + E_S \sin \theta, \\ E_{long} = -E_Q \sin \theta + E_S \cos \theta. \end{cases} \quad (6.27)$$

θ is the rotation angle of the transformation, corresponding to the angle of inclination of the ellipse. $E_{rotated}$ represents the coordinate along the minor axis of the ellipse, the 1-dimensional *rotated* energy which will be used for the final analysis. The secondary coordinate, E_{long} , is not used. The value determined for the rotation angle is related to the resolutions in the separate charge and light channels, as well as the ratio of charge to light produced in a given type of event. It is thus specific to this detector, as well as the type of energy deposit (α , β , or γ). Once the optimal rotation angles are identified, the rotated energy can be calculated, and the energy scale can be calibrated for single-site and multi-site events. Note that it is not possible to use the 511 keV peak in this calibration, as we done with the charge calibration in Section 6.5, since E_Q no longer distinguishes the individual clusters of multi-site events.

6.8.1 Rotation parameter

The rotation angle is determined using the 2615 keV full-absorption peak from strong ^{228}Th source calibration data between runs 2424-2448. The optimal angles, θ_{SS} and θ_{MS} are determined by scanning over many different values of θ and selecting the ones that give the best energy resolution. For each value, a spectrum

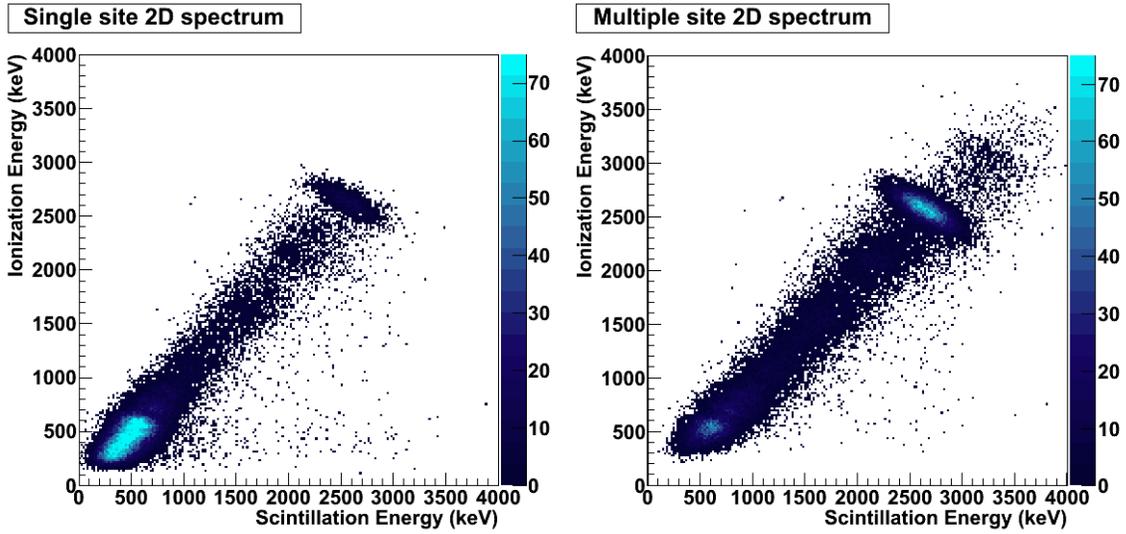


Figure 6.20: Two-dimensional charge-light spectrum of weak ^{228}Th calibration data from runs 2719, 737, 2754, 2761, and 2785, all taken with a 50 Hz solicited trigger. Separate single-site and multi-site spectra are shown. The elliptical “island” near the upper right of each spectrum shows the anti-correlation between charge and light. (An approximate calibration for each scale has been applied for graphical convenience only.)

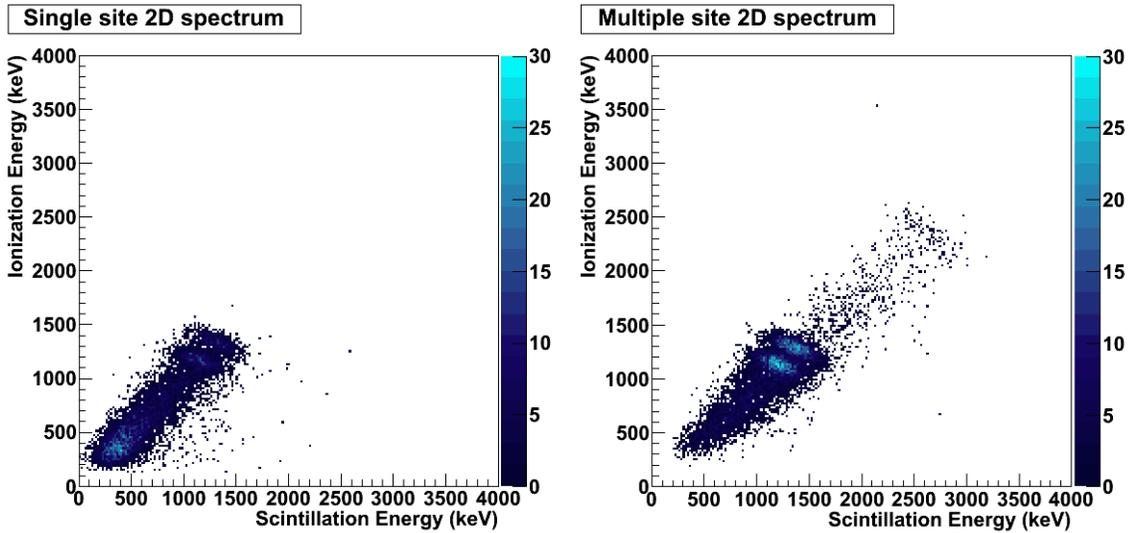


Figure 6.21: Two-dimensional charge-light spectrum of weak ^{60}Co calibration data from runs 3431 and 3433, taken with a 50 Hz solicited trigger. Separate single-site and multi-site spectra are shown. Two elliptical anti-correlated “islands” can be seen, one for each gamma peak. (An approximate calibration has been applied for graphical conveniences only.)

of $E_{rotated}(\theta)$ is formed and fit with the usual sum of a gaussian and erfc (Equation 6.12). The resolution is extracted and plotted as a function of rotation angle, and the resulting curve is fit to a parabola, shown in Figure 6.22. The angle that minimizes the resolution of the full-absorption peak is taken as the optimal angle; angles of $\theta_{SS} = 0.1814$ rad and $\theta_{MS} = 0.2050$ rad are found for the single and multi-site data, respectively. The angles are simply parameters and are used without uncertainties. The rotation can then be carried out, forming the energy variables for analysis.

$$E_{rotated} = \begin{cases} E_Q \cos \theta_{SS} + E_S \sin \theta_{SS} & \text{(single-site)} \\ E_Q \cos \theta_{MS} + E_S \sin \theta_{MS} & \text{(multi-site)} \end{cases} \quad (6.28)$$

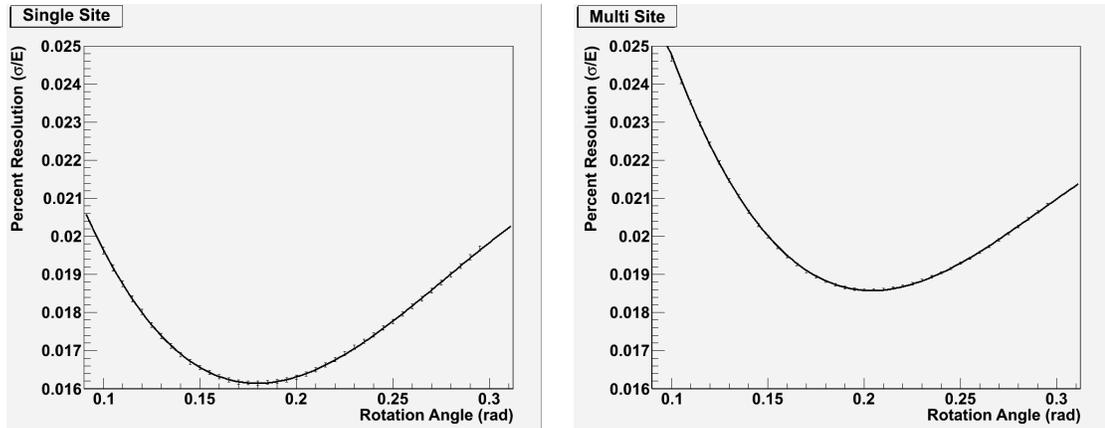


Figure 6.22: Variations in the energy resolution of the ^{228}Th full-absorption peak with the charge-light rotation angle, θ . Optimal values are found by fitting to a parabola (black curve) and finding the minimum. Angles are found separately for single-site and multi-site spectra to be $\theta_{SS} = 0.1814$ rad and $\theta_{MS} = 0.2050$ rad, respectively.

6.8.2 Initial rotated energy scale calibration

The energy scale is determined using the four γ peaks from the calibration sources, listed in Table 3.1. These span the energy range from near the low-energy threshold to above the Q-value of $\beta\beta$. Data from the strong sources is taken from the calibration campaign, using the following runs:

1. ^{228}Th : 2424, 2426, 2431, 2432, 2433, 2434, 2447, 2448
2. ^{60}Co : 2526, 2538, 2543, 2555, 2566, 2578, 2596, 2608, 2620, 2634, 2635, 2640, 2646, 2653, 2667, 2683, 2689, 2708
3. ^{137}Cs : 2450, 2469, 2473.

Spectra of $E_{rotated}$ are formed and the full-energy peaks are fit with the usual gaussian and erfc fit functions to extract the mean energy. The four peak energies are fit to a quadratic function to yield a preliminary energy scale calibration, and this calibration is applied to each source spectrum.

6.8.3 Rotated energy resolution

The detector energy resolution must be characterized for accurate PDF generation. In general, the resolution can be characterized by a three-parameter function of energy, given by Knoll [80].

$$\begin{aligned}\sigma^2(E_{rotated}) &= \sigma_{stat}^2 + \sigma_{noise}^2 + \sigma_{drift}^2 \\ &= (\sqrt{wFE_{rotated}})^2 + (w\sigma_e)^2 + (kE_{rotated})^2.\end{aligned}\tag{6.29}$$

The first term, σ_{stat} , is due to statistical fluctuations in the number of electrons and photons collected; these numbers are proportional to the energy deposited, so the fluctuations scale as $\sqrt{E_{rotated}}$. The next term, σ_{noise} , is due to electronic noise, and is independent of deposit energy. The last term, σ_{drift} , arises from time variations in electronic gain that tend to broaden the resolution as $\sigma_{drift} \sim E_{rotated}$. These components are independent of one another and add in quadrature. The third form is expressed in terms of related physical parameters: $w = 0.0187$ keV is the W-value, or average energy required to liberate a single electron from liquid xenon; F is the Fano factor, which enters into the statistical fluctuations term; $\sigma_{E_{rotated}} = 800$ is the number of electrons produced by noise; and k is a constant parameterizing the amount of electronic drift.

To characterize the detector, the calibrated rotated spectra are refit with the same fit function. The best fit gaussian width is identified with the resolution, and the four resolutions are fit to a functional form derived from Equation 6.29:

$$\sigma(E_{rotated}) = \sqrt{r_0^2 E_{rotated} + r_1^2 + r_2^2 E_{rotated}^2}. \quad (6.30)$$

The measured resolutions for each radioactive source are shown in Table 6.5, and the best fit values for the parameters, r_i , are given in Table 6.6. Figure 6.23 shows the measured resolutions plotted with the best fit curve. In the ultimate fits to the low background spectrum, these parameters will be allowed to float, constrained by these measured uncertainties. Thus the uncertainty in the resolution measurement is treated as a systematic uncertainty in the final fits.

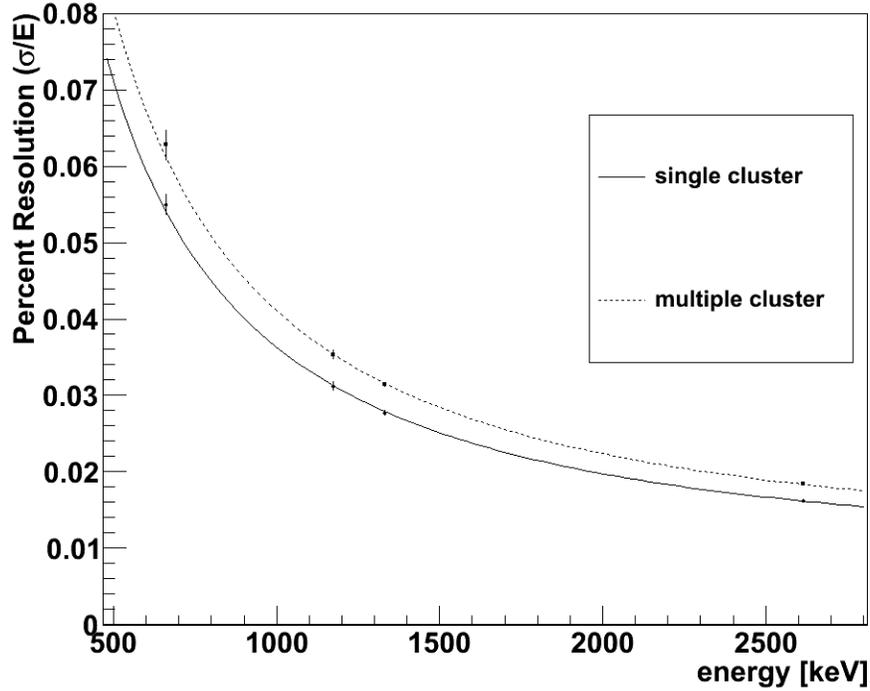


Figure 6.23: Measured energy resolutions of radioactive calibration sources. These are used to measure the detector energy resolution as a function of energy. The data are fit to a physically motivated functional form, given in Equation 6.29. Single-site and multi-site resolutions are measured separately. After [4].

Table 6.5: Measured values of the rotated energy resolution for each calibration peak.

| Source | σ/E (SS) (%) | σ/E (MS) (%) |
|----------------------|---------------------|---------------------|
| ^{137}Cs | 5.50 | 6.23 |
| ^{60}Co (1) | 3.12 | 3.53 |
| ^{60}Co (2) | 2.77 | 3.15 |
| ^{228}Th | 1.62 | 1.84 |

Table 6.6: Best fit values for the rotated energy resolution ($\sigma(E_{rotated})$) parameters.

| Number of Sites | r_0 | r_1 | r_2 |
|-----------------|-------|------------------|-------------------------------------|
| Single | 0 | 35.16 ± 0.34 | $8.95\text{e-}3 \pm 3.78\text{e-}4$ |
| Multi | 0 | 39.7 ± 0.39 | $1.03\text{e-}2 \pm 3.38\text{e-}4$ |

6.8.4 Energy calibration fitting bias

The gaussian/erfc fit function in Equation 6.12 was chosen as a phenomenological model of a Compton shelf with a full-absorption peak. Non-gaussianities as well as differences between the shape of the erfc and the true shape of the Compton shelf lead to a potential bias in the measured mean energy of the full-absorption peak.

The bias is searched for using MC energy spectra. An arbitrary energy resolution can be imposed by convolving a spectrum with gaussians of a certain width. We can compare with the resolution returned by the fit; any discrepancy indicates a bias. By varying this width, we can generate a sample of spectra at known resolutions, and measure the bias at each width.

This test is run using MC calibration source decays. In fact, a systematic bias does appear at about 0.5%, with some variation with energy resolution. We can use the trend to infer the bias at any given resolution; for each source, the trend is described. The difference between the measured and known full-absorption peak energy is taken as the bias, given for each source in Table 6.7.

6.8.5 Final rotated energy scale calibration

A final energy scale calibration is performed to incorporate the bias. Each of the four fitted full-absorption peak energies is corrected by adding the fit function bias, and the new peak energies are fit again to a quadratic. The calibrated energy,

$E_{rotated}^{cal}$, can then be found by simply applying the calibration:

$$E_{rotated}^{cal} = a_0 + a_1 E_{rotated} + a_2 E_{rotated}^2, \quad (6.31)$$

where a_i are the parameters of the fit. The best fit parameters are shown in Figure 6.24 and Table 6.8; the quadratic term is quite small. These calibrations are applied to low background data and will be used in generating PDFs. Calibrated, rotated spectra for calibration sources and low-background data are shown in Figures 6.25 and 6.26.

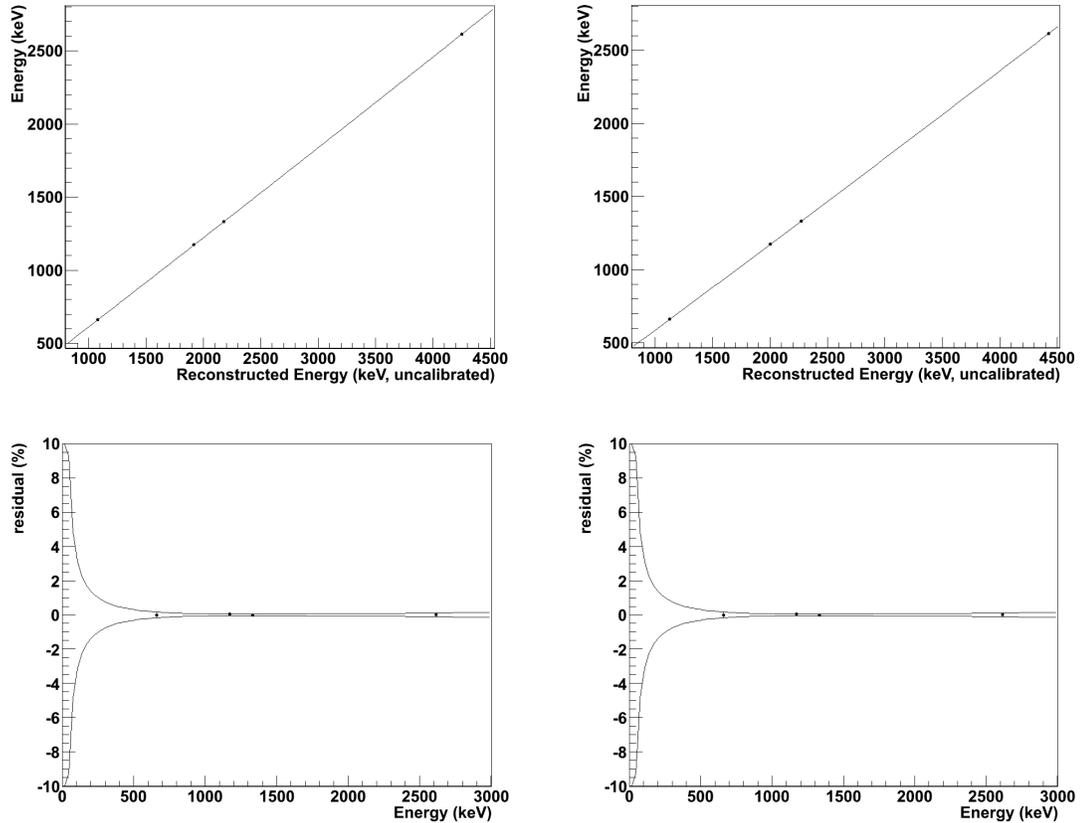


Figure 6.24: Second order polynomial rotated energy calibration curves for the four calibration source gammas, single-site on the left and multi-site on the right. Residuals are shown with error bands. Best fit parameters and uncertainties are given in Table 6.8.

Table 6.7: Fit function bias for each source in single and multisite. Biases are all given in keV.

| No. of Sites | ^{137}Cs | ^{60}Co (1) | ^{60}Co (2) | ^{228}Th |
|--------------|-------------------|----------------------|----------------------|-------------------|
| Single | 13.19 | 16.13 | 13.19 | 14.27 |
| Multiple | 18.69 | 18.6 | 17.86 | 18.32 |

Table 6.8: Quadratic energy scale calibration parameters. The calibrated energy, $E_{rotated}^{cal} = \sum_i a_i E_{rotated}^i$. Single-site and multi-site are calibrated separately.

| No. of Sites | a_0 | a_1 | a_2 |
|--------------|------------------|---------------------|--------------------------------------|
| Single | 9.4 ± 3.3 | 0.6010 ± 0.0032 | $2.516\text{e-}6 \pm 6.47\text{e-}7$ |
| Multiple | 11.56 ± 3.91 | 0.5734 ± 0.0031 | $3.257\text{e-}6 \pm 5.74\text{e-}7$ |

6.8.6 Effect of bias on β energy scale

The fit function bias may partially explain the observed energy scale discrepancy in the charge calibration (see Section 6.5.1). Since the double escape peak is not a photon interaction, it has no Compton shelf, meaning the bias does not apply to it. The bias parameters can also be measured using the charge only spectrum. Applying the corrections, we find that the double-escape peak shifts closer to the γ s, offset by 1.4% from the best fit calibration. This can be compared to 1.9% before the correction.

We observe that β and γ deposits may have different spatial extent, which would lead to differences in the induction signals. We can limit induction effects by only consider charge clusters collected on a single U-wire, for which only one induction signal is relevant. Combining this with the bias correction, we find that the double-escape peak becomes consistent with the γ calibration, as shown in Figure 6.27. It may therefore be possible to resolve the discrepancy of the energy scales by

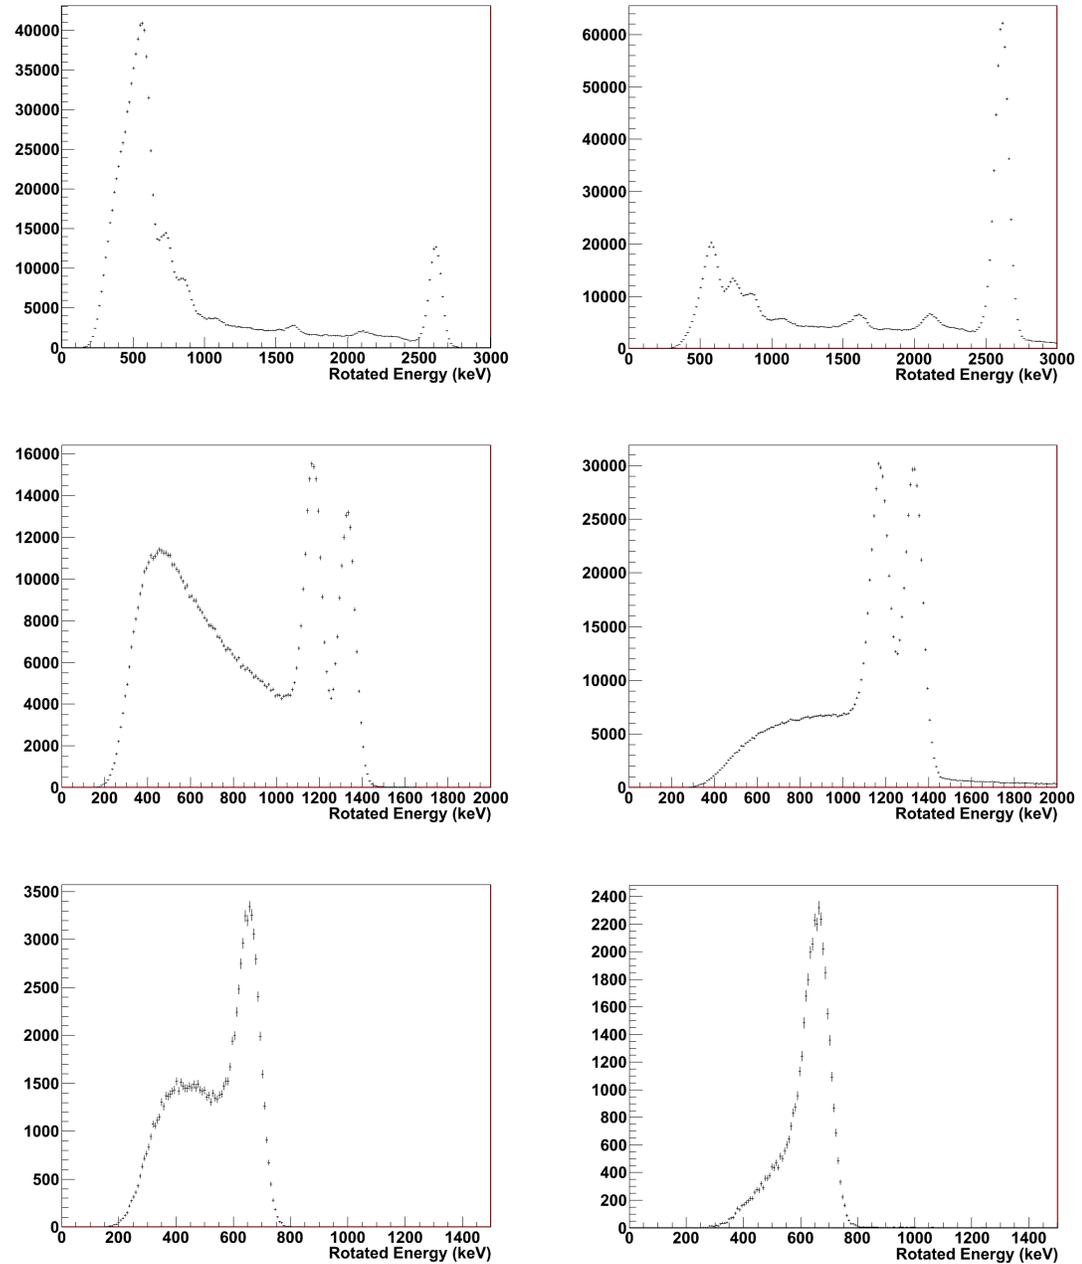


Figure 6.25: Optimized rotated energy spectra for the three strong radioactive calibration sources. From top to bottom: ^{228}Th , ^{60}Co , and ^{137}Cs Single-site (left) and multi-site (right) spectra are calibrated separately.

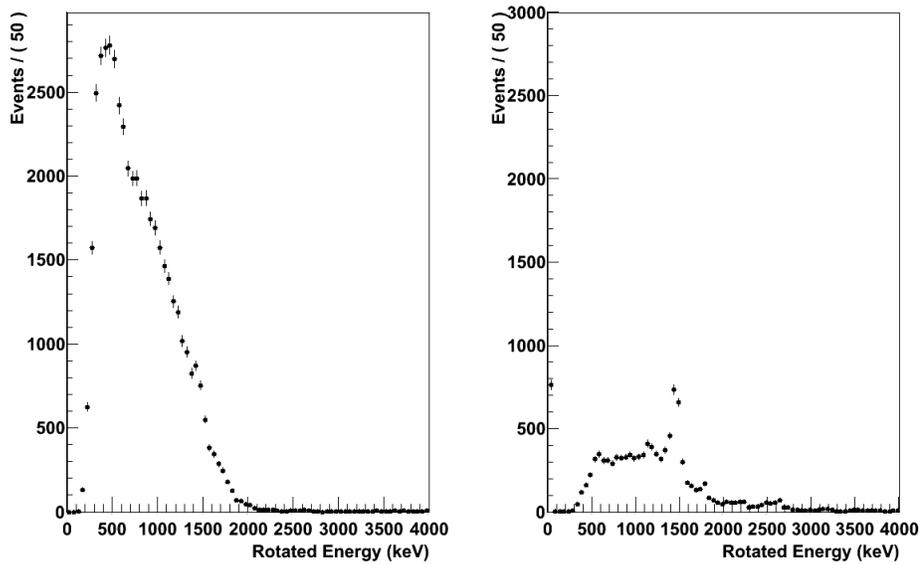


Figure 6.26: Rotated energy spectra for low-background data, with single-site on the left, and multi-site on the right. Data selection criteria have not yet been applied.

making a correction for signals collected on two or more U-wires. For the present analysis, we must retain multiple U-wire signals to increase statistics, so we attempt to account for the energy scale differences as a systematic.

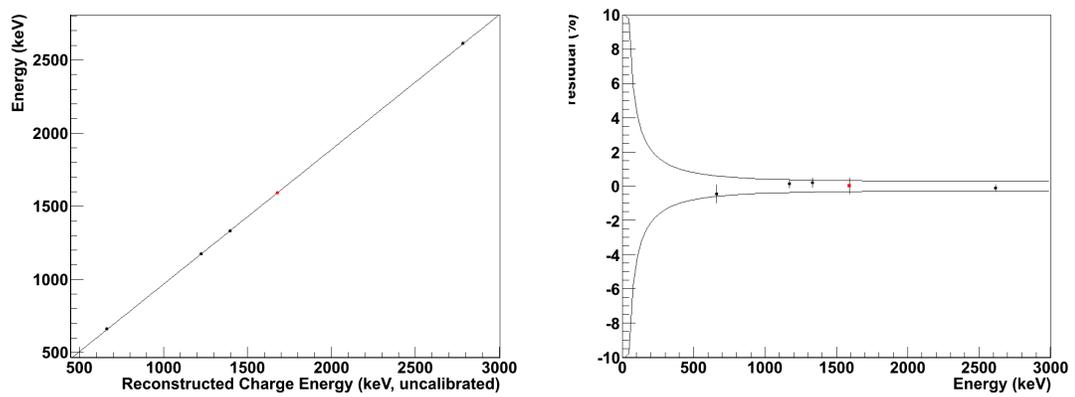


Figure 6.27: Linear calibration curves for the charge energy, created using the four calibration source gammas. The peak-fitting bias correction has been applied, and only single U-wire clusters are allowed. The best fit single-site energy scale is on the left, and the residuals are on the right. The red points indicate the ^{228}Th double-escape peak, which now aligns with the other sources.

Chapter 7

Fitting

Maximum likelihood fits to energy-based MC PDFs are used to search for $\beta\beta 0\nu$ in the low-background data. We describe the generation of these PDFs from the MC data sets. In order to incorporate the measured model of the detector energy resolution, the PDFs are generated from the PCD spectra¹, rather than from reconstructed spectra. Selection criteria for the low-background and MC datasets are applied to both sets when possible to ensure similarity in spectral shape. Special care is taken to exclude energy from clusters in the dead region, a process we will refer to as *culling*. The entire analysis process, including reconstruction, calibration, and data selection, is validated by fitting calibration source spectra to the corresponding MC spectra. Uncertainties can be estimated, either by taking profiles of the final likelihood function or by specialized means. Fitting is carried out using the ROOFIT framework, a package based on the ROOT data analysis framework [81, 82].

7.1 Culling of charge clusters

EXOSIM does not contain a complete three-dimensional model of the electric field in the TPC, instead using a two-dimensional uniform electric field to generate signals. In the physical TPC, regions at large radius may have non-uniform electric

¹Pixelated Charge Deposit; see Chapter 5

field. Charge deposited in this *dead region* is either not collected or only partially collected (see Section 3.3.3). To ensure that the MC accurately models the charge collected by the TPC, we remove, or *cull*, from the simulated datasets all clusters deposited in the dead region. Not all clusters in the dead region of the physical TPC are missed, so the MC *culling* is somewhat stricter than the charge collection loss in the physical TPC. For consistency, then, the low-background data is also *culled*. For multisite events, only the clusters in the dead region are culled, and the rest are unaffected. Culled clusters are considered removed from the event; no cuts reference culled clusters, and their energy is removed from the total event energy. We define the *culled energy*, E , to be the sum energy of either all PCDs (for MC) or all charge clusters (for data) that pass culling.

$$E = \sum_{\text{culled PCDs}} E_{PCD} \quad (\text{MC}) \quad (7.1)$$

$$E = \sum_{\text{culled clusters}} E_{rotated}^{cal} \quad (\text{data}) \quad (7.2)$$

We cull PCDs from the MC and charge clusters from TPC data, but the reconstructed position of the deposit is used as the selection criterion in either case. This is simply for uniformity. For selection purposes, a PCD is identified with a charge cluster if both have the same collection channel and the collection times differ by less than $3.5 \mu\text{s}$. For this analysis, clusters were culled that were at larger radius than the Teflon reflectors ($r > 183 \text{ mm}$), where it is clear that the electric field will be non-uniform. As described in Section 3.3.3, there are regions inside

the reflectors but outside the hexagon defined by the collection wires, which are candidates for inclusion in the dead region. However, doing so results in poorer agreement between calibration source data and MC, so clusters in these regions were not culled for the final analysis. Three-dimensional simulations of the TPC electric field are required to fully understand charge collection in these regions.

7.2 Event selection

7.2.1 Low-background data set

The low-background data used in this analysis were taken in a dedicated campaign between September 22, 2011 to April 15, 2012. These dates correspond to runs 2379-3563 and were selected for detector stability. The TPC was filled with enriched xenon for the duration. During this time, the detector was operated with a 0.1 Hz forced trigger. The livetime can then be determined by counting the number of triggers, and was found to be 2896.6 hours. The average trigger rate was ~ 0.2 Hz.

To ensure that only high-quality data was included, these runs were further selected. Microphonics caused by loud noises in or near the clean rooms caused high TPC event rates, so runs with unusually high events rates were rejected. Periodic fire alarm testing occurred during some runs; the duration of the alarm was recorded and events during those times were rejected, but the rest of the data was used.

7.2.2 Data masking

To reduce the potential for analyzer bias, the low-background data set was *masked* during the analysis, so that only one-third of the data was available. In this way, the information about events in the region near the $\beta\beta$ Q-value was reduced. Calibration data was not masked, since any expected $\beta\beta$ signal is completely dominated by the source rate.

For every low-background data file, a new, masked file was created, and analyzers only accessed the masked files while analysis was ongoing. A masking algorithm selected the events from the first 10 minutes out of every 30 minutes of the original data set to pass to the masked file. This choice allowed for potential time-varying effects to be identified, even in the masked data.

Once work on the analysis was “frozen”, including all reconstruction, data selection, fitting, and any other algorithms, the data were “unmasked” and the analysis was applied to the original low-background data files. None of these algorithms were changed after unmasking.

7.2.3 Fiducial cut

Due to the radioactive content of the detector materials, especially the signal cables (see Section 3.10), the bulk of background events are located at large radius or near the electrodes. The fiducial cut removes charge clusters outside a hexagonal solid region with apothem $r = 163$ mm. Clusters from the regions around the cathode and anode are also removed, keeping ones with $5 \text{ mm} < |z| < 182 \text{ mm}$.

This fiducial volume corresponds to $3.258 \cdot 10^4 \text{ cm}^3$. The density of the enriched xenon is 3.023 g/cm^3 at the operating temperature, giving a fiducial mass of 98.5 kg of active enriched LXe, or 79.4 kg of ^{136}Xe (accounting for the 80.6% enrichment factor). Multisite events are classified as non-fiducial if any of the associated charge clusters are found outside the fiducial region. Note that the entire event is removed, as opposed to cluster culling, in which only single clusters are excised. Culled clusters are considered omitted and cannot cause an event to fail the fiducial cut.

Clusters are also considered as non-fiducial if any one position variable could not be determined by reconstruction. The bulk of clusters failing for this reason are due to missing V-wire positions. The efficiency for V-position reconstruction was discussed in Section 4.2.1 (Figure 4.9) and shown to be energy dependent, with a low-energy threshold. This threshold is not seen in the MC, so in order to more accurately model the spectral shape of the data, we impose one. MC charge clusters are randomly labeled as having invalid V-position, with a probability given by the efficiency at the energy of the cluster. Events containing these clusters are then removed by the fiducial cut.

7.2.4 Muons tagged by veto panel

As discussed in Section 3.7, EXO-200 incorporates muon veto panels to identify events due to cosmic rays. Muon panel event rates were constantly monitored to ensure stability. Events in coincidence with muon panel triggers are cut; muon panel triggers up to 25 ms before the event or up to 1 ms after the event constitute

a coincidence.

7.2.5 Muons tagged by TPC

A second approach to tagging muons uses TPC data. We can exploit the fact that muons depositing energy in the TPC typically leave long linear tracks, with many wires collecting charge sequentially in time. An example muon track can be seen in Figure 7.1. Muons also typically have greater energy than radioactive decays. TPC events with large collection signals are flagged, and the charge deposit locations are subjected to a linear Hough transform to identify tracks. Events with strongly linear charge deposits are tagged as muons. Events up to 60 s after a muon track or up to $1 \mu\text{s}$ before a track are removed.

$96.0 \pm 0.5 \%$ of TPC-tagged muons are coincident with veto panel signals, so this is taken as the panel efficiency [47]. Simulations show a comparable inefficiency due solely to the geometry of the panels [65].

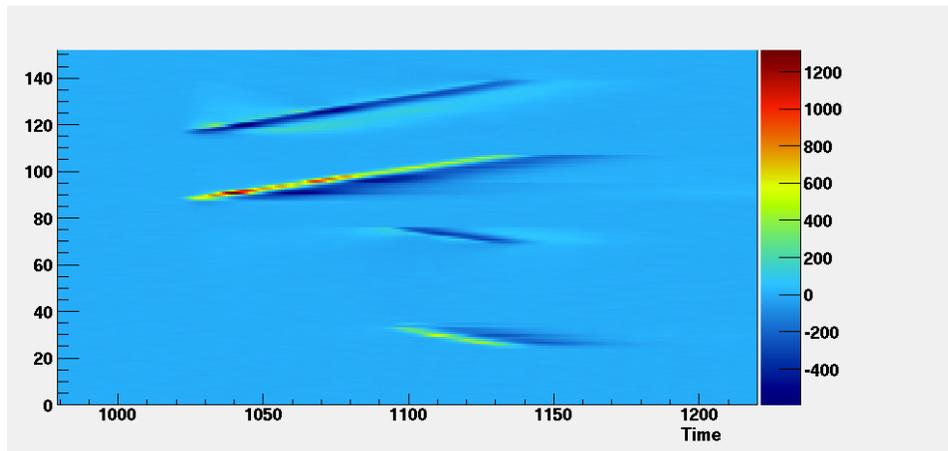


Figure 7.1: A typical muon event in the TPC. Four bands can be seen, each indicating the signal seen on either U or V wires in one half-TPC.

7.2.6 Correlated noise

Correlated noise is identified in the detector across a wide frequency spectrum. In order to eliminate noise events, a 1 s TPC event/TPC event coincidence cut is imposed. As the low-background trigger rate is only ~ 0.2 Hz, the dead time introduced is not significant.

Table 7.1: List of cuts based on coincidence of TPC events with muons or other TPC signals.

| Coincidence event | Coincident time |
|-------------------|-----------------|
| Muon veto panel | 25 ms |
| TPC muon track | 60 s |
| TPC event | 1 s |

7.2.7 Scintillation/ionization ratio

The ratio, R , of scintillation to ionization energy is determined by the recombination physics, so events of the same decay type (β , γ , etc.) should generally appear with the same value of R .

$$R = E_S/E_Q, \tag{7.3}$$

where E_S and E_Q are the scintillation and ionization energies, respectively. Two dimensional calibration spectra consist mostly of events with similar ratios R , but also show a population of events with anomalously low charge energy relative to scintillation energy. The weak Th spectrum in Figure 7.2 indicates this population below and to the right of the red diagonal line. These events may have incomplete

charge collection, which may occur when one or more charge clusters are not properly associated with the scintillation cluster, for example. Missing charge energy would lead to falsely low values for the rotated energy, so these events are removed from the low-background data set. Note that no population of events with relatively high ionization energy is seen, so no cut is made.

The cut criteria are determined using the rotated energy spectrum from each calibration source, as follows. The full-energy peak is fit to a Gaussian, and the width, σ_R , and mean, μ_R , are measured. Events with energy near the mean, in the range $(\mu_R - 4\sigma_R, \mu_R + 4\sigma_R)$, are selected, and the scintillation energy distributions of these events is then formed. This distribution is also fit to a Gaussian to obtain the width and mean, σ_S and μ_S . Events with scintillation energy deviating more than $4\sigma_S$ from the mean are considered anomalous:

$$E_S^i > E_{S_{max}}^i = \mu_S + 4\sigma_S \quad (7.4)$$

and $0 < i < 4$ ranges over each source. Combining the values of $E_{S_{max}}^i$ for each source gives a line in the two-dimensional ionization/scintillation spectrum. It is convenient to make the cut on the raw, uncalibrated energies, E_S and E_Q , since these are ultimately used in processing. Note, however, that the calibration in Section 6.5 is applied to make Figure 7.2. The cut line is given by

$$E_S = \begin{cases} 2600.67 + 3.405 \times E_Q & \text{(single-site)} \\ 2876.03 + 3.630 \times E_Q & \text{(multi-site)} \end{cases} \quad (7.5)$$

Events with E_S above this line are removed.

A secondary effect of this cut is to remove remaining α decays. As discussed in Section 3.1, α decays result in much more electron/ion recombination, and hence much larger scintillation/ionization ratios, than β s or γ s. This can be seen in the single-site low background spectrum in Figure 7.3.

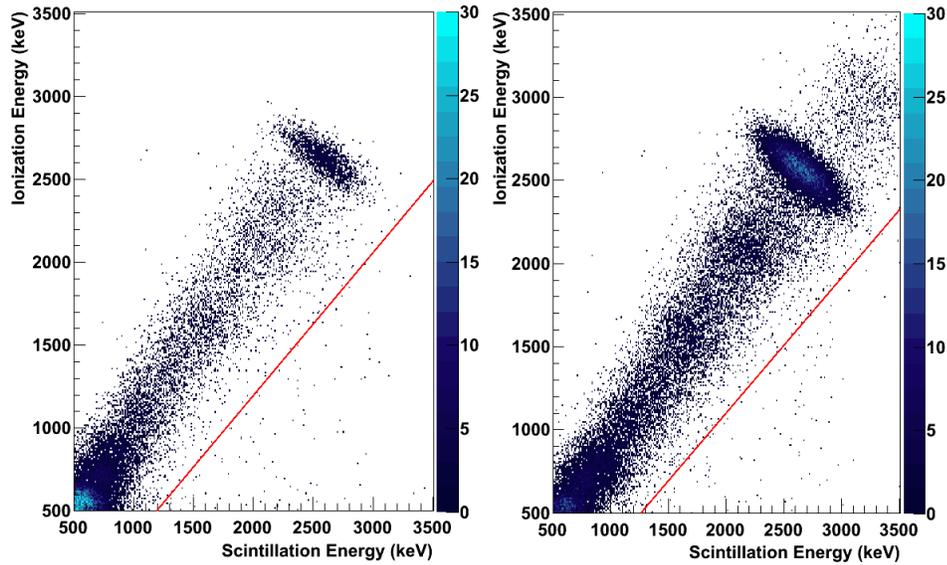


Figure 7.2: Weak thorium calibration data showing examples of events with anomalously high ratios of scintillation to ionization energy. These events are below and to the right of the red line, which indicates a cut placed on low-background data. The cut is slightly different for single-site (left) and multi-site (right). All selection cuts have been applied except the scintillation/ionization ratio cut.

7.2.8 Selection results

The effects of each cut on the single-site low background spectrum can be seen in Figure 7.4. The final spectrum used for fitting is shown in blue. The cuts are applied sequentially; the number of events removed from the data set when each cut is applied is shown in Table 7.2. Most of the cut events fail the fiducial volume

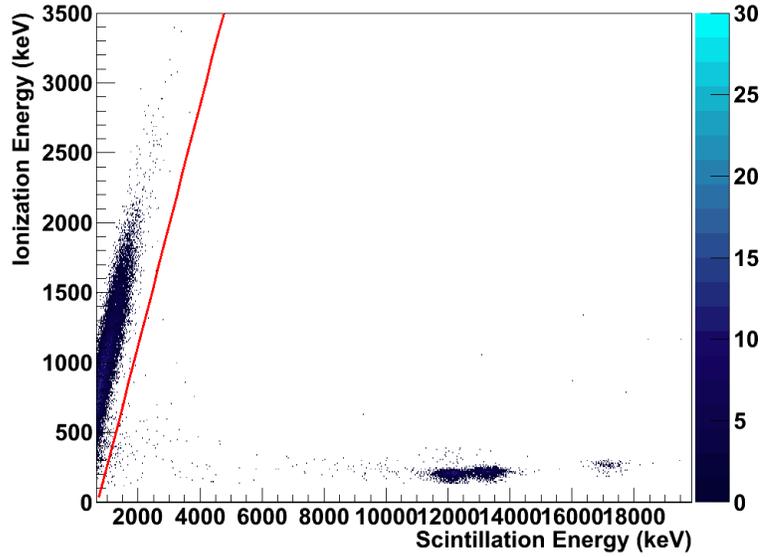


Figure 7.3: Single-site low background spectrum emphasizing α decays in the bottom right. These decays have much greater recombination and yield much stronger scintillation and weaker charge signals than the β s or γ s at the left of the figure. Hence they are removed by the cut indicated by the red line.

cut, providing evidence that most backgrounds are on the surface of the electrodes or originate outside the fiducial region. Since detector surfaces are expected to contain both α and β or γ emitting particles, we can compare the effects of the scintillation/ionization ratio cut with the fiducial cut to identify which types of events are removed.

In Figure 7.5, we can see that the the fiducial cut is more strict. Above 1500 keV, most events that are removed fail both cuts, indicating that α decays predominate in this region. This is consistent with α decay energies of known backgrounds, which are generally above 3000 keV, and which lose energy rapidly in the surrounding material. Below 1500 keV, the ratio cut is largely ineffective, consistent with β and γ decays. Still, close examination of the energy range near the $\beta\beta 0\nu$ Q-value of 2458 keV, seen in the inset of 7.5, shows that both cuts are

removing distinct sets of events.

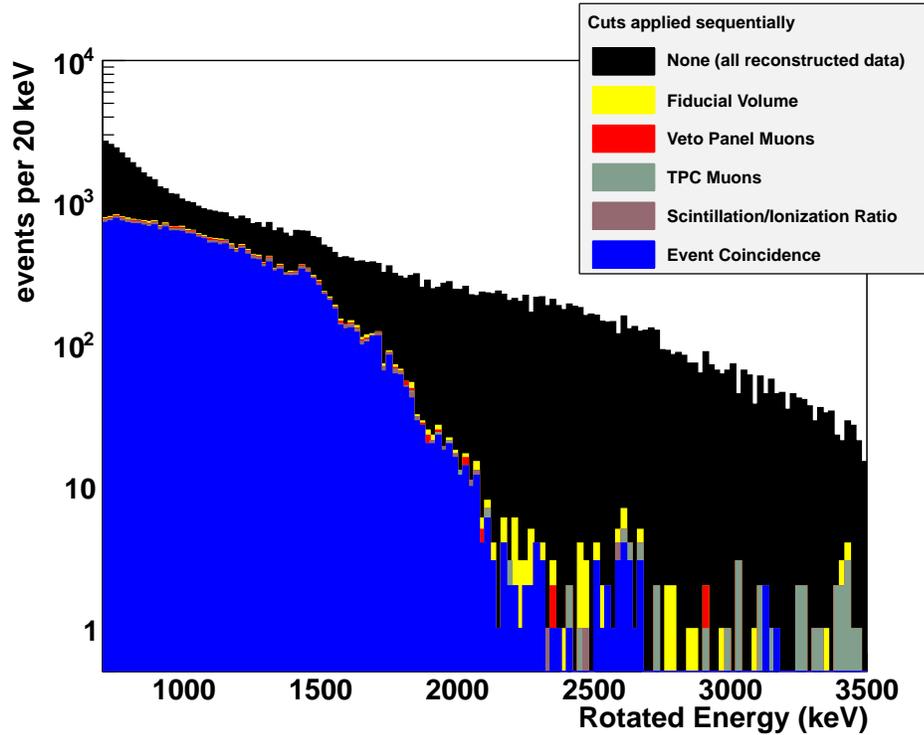


Figure 7.4: Single-site low background spectrum showing the effects of each cut. Cuts are applied sequentially; labels indicate that cut applied in addition to the previous cuts. Most removed events fail the fiducial cut, as most backgrounds are on detector surfaces or are external to the TPC. The spectrum with all cuts applied (in blue) is used for fits.

7.3 Efficiencies

Detection efficiencies for each of the various signal and background processes can be determined from the MC spectra. Using simulated data allows us to include efficiencies for rare processes that cannot be directly measured, as well as discriminating between decay locations. The efficiency for a process is given by the ratio

$$\epsilon = N_{spec}/N_{gen} \tag{7.6}$$

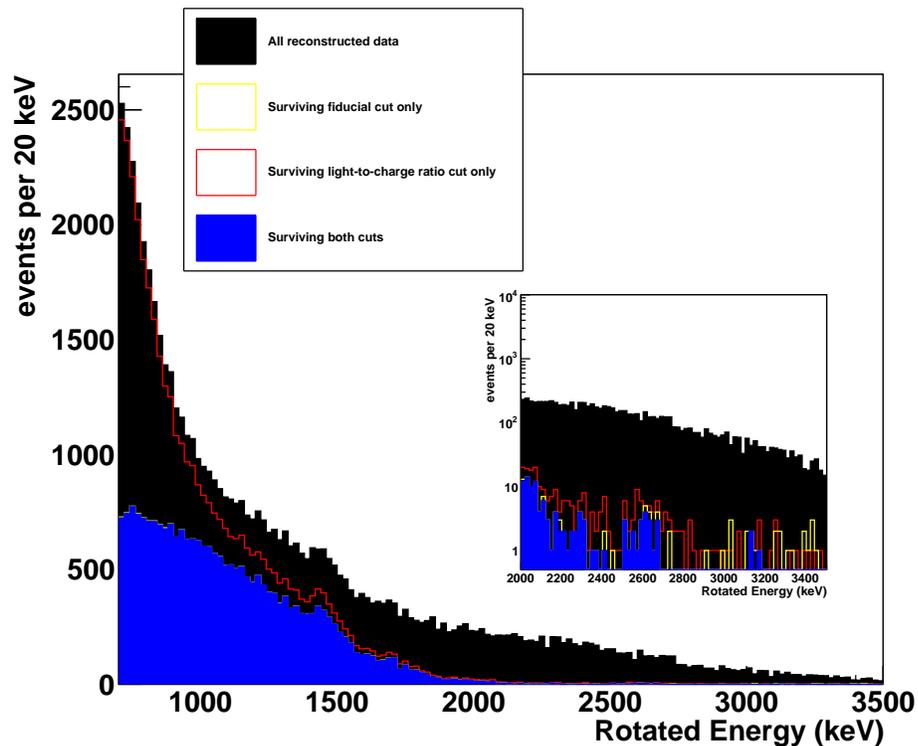


Figure 7.5: Comparison of the effects of scintillation/ionization ratio cut with the fiducial cut on the single-site low background spectrum. A significant number of events fail both cuts, especially above 1500 keV; these are due to α decays on detector surfaces. Below 1500 keV, most events only fail the fiducial cut indicating β s on surfaces or external γ s. The inset emphasizes the $\beta\beta 0\nu$ region of interest. The veto and coincidence cuts are applied in all four spectra shown.

Table 7.2: Summary of events removed from the data set by each cut. The first column gives the number of events removed sequentially from the each histogram in Figure 7.4. The second column gives the total number of events in the low-background data set failing that cut. Since cuts overlap, the second column sums to more than the total number of events.

| | | | |
|--------------------------------|----------------------|-------------|----------------------|
| Initial data set (no cuts) | 62475 events | | |
| Cut (number of events) | Removed sequentially | Failing cut | Passing sequentially |
| Fiducial volume | 37548 | 37412 | 24297 |
| Veto Panel Muons | 472 | 871 | 24455 |
| TPC Muons | 429 | 1005 | 24026 |
| Scintillation/Ionization Ratio | 102 | 20429 | 23924 |
| Event Coincidence | 828 | 2170 | 23096 |

of the number of events in the final spectrum after data selection, N_{spec} , to the total number of events, N_{gen} .

7.4 PDF generation

After culling and data selection, PDFs for the various MC energy spectra can be created. PCD energies are used to form initial spectra. By using the MC generated energies, and not the reconstructed energies, we can validate reconstruction at the same time as the simulated detector construction. This approach requires applying an energy resolution model to the MC spectra. We explicitly include energy calibration parameters in the PDFs so that they may float in the final fits. It is for this reason that the PDFs are generated from the PCD spectra, rather than reconstructed spectra, as the reconstruction includes a simplified resolution model which does not float. Similarly, the energy scale calibration parameters are included in the PDF. This enables more robust spectral fits between data and MC, and allows for studies of the systematic uncertainties associated with the energy resolution determination.

7.4.1 Applying energy resolution

To apply the effects of resolution, a gaussian PDF, $g(E)$, is defined with an energy-dependent width given by $\sigma(E)$ in Equation 6.29. This gaussian is convolved with the initial MC spectra, but the energy-dependence of σ requires a modification.

To illustrate the difficulties, consider a gaussian

$$g(E) = e^{-E^2/2\sigma^2(E)}; \quad (7.7)$$

the ordinary convolution with an arbitrary PDF $h(E)$ would be

$$\begin{aligned} (g \star h)(E) &= \int g(Q)h(E - Q)dQ \\ &= \int e^{-Q^2/2\sigma^2(Q)}h(E - Q)dQ, \end{aligned} \quad (7.8)$$

where Q is some dummy variable. However, in this case the effects of the width are integrated out: the functional form of σ over the entire range of Q affects the final result. Instead, only the form of σ near the energy E where the PDF is evaluated should be relevant. For example, if we select the PDF to be an arbitrarily narrow peak at some energy, E_0 , $h(E) = \delta(E - E_0)$, then we find

$$\begin{aligned} (g \star h)(E) &= \int e^{-Q^2/2\sigma^2(Q)}\delta(E - E_0 - Q)dQ \\ &= e^{-(E-E_0)^2/2\sigma^2(E-E_0)}. \end{aligned} \quad (7.9)$$

Here, σ is evaluated not at the energy of interest, E , but at the relative energy $E - E_0$.

One possible solution is to use a conditional PDF with a second energy variable.

If we define a gaussian, $g(E|\mu)$ as

$$g(E|\mu) = e^{-E^2/2\sigma^2(\mu)}, \quad (7.10)$$

where μ is the energy at which the resolution is evaluated, then the convolution with $h(E)$ would be

$$\begin{aligned}(g \star h)(E|\mu) &= \int g_\mu(Q)h(E - Q)dQ \\ &= \int e^{-Q^2/2\sigma^2(\mu)}h(E - Q)dQ,\end{aligned}\tag{7.11}$$

This form properly specifies the width separately from the integration variable. However, as a conditional PDF, it represents a family of PDFs for any value of μ , and a second PDF specifying the distribution of μ is required. In our case, μ is directly related to E , so this leads to circularities.

The implementation we use is to include the dummy variable, Q , directly in our definition of the gaussian:

$$g(E) = e^{-Q^2/2\sigma^2(E)}.\tag{7.12}$$

Then the convolution with $h(E)$ becomes

$$(g \star h)(E) = \int e^{-Q^2/2\sigma^2(E)}h(E - Q)dQ.\tag{7.13}$$

Now, if we convolve with an arbitrarily narrow peak, $h(E) = \delta(E - E_0)$, we find

$$\begin{aligned}(g \star h)(E) &= \int e^{-Q^2/2\sigma^2(E)}\delta(E - E_0 - Q)dQ \\ &= e^{-(E-E_0)^2/2\sigma^2(E)}.\end{aligned}\tag{7.14}$$

This is a gaussian with width $\sigma(E)$ and mean equal to the mean of the δ -function,

as required. The modified convolution is carried out via Fourier Transforms using the interface of ROOT to the FFTW3 package [83]. Figure 7.6 shows an example of a MC spectrum and its convolution.

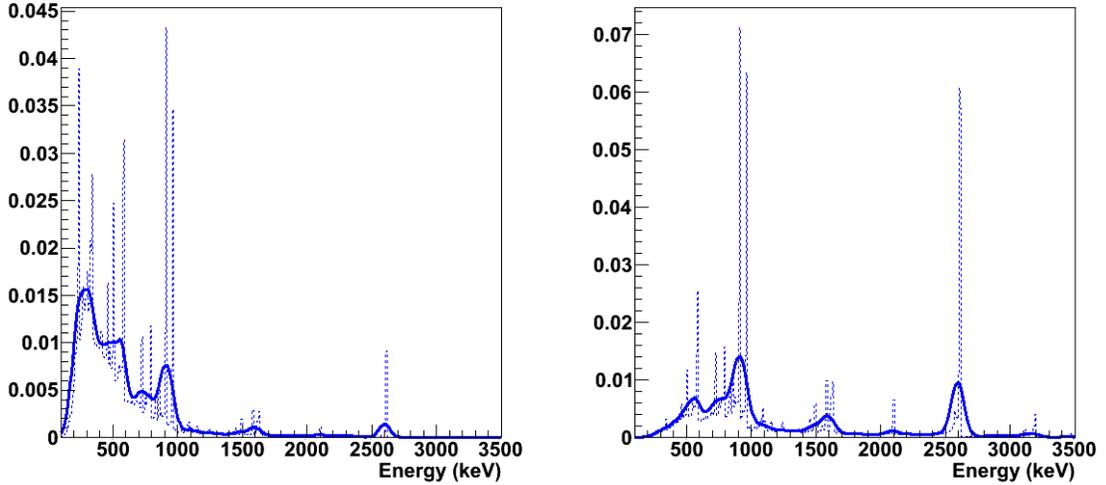


Figure 7.6: Simulated energy spectrum for ^{228}Th located in the copper TPC vessel (left: single site, right: multi-site). The dashed curve represents the PCD spectrum (E_{PCD}), and the solid curve shows the energy spectrum (E) after convolution with an Gaussian energy resolution of varying width.

7.4.2 Applying energy scale

The floating energy scale is introduced by modifying the energy term of the PDF. Corresponding to the quadratic calibration curve in Equation 6.31, three new terms, b_i , are used to indicate deviations of the calibration scale from the best fit parameters:

$$E_{floating} = E + (b_0 - a_0) + (b_1 - a_1)E + (b_2 - a_2)E^2. \quad (7.15)$$

Here, $E_{floating}$ is the modified energy, and the constant parameters a_i correspond to the best quadratic fit to the energy scale, shown in Table 6.8. The b_i are the floating parameters in the final fit; the form of Equation 7.15 allows for natural comparison of the best fit b_i to the initial parameters.

7.5 Monte Carlo validation

The MC can be validated with several comparisons are made between the generated PDFs and source calibration data. Comparisons of spectral shape show both that the Monte Carlo contains the correct physics processes and that the detector energy resolution is applied accurately. At the same time, reconstruction is also validated: the reconstructed source data is compared to convolved PCD spectra that have not been processed by reconstruction. By comparing the number of events in the single-site and multi-site spectra, we can check that the MC digitized waveforms are similar to the ones recorded in the TPC. Finally, the calibration source rates are compared to the MC. This validates the simulated detector geometry, since that strongly affects the rate at which source events are detected in the TPC.

7.5.1 Spectral Shape

The shape of the spectrum is validated by binned maximum-likelihood fits of the rotated weak source spectra to the corresponding PDF. Fits are shown in Figures 7.7 and 7.8 for ^{60}Co and ^{228}Th . We can evaluate the quality of the fit based on the reduced χ^2 (χ^2 per degree of freedom) of the fit. Although this only accounts

for statistical errors, which depends on the choice of binning, we find reasonable agreement for the binnings shown, with χ^2 values less than or close to 1.

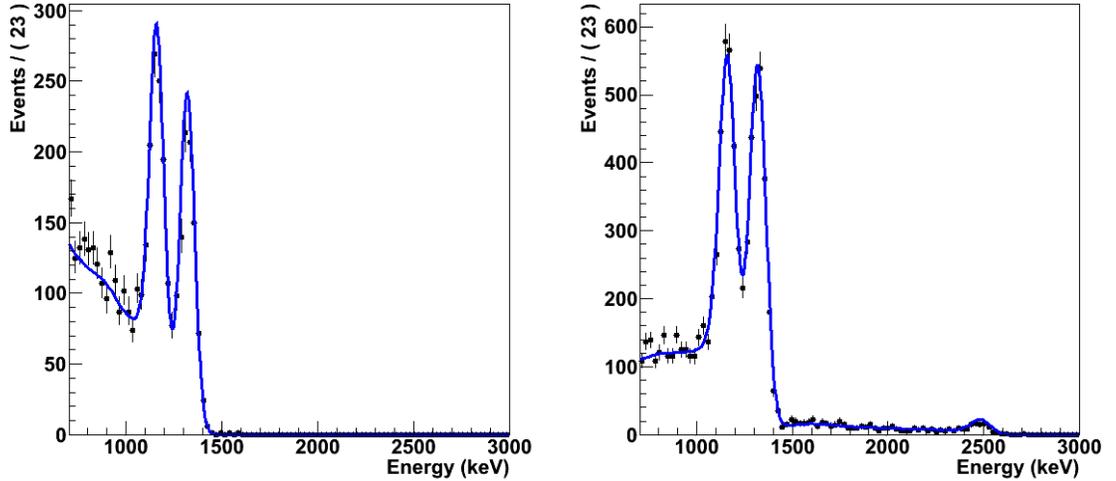


Figure 7.7: Calibration data taken with weak ^{60}Co source, fit to the ^{60}Co PDF (blue curve). Single site (left) and multi-site (right) spectra are fit simultaneously. The fit reduced χ^2 is found to be 0.5060 and 1.02 for single-site and multi-site, respectively. From [78].

7.5.2 Multiplicity

If the simulated signal shapes are different in the MC than in data, then reconstruction may classify a larger fraction as single-site or multi-site. To measure this, we compute the single-site fraction

$$f_{SS} = N_{SS}/(N_{SS} + N_{MS}) \quad (7.16)$$

for both data and MC spectra of ^{60}Co and ^{228}Th . Here N_{SS} and N_{MS} are the total number of events in the single and multi-site spectrum. Since discrepancies in different regions of the energy spectrum may offset, we compute f_{SS} for various

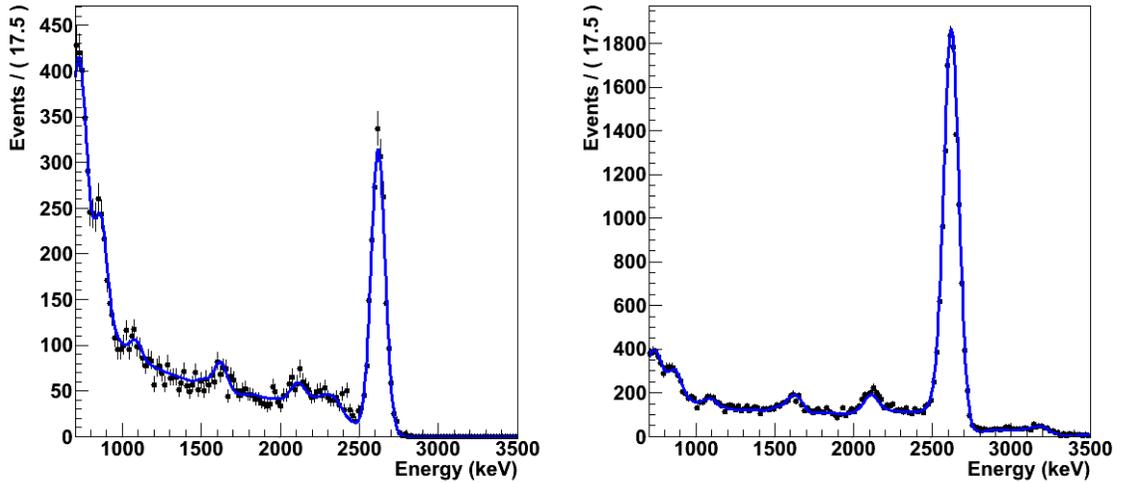


Figure 7.8: Calibration data taken with weak ^{228}Th source, fit to the ^{228}Th PDF (blue curve). Single site (left) and multi-site (right) spectra are fit simultaneously. The fit reduced χ^2 is found to be 0.8405 and 1.445 for single-site and multi-site, respectively. After [4].

energy ranges, shown in Table 7.3. This largest discrepancy will be used to constrain the ultimate fits.

7.5.3 Source rate

The rate of calibration source events observed is compared to the known source activities, which at the beginning of the calibration runs were $\Gamma_{Th}^0 = 643.17$ Bq for ^{228}Th and $\Gamma_{Co}^0 = 379.1$ Bq for ^{60}Co . This comparison validates the MC geometry and the reconstruction efficiency. The rate observed is calculated from the number of events, N , in the data set above threshold:

$$\Gamma_X = \alpha \frac{N}{\epsilon T}, \quad (7.17)$$

Table 7.3: Comparison of the single-site fraction, f_{SS} between data and MC. The comparison is made in several energy range to rule out cancellations between ranges. The maximum difference is 7.52 %.

| Source | Range (keV) | f_{SS} - MC | f_{SS} - Data | Difference |
|-------------------|-------------|---------------|-----------------|------------|
| ^{60}Co | 500-750 | 0.625 | 0.628 | 0.48% |
| | 750-1000 | 0.481 | 0.482 | 0.17 % |
| | 1000-1250 | 0.341 | 0.318 | 7.14% |
| | 1250-1500 | 0.289 | 0.270 | 6.88% |
| ^{228}Th | 500-1000 | 0.591 | 0.585 | 1.0% |
| | 1000-1500 | 0.378 | 0.357 | 5.74% |
| | 1500-2000 | 0.307 | 0.286 | 7.52% |
| | 2000-2500 | 0.243 | 0.236 | 2.80% |
| | 2500-3000 | 0.134 | 0.134 | 0.13% |

where α is a fraction accounting for a cut requiring only single scintillation clusters, ϵ is the detection efficiency, T is the livetime of the run, and the label X indicates the source. An energy threshold of $E > 700$ keV is used, and based on these studies, it is determined that this should be the threshold for the low-background fits. We can then calculate the rate agreement, ρ_X , as a fraction:

$$\rho_X = (\Gamma_X - \Gamma_X^0)/\Gamma_X \quad (7.18)$$

This gives $\rho_{Co} = 8.9\%$, and $\rho_{Th} = 9.4\%$.

7.6 Background model

The background model consists of PDFs of the single-site and multi-site energy spectra for each background. The materials selection campaign (see Section 3.10) informs the choice of PDF. The primary backgrounds are ^{238}U , ^{232}Th , and ^{40}K traces in the detector materials, as well as ^{222}Rn dissolved in LXe or in the air.

The spectral shape can vary depending on the original location of the decay, so for several backgrounds, PDFs are included for multiple locations. Rates for ^{222}Rn decays can be measured independently, leading to useful constraints. A full list of PDFs is shown in Table 7.4. PDFs having a primary β -like decay are identified for applying energy scale effects; see also Section 7.7.1.

Table 7.4: List of backgrounds included in the PDF model. The primary decay type of interest is noted to allow for separate treatment of the β -decay energy scale.

| Background | Location | Primary decay type |
|-------------------|-----------------|--------------------|
| ^{238}U | TPC vessel | γ |
| ^{232}Th | TPC vessel | γ |
| ^{60}Co | TPC vessel | γ |
| ^{65}Zn | TPC vessel | γ |
| ^{54}Mn | TPC vessel | γ |
| ^{40}K | TPC vessel | γ |
| ^{222}Rn | Active LXe | γ |
| ^{135}Xe | Active LXe | β |
| ^{222}Rn | Inactive LXe | γ |
| ^{214}Bi | Air gap | γ |
| ^{214}Bi | Cathode surface | γ |

7.6.1 Radioactive impurities in detector components

The material selection campaign found that ^{238}U , ^{232}Th , and ^{40}K impurities present in the copper TPC vessel, LAAPDs, and read-out cables may significantly contribute to the background rate. Simulations show that the shape of the energy spectrum these decays deposit is not strongly dependent on the location of the source. Thus, a single PDF generated from simulations in the copper TPC vessel is taken to include all these detector components.

7.6.2 Cosmogenic Activation

Cosmogenic activation of the copper TPC vessel is negligible when the TPC is shielded underground, but several isotopes with half-lives of ~ 1 yr could have been produced when the TPC was on the surface. To account for these, PDFs for ^{65}Zn , ^{54}Mn , and ^{60}Co are included in the model.

Activation of the xenon is also possible. Even short lived isotopes must be considered, since there is little shielding and no considerable loss of detection efficiency for events produced in the LXe. Accordingly, ^{135}Xe is incorporated in the background model.

7.6.3 Backgrounds external to the cryostat

Computational requirements limit the accuracy of simulations of sources outside the cryostat, but distance and the lead shielding limit their importance. The only one that is considered is the important ^{222}Rn in the air gap between the lead wall and the cryostat. To simplify computation, only the daughter ^{214}Bi , which emits the 2448 keV γ , is simulated. The decay rate of ^{222}Rn outside the cryostat is also measured with a Rad7 instrument and an upper limit of 42 ± 6.5 Bq/kg is set.

7.6.4 Dissolved Rn

We expect a small amount of ^{222}Rn to be emitted from Xe gas handling components. Since Rn is a noble gas, not removed by the purifiers, and since the LXe is constantly warmed and recirculated as a gas, we expect some Rn to constantly be

dissolved in the LXe. With a half-life of 3.8 days, ^{222}Rn is the isotope of concern. The ^{222}Rn decay is accompanied by coincident β - α decays of its daughters, ^{214}Bi and ^{214}Po . α decays are tagged by their large scintillation to ionization ratio, so the rate of ^{222}Rn decay can be measured by searching for coincidences. The overall rate, $\Gamma_{222\text{Rn}}$, is determined to be

$$\Gamma_{222\text{Rn}} = 360 \pm 65 \mu\text{Bq}. \quad (7.19)$$

The ^{222}Rn is mobile in the liquid; it is useful to further categorize its decays by location, to learn about the source of Rn. Decaying ^{222}Rn in the bulk of the TPC leaves behind positively charged daughter ions, which drift toward the cathode under the influence of the electric field. As the decay rate is significantly longer than the drift time (3.8 days versus $117 \mu\text{s}$), it is expected that most of the daughter ions will reach the cathode and subsequently decay there. 70% of measured Bi-Po coincidences do originate at the cathode. Accounting for losses due to decays into the cathode, it is estimated that 83% of ^{222}Rn decay in fact occur there. Consequently, 17% of decays occur in the bulk, active LXe. For precision, we can use two PDFs, one for Rn in the bulk, and one for Rn on the cathode. To simplify the interaction of the decay particles with the cathode, only the ^{214}Bi decay is simulated for the latter PDF.

A third PDF is used for the rate of ^{222}Rn in the inactive LXe, outside the teflon tiles. This can be inferred from the rate in the active Xe by using the mass ratio. The active LXe contains 129.8 kg, and the inactive 30.23 kg, so the rate is

constrained at $0.23 \Gamma_{222Rn}$.

7.6.5 ^{85}Kr

^{85}Kr is another radioactive noble gas that may remain dissolved in the LXe. However, the β Q-value of 687 keV is below the 700 keV threshold of the analysis, so ^{85}Kr is not included in the model.

7.6.6 Constraints

Various parameters in the PDFs represent measured quantities. These measurements can be used to constrain the overall PDF, typically with a multiplicative Gaussian constraint. For a general parameter q known with uncertainty $\pm\delta q$, the constraint factor, C_q , takes the form

$$C_q = e^{-(q-q_0)^2/2(\delta q)^2}, \quad (7.20)$$

where q_0 is the measured value of q .

Following is a list of constraints included.

1. ^{214}Bi in the air gap between lead wall and cryostat
2. Single-site/multi-site ratio
3. Energy scale (single-site and multi-site)
4. Energy resolution (single-site and multi-site)

The single-site/multi-site ratio is constrained separately for each component of the fit. The central value is found from the MC by taking the ratio of events in each spectrum above threshold. The relative uncertainty on the ratio for each component is the same, and is taken as the maximum discrepancy found from the studies of the calibration source multiplicity in Section 7.5.2.

The energy scale and energy resolution constraints are applied using the values and uncertainties determined in Section 6.8. Finally, the rate of ^{214}Bi in the air gap is constrained by the limit placed by the Rad7 measurement; see 7.6.3.

7.7 Maximum likelihood fit to low background data

The background model and the $\beta\beta$ PDFs are combined into a single PDF to be used in an extended maximum likelihood fit. Single-site and multi-site spectra are fit simultaneously. In total, 32 floating parameters are used: single-site and multi-site rate for 11 backgrounds and $\beta\beta 2\nu$; $\beta\beta 0\nu$ rate and single-site/multi-site ratio; one energy scale term and two energy resolution terms. Constraints are applied.

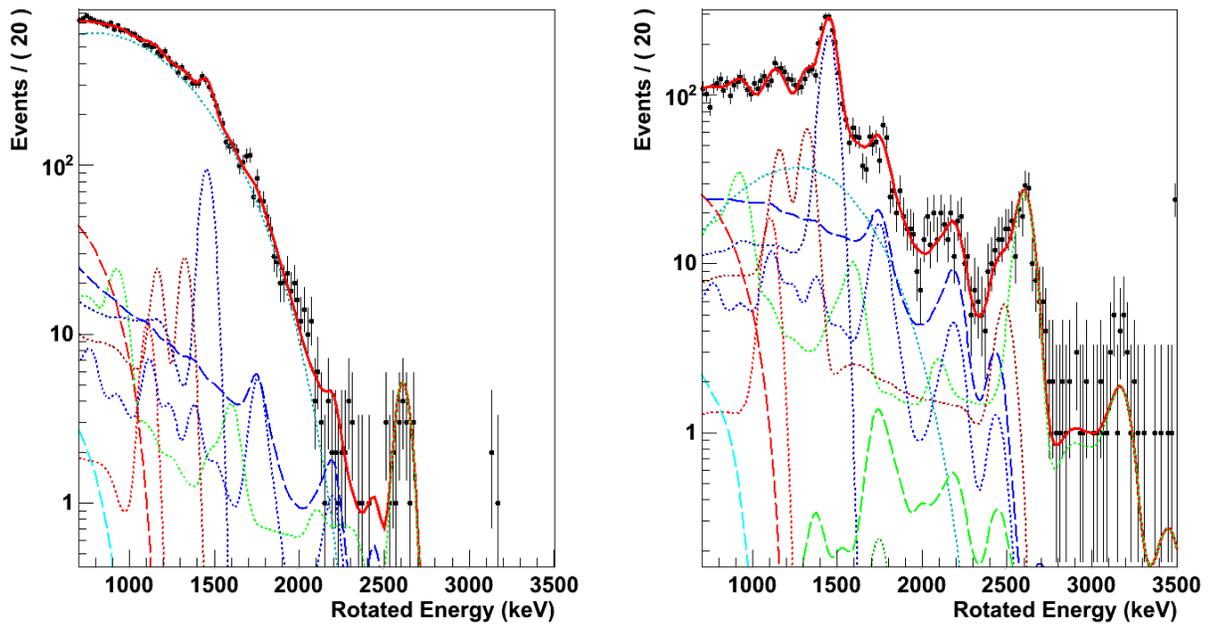
The Fano factor term of the energy resolution is not allowed to float, since it was measured to be vanishing. Only one term of the energy scale parameterization is floated due to difficulties in implementing the convolution. It will be seen that the energy scale is not largely affected by the fit and plays only a minor role in the overall systematic uncertainty.

7.7.1 β energy scale determination

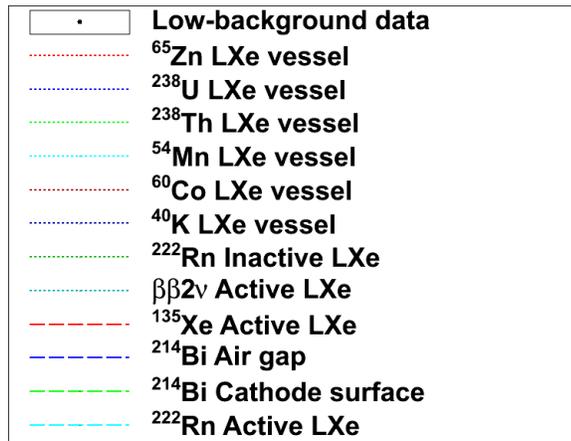
To address the concern about a possible discrepancy from the γ scale (see Section 6.5.1) the best relative β scale was determined using a profile likelihood scan. Since the β scale is not a floating parameter of the fit, a new PDF was created for each β -like model component by applying a fixed, overall multiplicative factor, η to the energy scale, including $\beta\beta$ of both types, as well as ^{137}Xe decay. The γ scale was not modified. An extended maximum likelihood fit of the model to the low-background data was made using the new PDFs, and the procedure was repeated for several values of η . The maximum likelihood was obtained for $\eta = 0.994$, with uncertainty ± 0.004 . This is sufficiently different from unity that we consider this systematic again in Section 7.9.2.

7.7.2 Fit results

A final maximum likelihood fit was made using the value $\eta = 0.994$ for the β energy scale. Figure 7.9 shows the low background spectrum with the fitted PDFs. Peaks from ^{232}Th and ^{40}K stand out strongly at 2615 keV and 1461 keV. In the region of the $\beta\beta 0\nu$ Q-value of 2458 keV, ^{214}Bi in the air gap can be seen to be the background that contributes most significantly. Within $1\text{-}\sigma$ of the Q-value, the fit predicts 1.9 ± 0.2 events. The best fit values for all the floating parameters, including the number of events in each component, are given in Table 7.5. The reduced χ^2 for the single-site fit is 0.668, and for the multi-site fit it is 1.238.



(a)



(b)

Figure 7.9: Low background energy spectra with fitted to the model of signal and background. Left is single-site; right is multi-site.

Table 7.5: Best fit values for all observables of the PDF. Calibration-related parameters are location-independent. Parameters are listed for single-site (SS) and multi-site (MS), except where noted.

| Location | Source | Multiplicity | Number of events | Error |
|---------------------|-------------------------|--------------|------------------|---------|
| Active LXe | ^{222}Rn | MS | 1.73e1 | 2.79 |
| Active LXe | ^{222}Rn | SS | 1.65e1 | 8.24e-3 |
| Active LXe | ^{135}Xe | MS | 2.50e2 | 4.15e1 |
| Active LXe | ^{135}Xe | SS | 3.90e2 | 1.02e2 |
| Air gap | ^{214}Bi | MS | 1.26e3 | 1.38e2 |
| Air gap | ^{214}Bi | SS | 6.04e2 | 7.27 |
| Cathode surface | ^{214}Bi | MS | 3.16e1 | 3.21 |
| Cathode surface | ^{214}Bi | SS | 5.84 | 8.52e-3 |
| Inactive LXe | ^{222}Rn | MS | 5.91 | 7.33e-1 |
| Inactive LXe | ^{222}Rn | SS | 2.72 | 2.07e-3 |
| LXe vessel | ^{60}Co | MS | 9.18e2 | 8.06e1 |
| LXe vessel | ^{60}Co | SS | 4.20e2 | 5.96e1 |
| LXe vessel | ^{40}K | MS | 1.70e3 | 6.53e1 |
| LXe vessel | ^{40}K | SS | 8.40e2 | 6.66e1 |
| LXe vessel | ^{54}Mn | MS | 2.20e-1 | 3.26e-2 |
| LXe vessel | ^{54}Mn | SS | 1.90e-1 | 2.79e-2 |
| LXe vessel | ^{232}Th | MS | 8.22e2 | 5.14e1 |
| LXe vessel | ^{232}Th | SS | 4.11e2 | 4.53e1 |
| LXe vessel | ^{238}U | MS | 4.82e2 | 1.20e2 |
| LXe vessel | ^{238}U | SS | 2.31e2 | 6.05e1 |
| LXe vessel | ^{65}Zn | MS | 1.25e2 | 4.18e1 |
| LXe vessel | ^{65}Zn | SS | 7.86e1 | 2.78e1 |
| Active LXe | $\beta\beta 0\nu$ | N/A | 1.01e-7 | 1.24 |
| Active LXe | $\beta\beta 0\nu$ SS/MS | N/A | 8.28e-1 | 8.44e-2 |
| Active LXe | $\beta\beta 2\nu$ | MS | 1.58e3 | 1.80e2 |
| Active LXe | $\beta\beta 2\nu$ | SS | 2.01e4 | 2.27e2 |
| Parameter | Label | Multiplicity | Best fit value | Error |
| Energy scale offset | b_0 | MS | 7.56 | 1.71 |
| Energy scale offset | b_0 | SS | 7.87 | 2.34 |
| Energy res. (noise) | r_1 | MS | 3.98e1 | 2.29e-1 |
| Energy res. (noise) | r_1 | SS | 3.53e1 | 3.22e-1 |
| Energy res. (drift) | r_2 | MS | 1.05e-2 | 2.00e-4 |
| Energy res. (drift) | r_2 | SS | 9.19e-3 | 3.61e-4 |

7.7.3 Fitting systematics

Several PDFS have a preponderance of events at lower energies, near the threshold of 700 keV, while also contributing near the $\beta\beta$ Q-value. For these PDFS, any sensitivity of the fitted event rate to the chosen threshold will lead to a systematic shift in the $\beta\beta 0\nu$ rate. In particular, this affects ^{214}Bi , ^{238}U , ^{60}Co , with peaks respectively at 2448 keV, 2615 keV, and 2506 keV (sum peak). Additionally, background processes should be more tightly constrained by the multi-site fit, whereas $\beta\beta 0\nu$ should mainly appear in the single-site spectrum. If a discrepancy should exist between the single and multi-site energy scales, single-site events may be misidentified.

We investigate these effects by carrying out the simultaneous maximum likelihood fit with various threshold values. We also carry out fits with single-site and multi-site fits separately by only including the appropriate data set in the likelihood determination. Four threshold values were used: the analysis value of 700 keV; 900 keV; 1100 keV; and 1300 keV. These values were chosen to exclude significant portions of the low-energy spectra. In particular ^{60}Co peaks at 1173 and 1332 keV are partially or completely removed, allowing us to evaluate the effects of these peaks on the overall fitted rate.

The fit results can be seen in Figures 7.10 (simultaneous fits), 7.11 (single-site only), and 7.12 (multi-site only). In all cases, it can be seen that the quality of fit, as measured by the reduced χ^2 , is not strongly affected by the changing threshold. For each threshold value, the simultaneous fit values were compared to the separated

fits. Figure 7.13 shows that for several PDFs of ^{214}Bi and ^{222}Rn , the single-site only fit returns significantly fewer events than the simultaneous fit, with about a 25% effect. Conversely, ^{238}U and ^{60}Co return $\sim 75\%$ higher values in the single-site only fit, although the significance is much less. These effects tend to become slightly more significant as the threshold is increased, as can be seen in Figure 7.14.

7.8 $\beta\beta 0\nu$ half-life limit

The best fit number of $\beta\beta 0\nu$ events is seen in Table 7.5 to be consistent with 0. In total, we can find 4.1 ± 0.3 counts due to background in the $1\text{-}\sigma$ region of interest, and 7.5 ± 0.5 counts in the two σ region of interest of the single-site spectrum. Here $\sigma = 41.5$ keV is the single-site energy resolution at the Q-value of 2458 keV.

We can use the maximum likelihood function to set an upper limit on the number of $\beta\beta 0\nu$ events seen in the detector by taking a profile likelihood. The limit obtained after accounting for systematics (discussed in Section 7.9) is

$$n_{\beta\beta 0\nu} < 2.8(1.1) \quad \text{at } 90\% \text{ C.L. (68\% C.L.)} \quad (7.21)$$

We can then convert this to a limit on the $\beta\beta 0\nu$ half-life, $T_{1/2}$, with the formula

$$T_{1/2} = \ln(2) \frac{N(^{136}\text{Xe})T\epsilon}{n_{\beta\beta 0\nu}}, \quad (7.22)$$

where $N(^{136}\text{Xe})$ is the number of ^{136}Xe atoms in the fiducial volume, $T = 2896.6$ h is the livetime of the experiment, and $\epsilon = 71\%$ is the efficiency for $\beta\beta 0\nu$ detection.

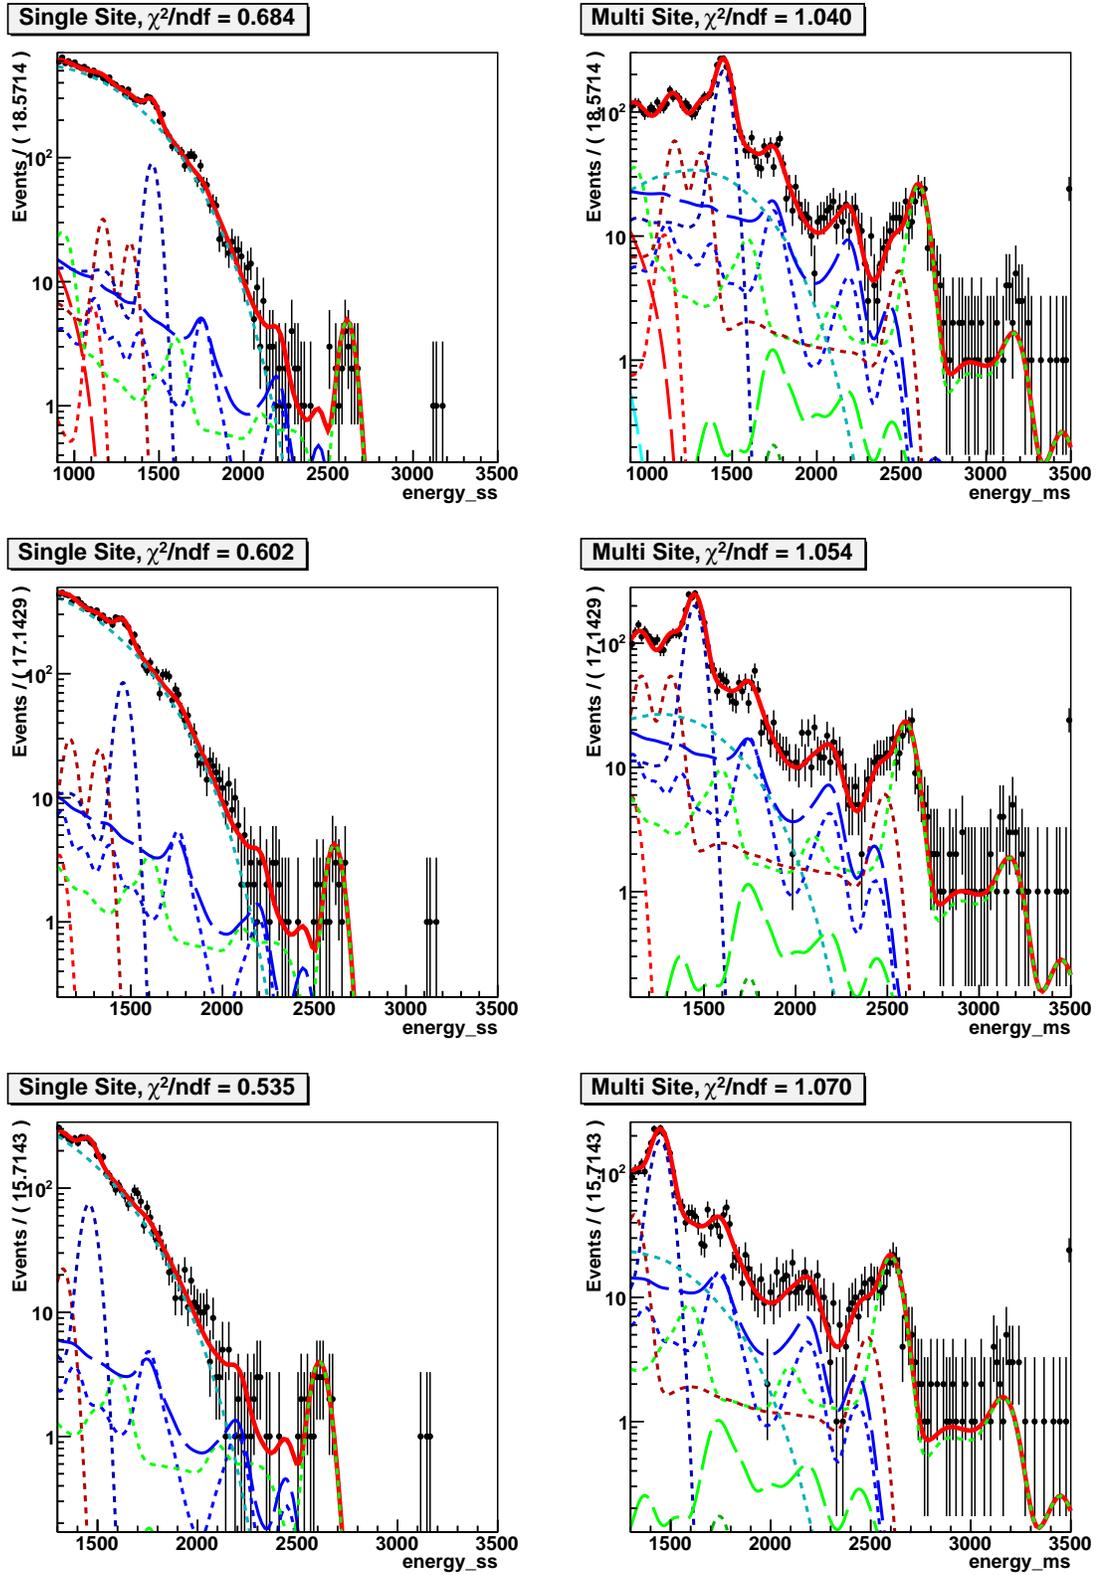


Figure 7.10: Simultaneous single-site and multi-site fits to low background data using non-standard thresholds of 900 (top), 1100 (middle), and 1300 keV (bottom).

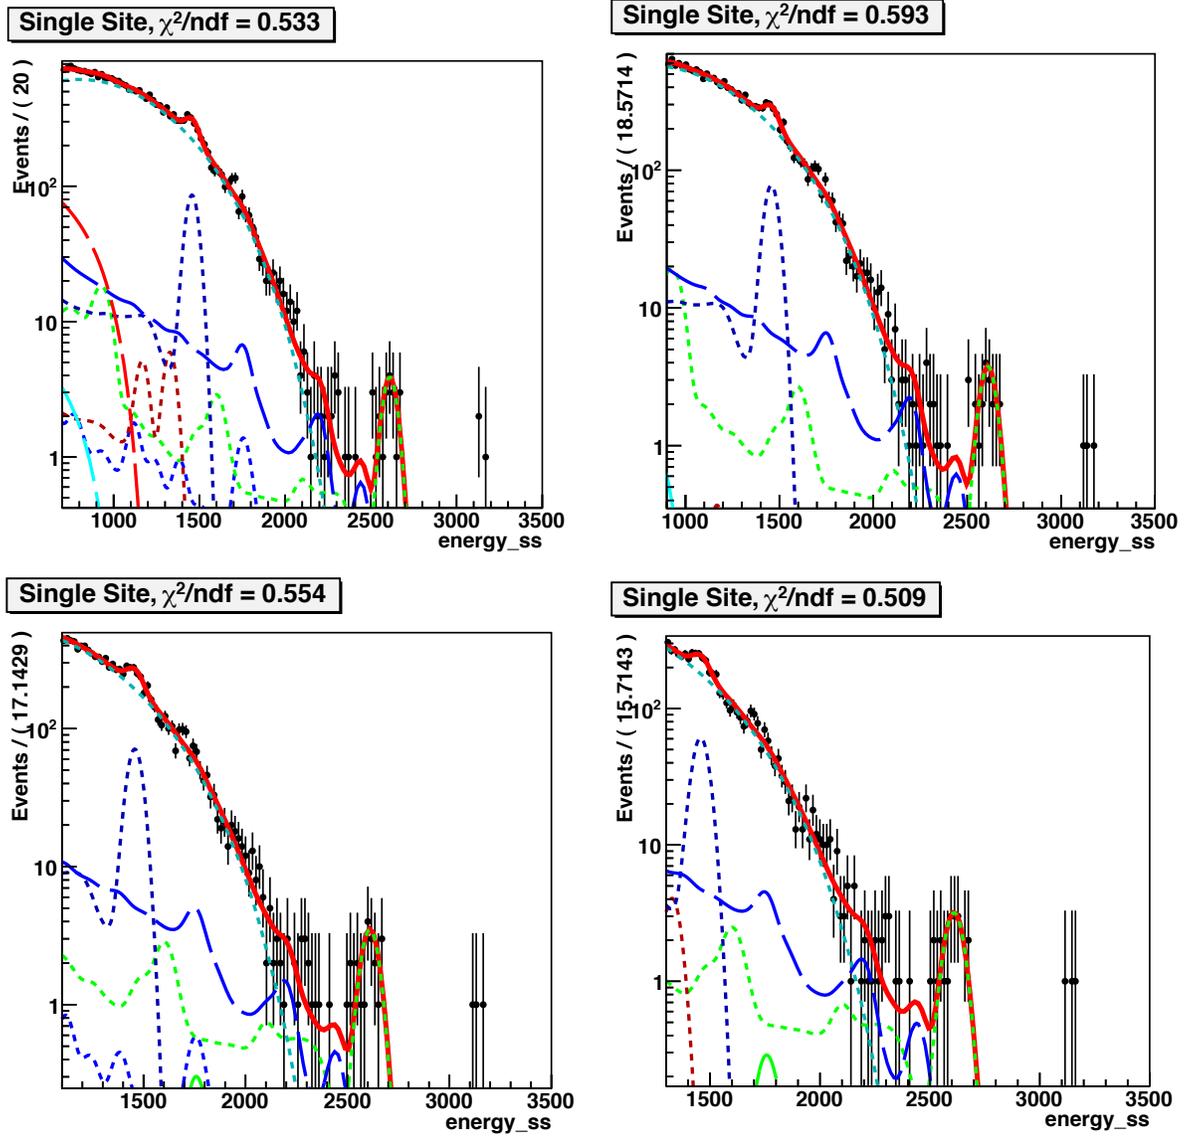


Figure 7.11: Fits to only single-site low background data using thresholds of 700 (top left), 900 (top right), 1100 (bottom left), and 1300 keV (bottom right).

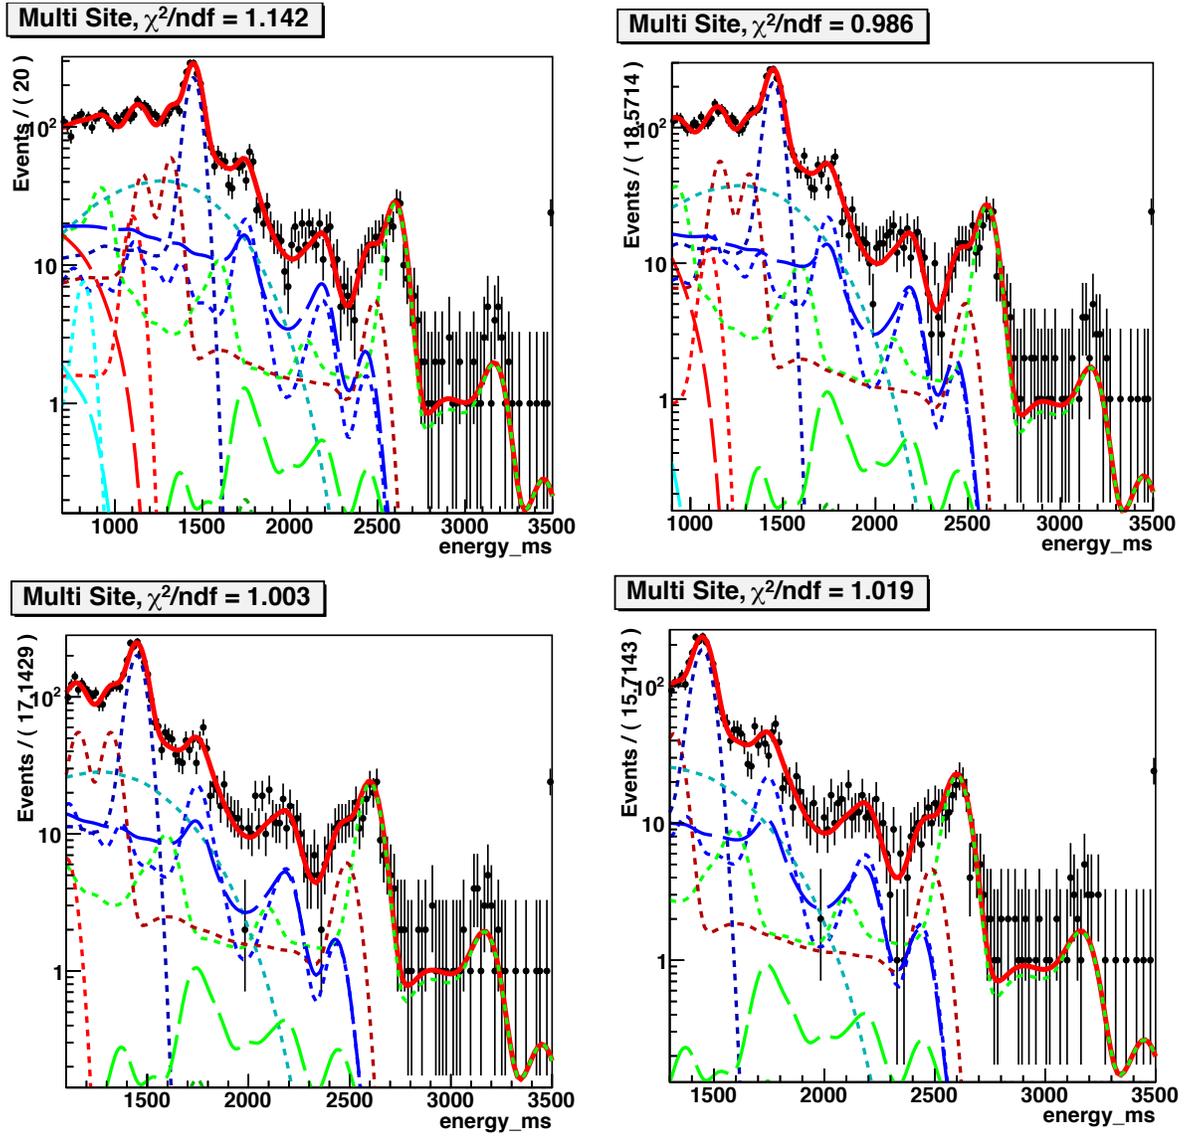


Figure 7.12: Fits to only multi-site low background data using thresholds of 700 (top left), 900 (top right), 1100 (bottom left), and 1300 keV (bottom right).

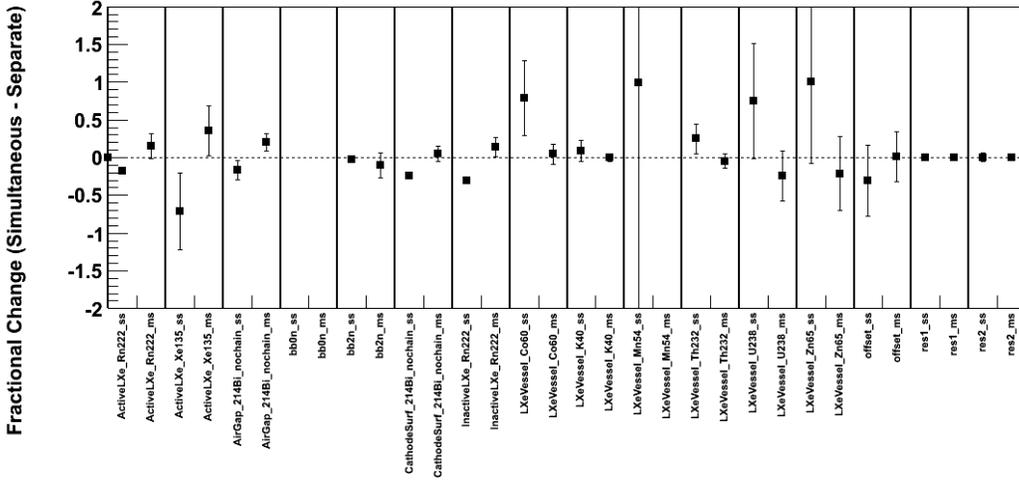


Figure 7.13: Comparison between fitted number events using simultaneous single-site/multi-site fit and individual single-site (SS) or multi-site (MS) fits. The difference is shown as a fraction of the fitted number of events in the simultaneous fit. The default threshold value of 700 keV was used.

From the fiducial mass of ^{136}Xe , $M_{Xe} = 79.4$ kg, we can find

$$N(^{136}\text{Xe}) = \frac{M_{Xe} N_A}{M}, \quad (7.23)$$

where $M = 135.6$ g/mol is the molar mass of the enriched Xe, and N_A is Avogadro's number. Thus $N(^{136}\text{Xe}) = 3.53 \cdot 10^{26}$ atoms. This corresponds to a lower limit on the rate of $1.6 \cdot 10^{25}$ yr at 90% C.L. ($4.6 \cdot 10^{25}$ yr at 68% C.L). Depending on the nuclear model used to calculate the matrix element, we can set an upper limit on the Majorana neutrino mass of 0.14-0.38 eV. The extreme cases are the GCM model [26] and the QRPA mode [28] for 0.14 eV and 0.38 eV, respectively.

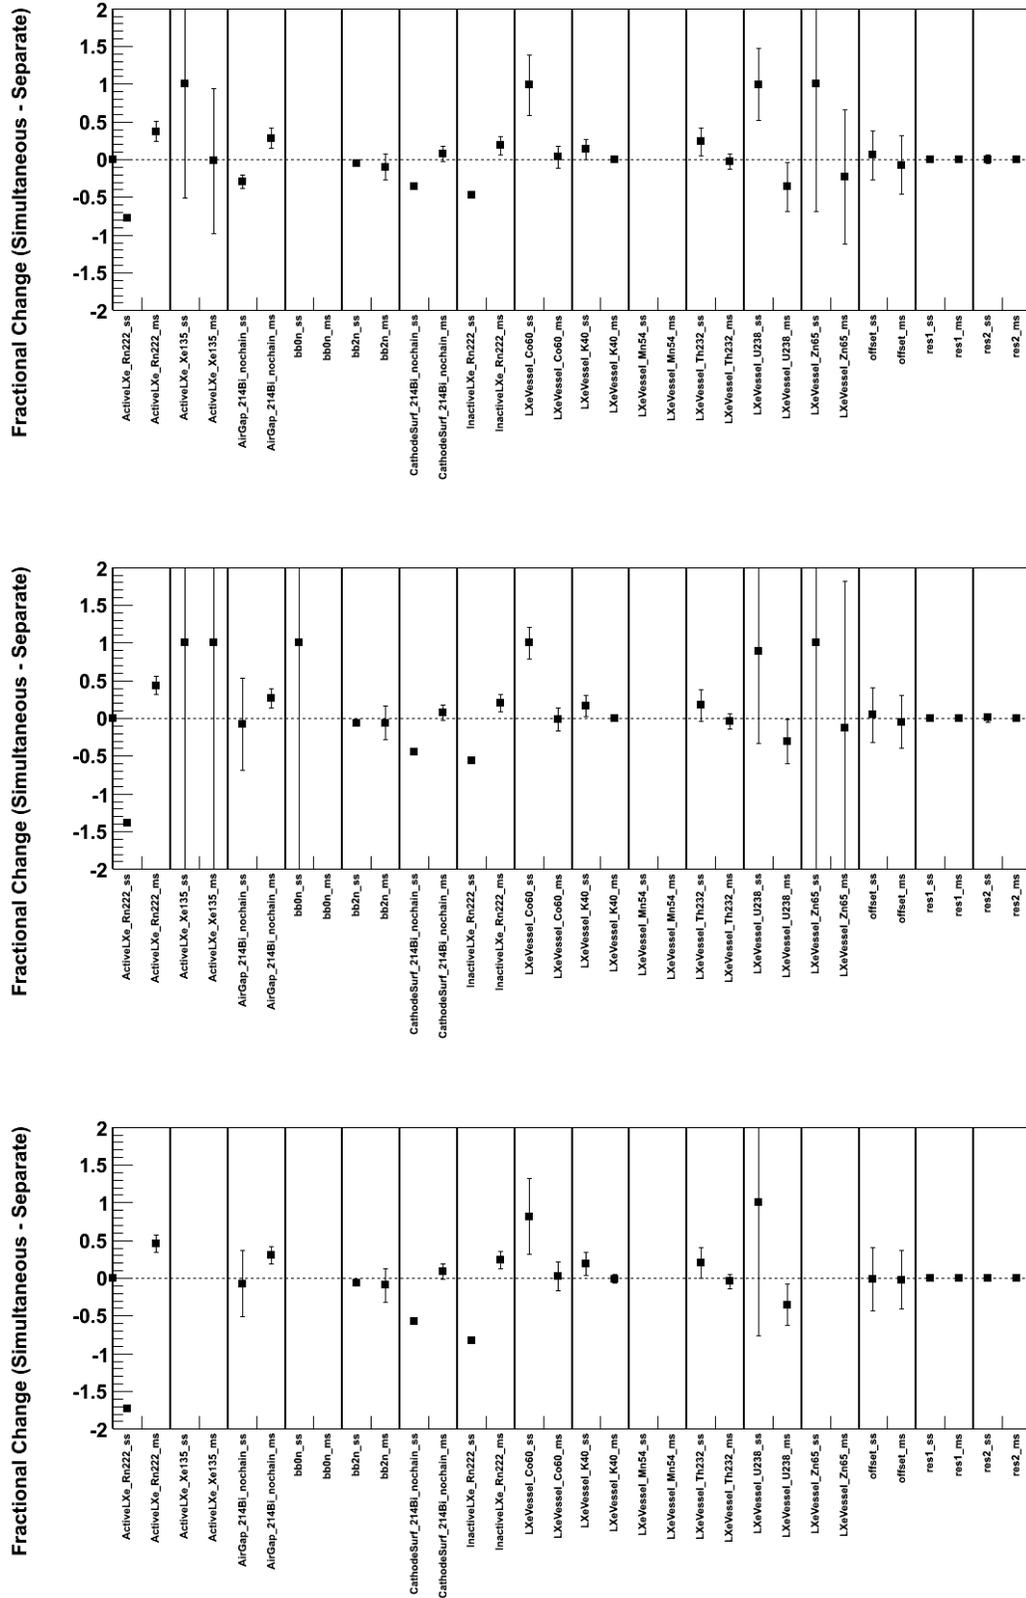


Figure 7.14: Comparison between fitted number events using simultaneous single-site/multi-site fit and individual single-site (SS) or multi-site (MS) fits for threshold values of 900 keV (top), 1100 keV (middle) and 1300 keV (bottom) were used.

7.9 Systematic uncertainties

Most of the assigned systematic uncertainties are found by taking profile likelihood scans over the relevant fit parameters. One exception is the fiducial volume uncertainty, which is estimated from source rate agreements, as discussed in Section 7.9.1 below. We also present an alternative method for estimating the β scale uncertainty, which we conservatively use to increase this uncertainty.

Each systematic reduces the lower limit set on $T_{1/2}$. The contribution is found by determining the lower limit on $T_{1/2}$ half-life limit that would result from not allowing the relevant parameter to float, and taking the fractional increase of that value over the limit with it floating. The results are shown in Table 7.6. The fiducial volume and β scale uncertainties clearly dominate.

Table 7.6: Summary of systematic uncertainties on the $\beta\beta 0\nu$ half-life limit. The associated values are the percentage increase in the lower limit on the $T_{1/2}$; a larger value indicates a weaker limit. The fiducial volume and β scale present the largest uncertainties.

| Systematic | % Increase in $T_{1/2}$ |
|----------------------------------|-------------------------|
| Fiducial Volume | 12.34 |
| β scale | 9.32 |
| Single-site/multi-site ratio | 0.93 |
| Energy scale calibration | 0.04 |
| ^{222}Rn in the air gap | 0.04 |
| Detector energy resolution | 0.00 |

7.9.1 Fiducial volume/source agreement

The fiducial volume in Section 7.2.3 is calculated from the dimensions of a hexagonal solid specified by height and apothem. However, the fiducial determina-

tion is made using reconstructed position, which may systematically differ from the true position. Thus, the region from which fiducial events are drawn may not align with that hexagonal solid. Since the size of the fiducial region determines the mass of decaying xenon, uncertainty on the size of the fiducial volume translates directly into an uncertainty on the $\beta\beta 0\nu$ rate.

To estimate the size of this uncertainty, we use the source rate agreement studies in Section 7.5.3. MC calibration data are used to predict the rate of TPC detections expected for a given source activity and fiducial volume; if the fiducial volume varies from the expected size, the rate of detections will vary in the same way. As described, the source rates do vary from the expected rates; we assign the maximum observed deviation of 9.4% as the uncertainty in the fiducial volume. This results in an decrease of 12.34% in the $\beta\beta 0\nu$ half-life limit.

7.9.2 β scale systematic

The β energy scale was found to differ from the γ scale by a factor $\eta = 0.994 \pm 0.004$, as in Section 7.7.1. While this measurement is nearly consistent with no discrepancy, the best fit value corresponds to a discrepancy of 15 keV at the $\beta\beta$ Q-value, or 35% of the $1\text{-}\sigma$ energy resolution. To better understand this and as an independent check, we constrain the energy scale discrepancy in another way.

We examine the possibility that the different β and γ energy scales are due to differing amounts of ion/electron recombination between the two types of deposits. Recombination controls the scintillation/ionization ratio, so under this scenario,

one type would produce less ionization and more scintillation, and vice versa for the other type. This would then necessarily lead to differing energy scales, since the projection to rotated energy weights the charge and light energy differently. However, it would always be possible to find a rotation angle, θ , such that the two scales aligned again. Thus, we can interpret differing recombination for different event types to mean that different event types should use different rotation angles.

It is helpful conceptually to consider the single-site low-background spectrum as shown in Figure 7.15, with black lines indicating constant ionization to scintillation ratio

$$E_S/E_Q = \text{constant} \quad (\text{black lines}), \quad (7.24)$$

and red lines indicating constant rotated energy

$$E_Q \cos(\theta) + E_S \sin(\theta) = \text{constant} \quad (\text{red lines}). \quad (7.25)$$

Here θ is the rotation angle defined in Equation 6.27. The γ and β -like events appear in a band at the the lower left of the figure, while α decays appear in the lower right. The lines shown span the space, indicating that an event can be parameterized by its light/charge ratio and its rotated energy. However, the rotated energy scale is not defined by this space; we are measuring the rotation angle, and different parts of the space may demand different scales. This leads to significant differences: Figure 7.16 shows lines of constant energy for two arbitrary different rotated energy scales, each with a different rotation angle. It is clear that for events with moderate

ionization/scintillation ratio, the difference in rotated energy between the two scales is small, but for the α decays, the energy may change by 1000 keV or more.

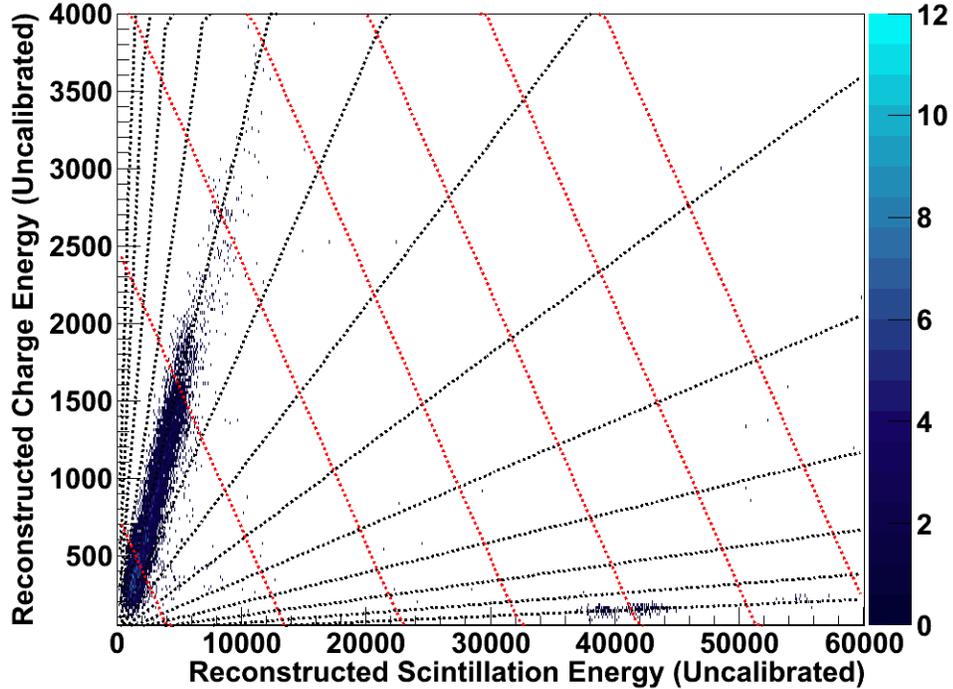


Figure 7.15: Low background ionization-scintillation spectrum showing lines of constant rotated energy (red) and constant ionization/scintillation ratio (black). α decays appear at lower right, with larger ratio than β s or γ s.

Since the effect of changing the energy scale/rotation angle is so much larger for α s, we should be able to determine an appropriate angle for α s and extrapolate the change in rotation angle to the β s and γ s. To do this we also need to estimate the possible difference in scintillation/ionization ratio between β s and γ s. Studies of the centroid distributions of events in low-background data and calibration show that the $\beta\beta 2\nu$ events are about 3% offset from the γ events in the two-dimensional ionization-scintillation plane; see Figure 7.17. Only events below the Q-value are studied, where both data sets have sufficient statistics and are continuous (not dominated by peaks

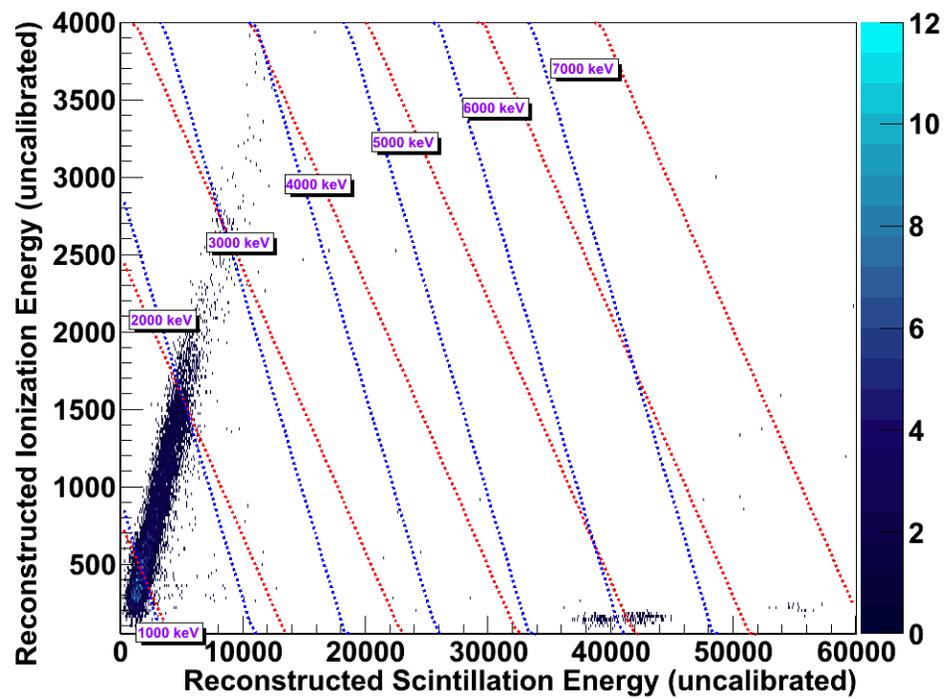


Figure 7.16: Low background ionization-scintillation spectrum showing lines of constant rotated energy corresponding to two different rotation angles. The rotated energy of the α s strongly depends on the rotation angle, while the effect on β s and γ s is much smaller.

of fixed energy). It is not clear whether this offset is in the rotated energy direction, or the scintillation/ionization direction, but we take 3% as a conservative upper limit on the offset in the ratio.

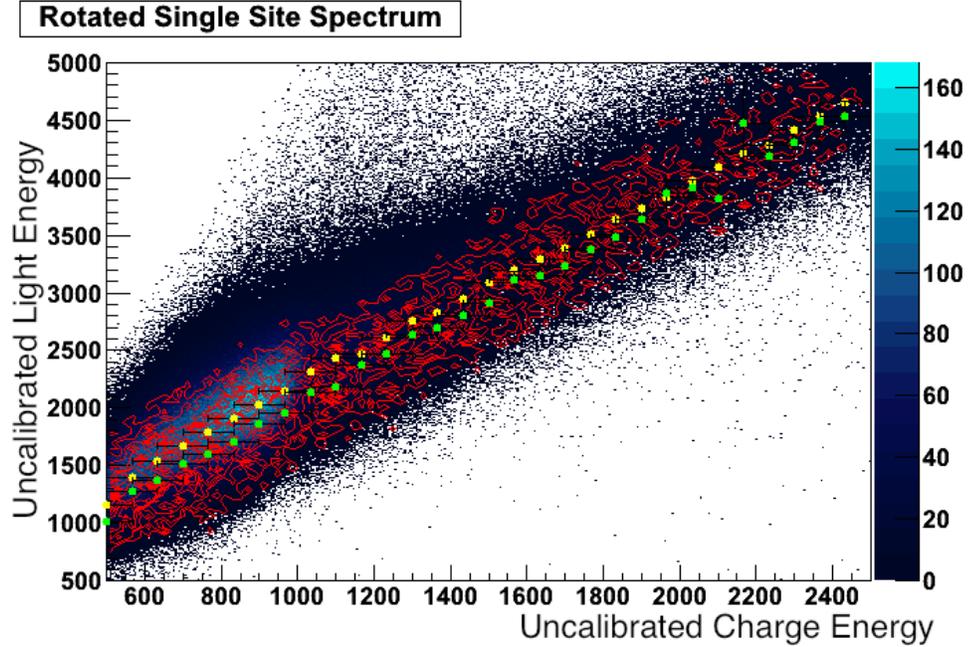


Figure 7.17: Low-background 2-D spectrum (in red contours) superimposed on thorium calibration spectrum. Points indicate fitted mean position after rotation of $\theta = 0.1802$; yellow is thorium, green is low-background. The means do not align, indicating an offset either in scintillation/ionization ratio, or in overall rotated energy.

From the decay chains of the known background sources, we identify three α decay peaks in the low-background spectrum. From the known backgrounds present in detector components, we take these to be three α s in the ^{222}Rn decay chain: ^{222}Rn itself, ^{218}Po , and ^{214}Po . These have well-known energies of at 5.49, 6.00, and 7.69 MeV, respectively. We can make gaussian fits of the α peaks in the rotated low-background spectrum produced by the best-fit rotation angle $\theta_{SS} = 0.18$. This is shown in Figure 7.18, where it is seen that the peaks lie energies well below the known values.

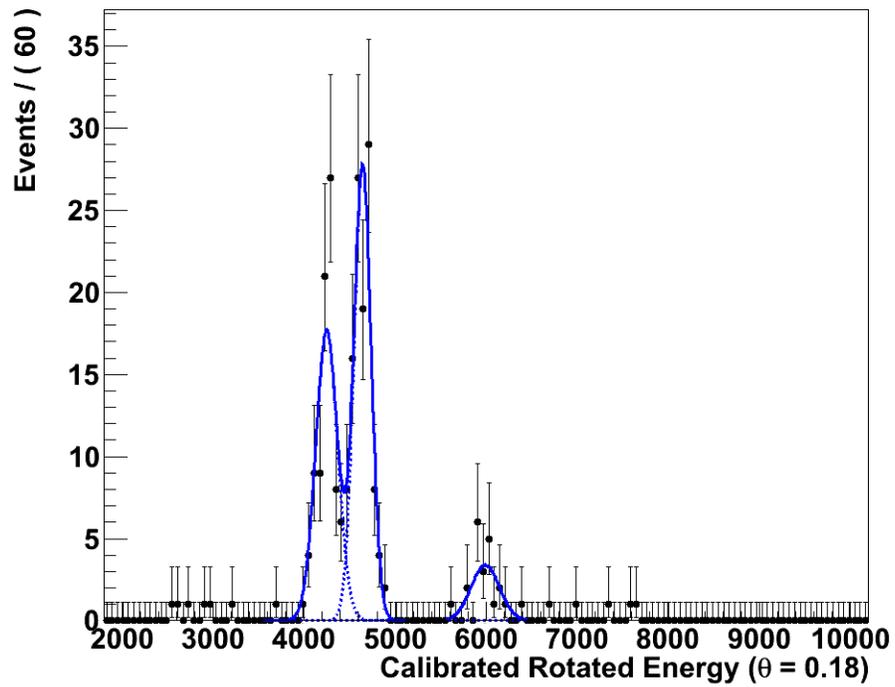


Figure 7.18: Rotated low-background spectrum focused on α decays. The blue line indicates the result of a maximum likelihood fit to the sum of three Gaussians. The best fit values for the peaks are 4.25 ± 0.02 MeV, 4.64 ± 0.01 MeV, and 5.99 ± 0.03 MeV, significantly below the known peak energies.

We can find the optimal rotation angle for calibration the α decays by scanning over various rotation angles. We recalibrate the rotated energy at each angle, and refit the α peaks. The resulting peak positions are shown in Figure 7.19, where we can see the scales align at $\theta_{SS}^\alpha = 0.26$. The fits for this parameter are shown in Figure 7.20. We can then geometrically determine the change in rotated energy of $\beta\beta 0\nu$ decays that would be caused by using $\theta_{SS}^\alpha = 0.26$ rather than $\theta_{SS} = 0.1802$. Combining this with the 3% offset, we find a 0.5% shift at the Q-value of 2458 keV.

Note that using the best fit rotation parameter from the α decays to constrain differences in the β and γ parameter is very conservative. β s and γ s are in fact very similar in scintillation/ionization ratio, and as one would expect the rotation parameter to vary continuously with the ratio, there is no reason to expect as large a difference between β s and γ s as between those and α s. However, in the absence of a fundamental recombination model, and given that a slight energy scale difference is observed, we choose to use this conservative limit. We combine in quadrature the 0.5% uncertainty from the rotation parameter with the 0.4% uncertainty from the β scale profile likelihood, giving a final value of $\eta = 0.994 \pm 0.006$. This systematic uncertainty results in an decrease of the final $\beta\beta 0\nu$ half-life limit of 9.32%.

7.9.3 Systematics with floating fit parameters

The other systematics, which all have parameters floating in the final fit, are sub-dominant to the ones already mentioned. Of these, the single-site/multi-site ratio is the most significant, contributing a 0.93% increase in the final result. We

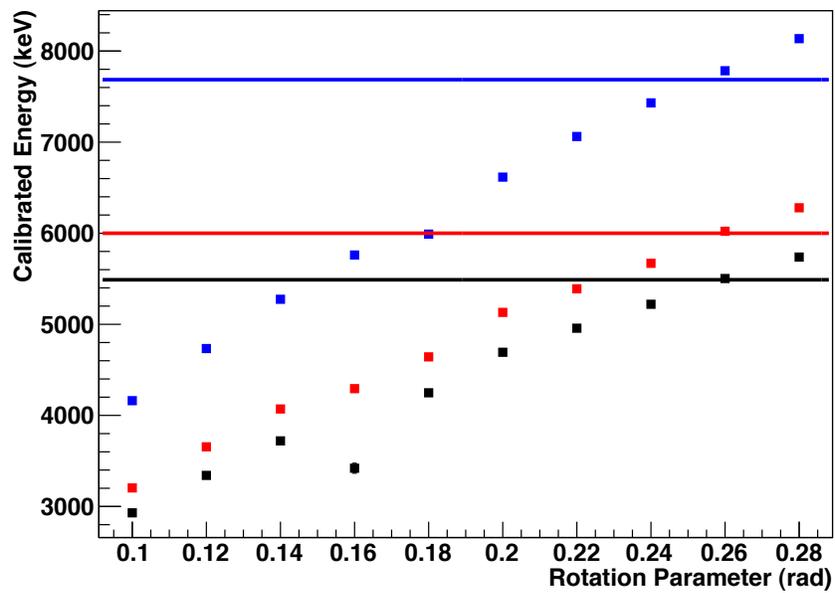


Figure 7.19: Best fit peak position for each α decay at various rotation angles. Black points indicate ^{222}Rn ; red indicate ^{218}Po ; and blue indicate ^{214}Po . The horizontal bars indicate the true decay energies. The α energy scale is calibrated for θ near 0.26.

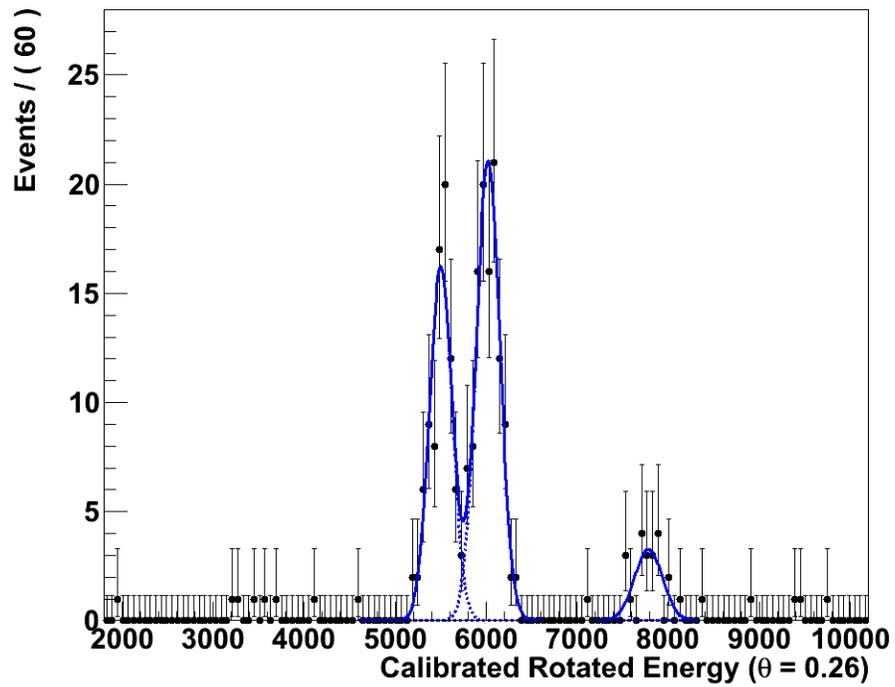


Figure 7.20: Rotated low-background spectrum of α_s using a rotation parameter $\theta = 0.26$. The blue line indicates the result of a maximum likelihood fit to the sum of three Gaussians. The best fit values for the peaks are 5.50 ± 0.02 MeV, 6.02 ± 0.01 MeV, and 7.78 ± 0.03 MeV, nearly consistent with the known peak energies.

should expect this to contribute significantly, since the uncertainty on the rate is substantial. The energy calibration uncertainties contribute negligibly. The systematic associated with the constraint on ^{222}Rn in the air gap is also negligible, although this may be expected since the constraint is only a weak upper limit.

7.10 Comparison with previous results

The claimed detection of $\beta\beta 0\nu$ in [32] reported a half-life of $T_{1/2} = (2.23_{-0.31}^{+0.44}) \cdot 10^{25}$ yr in ^{76}Ge . At 90% confidence level, this is in contradiction with the EXO-200 limit of $1.6 \cdot 10^{25}$ yr. The rates can be compared after taking into account the nuclear matrix elements (see Section 2.3), as shown in Figure 7.21. This shows the predicted half-life of ^{76}Ge and ^{136}Xe using several nuclear models for a range of $\langle m_{\beta\beta} \rangle$. In only in one model (QRPA-2) are the two half-lives consistent at 90% confidence level, and at 68% confidence level, they are not consistent for any model.

Since the data in this work were taken, the Kamland-Zen collaboration has also reported a new limit of $T_{1/2} > 1.9 \cdot 10^{25}$ yr on the $\beta\beta 0\nu$ decay half-life of ^{136}Xe [84]. As the isotopes are the same, the limit can be compared directly, and is a modest improvement on the EXO-200 result. In addition, those authors combine both results to obtain a limit of $T_{1/2} > 3.4 \cdot 10^{25}$ yr.

7.11 Prospects for $\beta\beta 0\nu$ searches

The values of $m_{\beta\beta}$ excluded by the present work fall in the degenerate neutrino hierarchy. (See Figure 2.4.) Sensitivity at the level of the inverted hierarchy requires

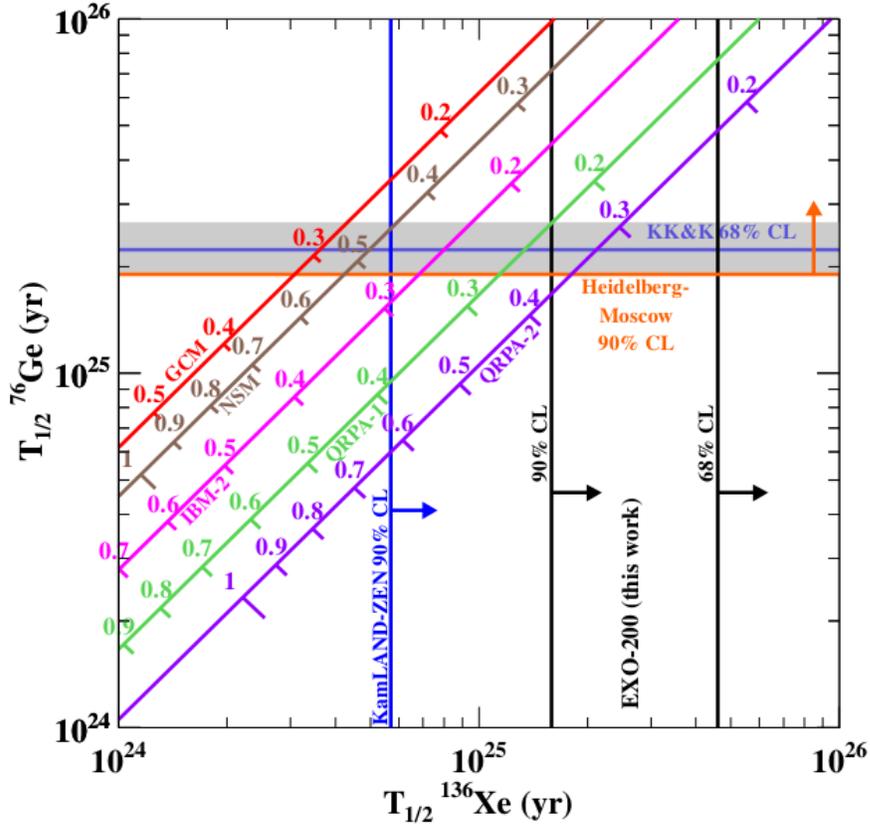


Figure 7.21: Comparison of measured and predicted $\beta\beta\nu$ half-lives for ^{76}Ge and ^{136}Xe . Vertical bands indicate limits for ^{136}Xe from this work (black) and preliminary Kamland-Zen results [41] (blue); for ^{76}Ge , horizontal bands indicate the Heidelberg-Moscow limit [31] (orange) and the claimed observation [32] (dark blue with grey bands indicating 68% confidence level). Diagonal bands show the predicted half-lives using the various nuclear models described in Section 2.3. The scale on these indicates the value of $m_{\beta\beta}$, which can be read off at the intersection of a diagonal band with another band. The allowed parameter space to the right of the limit in this work only includes the claimed observation for one nuclear model (QRPA-2) at 90% confidence, and no model is compatible at 68% confidence. Figure is from [4].

probing $m_{\beta\beta} \sim 50$ meV, a factor of 3 improvement. The sensitivity of an $\beta\beta 0\nu$ search improves with the run time, T , as $T^{1/4}$, too slow to expect EXO-200 or experiments of comparable scale to reach this level with additional running time. Instead, new detectors using larger masses of the active isotope must be created. A range of experimental efforts across different isotopes are under way to prepare such detectors. Many these efforts are taking data with or are planning to use a prototype on the scale of EXO-200, and larger scale detectors can be anticipated to start operating in the next decade. Detectors currently under construction are expected to achieve sensitivities of ~ 10 -50 meV [2], sufficient to investigate the inverted hierarchy (but not the normal hierarchy).

EXO EXO-200 is planned to continue taking data until it has been running for about 4 years. Assuming only statistical improvements, the ultimate mass sensitivity would be 86 - 233 meV. Some improvements to this can be made by reducing systematic uncertainties. In particular, the major uncertainty from ^{222}Rn daughter decays should be mitigated by removing Rn from the ambient cleanroom air; a charcoal filter is being installed to accomplish this.

A future experiment is planned using several tons of enriched Xe in a TPC. The experiment may incorporate “Barium tagging”, a technique in which the daughter of the $\beta\beta$ is identified using laser spectroscopy. This would allow for a zero-background experiment. The project is currently in the research and development phase.

Kamland-Zen The Kamland-Zen experiment employs a 3.08-m-diameter balloon filled with liquid scintillator and loaded with 179 kg (fiducial mass) of enriched ^{136}Xe . This balloon is suspended in the existing KamLAND detector. The experiment began to run in the summer of 2011, and the limits are similar to those of EXO-200, as discussed above. The limit obtained in the most recent run was made worse by a contamination of ^{110m}Ag , a β emitter with a Q-value of 3.01 MeV. The Xe/scintillator was removed from the balloon for removal of the ^{110m}Ag , and an investigation into the source of the contamination is ongoing. The current detector can be upgraded to accommodate more xenon, up to ~ 1 ton.

CUORE CUORE will consist of 988 TeO_2 crystal bolometers, each 5x5x5 cm, arranged into 19 towers. The active mass will be 204 kg of ^{130}Te . A prototype, CUORICINO, was able to achieve a limit of $m_{\beta\beta} < 190\text{-}680$ meV using a single tower [85]. The experiment will be operated in the Gran Sasso Underground Laboratory and is currently under construction.

GERDA GERDA is a ^{76}Ge experiment using bare high purity enriched Ge detectors [86]. These are placed in a cryostat directly immersed in highly purified liquid argon. The bare detectors have reduced background compared to clad detectors, as the cladding can be a source of background decays. Studies were made to show that the detector stability does not suffer in this state. The cryostat is surrounded by purified water, shielding the Ge from external backgrounds. Signal pulse shape analysis can be used to identify backgrounds. Phase I of the experiment, ongoing

since November 2011, uses 18 kg of enriched Ge. An additional 50 kg of enriched GeO₂ has been obtained for Phase II, and a ton scale experiment is also planned.

MAJORANA The MAJORANA collaboration will use high purity Ge diode detectors, enriched to 86% in ⁷⁶Ge. Currently, a 40 kg prototype, the MAJORANA Demonstrator, is under construction at the Sanford Underground Research Facility in Lead, SD. The Demonstrator should show that a background rate of 1 count per ton-year is achievable in a 4 keV window around the $\beta\beta$ Q-value. The collaboration will work with GERDA in developing a ton-scale ⁷⁶Ge experiment.

SNO+ SNO+ is an upgrade to the existing SNO kilo-ton scale liquid scintillator neutrino detector. A 12-m-diameter sphere filled with scintillator will be immersed in the water tank, and the scintillator will be loaded with a $\beta\beta$ decay source. The collaboration has decided to use ¹³⁰Te as the main $\beta\beta$ decay isotope [87]. The large shielding factors of the scintillator and water should yield low background rates. Detector installation is currently underway.

SuperNEMO The SuperNEMO experiment will use the same thin metal source foil technology that was used the NEMO-3 experiment [88]. Approximately 20 foils of approximately 5 kg each will be placed in a tracking chamber that can be used to record the path and energy of decay products. Possible isotopes of study include ⁸²Se, ¹⁵⁰Nd and ⁴⁸Ca. Construction is underway.

Chapter 8

Conclusions

We have presented a search for neutrinoless double beta ($\beta\beta 0\nu$) decay using 2896.6 hours of low-background data taken with $3.53 \cdot 10^{26}$ atoms of ^{136}Xe in the EXO-200 detector. We find no evidence for $\beta\beta 0\nu$ in ^{136}Xe , and set a lower limit on the $\beta\beta 0\nu$ half-life of $T_{1/2} > 1.6 \cdot 10^{25}$ yr at 90% C.L. ($4.6 \cdot 10^{25}$ yr at 68% C.L.). This corresponds to a Majorana neutrino mass upper limit of $\langle m_{\beta\beta} \rangle < 0.14 - 0.38$ eV. This limit probes the degenerate sector of the neutrino mass hierarchy; to probe the inverted mass hierarchy would require sensitivity to ~ 0.05 eV, requiring a larger detector with greater statistics and lower backgrounds.

For most nuclear models, the limit is not compatible with the only reported observation of $\beta\beta 0\nu$, in ^{76}Ge [32]. We improve upon the limit set in [31] of $\langle m_{\beta\beta} \rangle < 0.35$ eV, the leading limit in an isotope other than ^{136}Xe . Our result is also consistent with an independent limit of $1.9 \cdot 10^{25}$ yr set using the same isotope [89]. EXO-200 could achieve better sensitivity by reducing the systematic uncertainties, particularly in the fiducial volume and the energy scale. The largest background, ^{222}Rn in the air near the detector, can also be reduced through improvements in air handling systems.

Bibliography

- [1] S. M. Bilenky and C. Giunti. *Mod. Phys. Lett. A*, 27:30015, 2012.
- [2] A. S. Barabash. *ArXiv e-prints*, September 2012.
- [3] F. T. Avignone, S. R. Elliott, and J. Engel. *Rev. Mod. Phys.*, 80:481–516, Apr 2008.
- [4] M. Auger et al. *Physical Review Letters*, 109(3):032505, July 2012.
- [5] B. Aharmim, et al. *ArXiv e-prints*, September 2011.
- [6] P. Adamson, et al. *Phys. Rev. Lett.*, 106(18):181801, May 2011.
- [7] B. Aharmim, et al. *ArXiv e-prints*, September 2011.
- [8] F. P. An, et al. *Physical Review Letters*, 108(17):171803, April 2012.
- [9] R. Wendell, et al. *Phys. Rev. D*, 81(9):092004, May 2010.
- [10] V. M. Lobashev et al. *Nucl. Phys. Proc. Suppl.*, 91:280–286, 2001.
- [11] C. Kraus, et al. *Eur. Phys. J.*, C40:447–468, 2005.
- [12] J. Wolf. *Nucl. Instrum. Meth. A*, 623(1):442 – 444, 2010.
- [13] X. Wang, et al. *JCAP*, 11:18, November 2012.
- [14] U. Seljak, et al. *JCAP*, 10:14, October 2006.
- [15] M. Viel, et al. *JCAP*, 6:15, June 2010.
- [16] S. Hannestad. *Prog. Part. Nucl. Phys.*, 65(2):185 – 208, 2010.
- [17] K. Abazajian, et al. *Astropart. Phys.*, 35:177–184, November 2011.
- [18] E. Majorana. *Nuovo Cim.*, 14:171–184, 1937.
- [19] R. N. Mohapatra and G. Senjanović. *Phys. Rev. Lett.*, 44:912–915, Apr 1980.
- [20] T. Yanagida. *Prog. Theor. Phys.*, 64(3):1103–1105, 1980.
- [21] LBNL Isotopes Project Nuclear Data Dissemination Home Page. Retrieved March 11, 2002, from <http://ie.lbl.gov/toi.html>. LBNL Isotopes Project Nuclear Data Dissemination Home Page. Retrieved March 11, 2002, from <http://ie.lbl.gov/toi.html>. url.
- [22] N. Ackerman, et al. *Physical Review Letters*, 107(21):212501, November 2011.

- [23] A. S. Barabash. *Phys. Rev. C*, 81:035501, Mar 2010.
- [24] J. Suhonen. and O. Civitarese. *Phys. Rept.*, 300:123–214, 1998.
- [25] P. Vogel. *J. Phys. G. Nucl. Partic.*, 39(12):124002, December 2012.
- [26] T. R. Rodriguez and G. Martinez-Pinedo. *Phys. Rev. Lett.*, 105:252503, Dec 2010.
- [27] Menendez, J., et al. *Nuclear Physics A*, 818(3-4):139 – 151, 2009.
- [28] A. Staudt and K. Muto and H. V. Klapdor-Kleingrothaus. *Europhys. Lett.*, 13(1):31, 1990.
- [29] F. Šimkovic, et al. *Phys. Rev. C*, 79:055501, May 2009.
- [30] J. J. Gomez-Cardenas, et al. *JCAP*, 2011(06):007, 2011.
- [31] H. V. Klapdor-Kleingrothaus, et al. *Eur. Phys. J. A*, 12:147–154, 2001.
- [32] H. V. Klapdor-Kleingrothaus and I. V. Krivosheina. *Mod. Phys. Lett. A*, 21(20):1547–1566, 2006.
- [33] C. E. Aalseth, et al. *Mod. Phys. Lett.*, A17:1475–1478, 2002.
- [34] H. V. Klapdor-Kleingrothaus, et al. *Phys. Lett. B*, 586(3-4):198 – 212, 2004.
- [35] F. Feruglio, A. Strumia, and F. Vissani. *Nucl. Phys. B*, 637:345–377, August 2002.
- [36] S. Umehara, et al. *Phys. Rev. C*, 78(5):058501, November 2008.
- [37] A.S. Barabash and V.B. Brudanin. *Phys. Atom. Nucl.*, 74:312–317, 2011.
- [38] J. Argyriades, et al. *Nucl. Phys. A*, 847(3–4):168 – 179, 2010.
- [39] F. A. Danevich, et al. *Nucl. Phys. B*, 138(0):230 – 232, 2005.
- [40] E. Andreotti, et al. *Astropart. Phys.*, 34:822–831, June 2011.
- [41] A. Gando, et al. (KamLAND-Zen Collaboration). *Phys. Rev. C*, 85(4):045504, April 2012.
- [42] J. Argyriades, et al. *Phys. Rev. C*, 80(3):032501, September 2009.
- [43] S. R. Elliott and P. Vogel. *Ann. Rev. Nucl. Part. Sci.*, 52:115–151, 2002.
- [44] J. Jortner, et al. *J. Chem. Phys.*, 42:4250, 1965.
- [45] E. Aprile and T. Doke. *Rev. Mod. Phys.*, 82:2053–2097, Jul 2010.
- [46] E. Conti, et al. *Phys. Rev. B*, 68(5):054201, August 2003.

- [47] M. Auger, et al. *J. Inst.*, 7(05):P05010, 2012.
- [48] Advanced Photonix part SD630-70-75-500. www.advancedphotonix.com/ap_products/pdfs/SD630-70-75-500.pdf.
- [49] R. Neilson, et al. *Nucl. Instrum. Meth. A*, 608:68, 2009.
- [50] E.I. du Pont de Nemours and Co. Part No. TE-6472. <http://www.dupont.com/>.
- [51] F. Neves, et al. *IEEE T. Nucl. Sci.*, 52(6):2793 –2800, dec. 2005.
- [52] M. Yamashita, et al. *Nucl. Instrum. Meth. A*, 535(3):692 – 698, 2004.
- [53] National InstrumentsTM. <http://www.ni.com/labview/>.
- [54] J. Hodgson. Proposal to run EXO200 with a positive dP offset. *Internal Note*.
- [55] Fluitron model S1-15/70, Fluitron, Inc., 30 Industrial Drive, Ivyland, PA, 18974.
- [56] F. LePort, et al. *Rev. Sci. Instrum.*, 82:105114, 2011.
- [57] A. Dobi, et. al. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 620(2-3):594 – 598, 2010.
- [58] A. Dobi, et al. *Nucl. Instrum. Meth. A*, 659:215–228, December 2011.
- [59] 3M HFE-7000. <http://www.3m.com/>.
- [60] Brooks Automation, Part No. PFC-672HC. <http://www.brooks.com/>.
- [61] Sheldahl, Part No. G405377. <http://www.sheldahl.com/>.
- [62] Tetch Part No. CP-200. <http://www.tetch.com/Cold-Plate-Coolers/CP-200.html>.
- [63] Jetseal ®. <http://www.jetseal.com/umegaseal.php>.
- [64] E. I. Esch et al. *Nucl. Instrum. Meth. A*, 538:516, 2005.
- [65] D.S. Leonard. Simulation of Muon-Induced Backgrounds in EXO Prototype. *Internal Note*.
- [66] D.S. Leonard, et al. *Nucl. Instrum. Meth. A*, 591:490–509, July 2008.
- [67] A. Dobi, et al. *Nucl. Instrum. Meth. A*, 675:40–46, May 2012.
- [68] V. T. Jordanov and G. F. Knoll. *Nucl. Instrum. Meth. A*, 345(2):337 – 345, 1994.

- [69] S. Agostinelli et al. Geant4 - a simulation toolkit. *Nucl. Instrum. Meth. A*, 506(3):250 – 303, 2003.
- [70] A. S. Barabash. *Phys. Atom. Nucl.*, 73:162–178, 2010.
- [71] A.S. Barabash. Double beta decay to the excited states: Experimental review. In *AIP Conference Proceedings*, volume 942, pages 8–12, 2007.
- [72] NIST Photon Cross Sections Database. <http://www.nist.gov/pml/data/xcom/index.cfm>.
- [73] G. Schenter and P. Vogel. *Nucl. Sci. Eng.*, 83:393–396, 1983.
- [74] Cook, et al. Eliminating the Need for Photon Tracking Using Updated Scintillation Parameters for EXO Analysis. *Internal Note*.
- [75] V. Radeka. *Ann. Rev. Nucl. Part. Sci.*, 38(1):217–277, 1988.
- [76] O. Bunemann, et al. *Can. J. Res.*, 27a(5):191–206, 1949.
- [77] G. Bakale, U. Sowada, and W. F. Schmidt. *J. Phys. Chem.*, 80(23):2556–2559, 1976.
- [78] D. Auty, et al. Working Document for Run 2 EXO-200 Analysis.
- [79] O. Hilt and W. F. Schmidt. *J. Phys.-Condens. Mat.*, 6(47):L735, 1994.
- [80] G. Knoll. *Radiation Detection and Measurement*. John Wiley and Sons, Inc., 3 edition, 2000.
- [81] R. Brun and F. Rademakers. ROOT — An object oriented data analysis framework. *Nucl. Instrum. Meth. A*, 389(1–2):81 – 86, 1997. See also <http://root.cern.ch/>.
- [82] W. Verkerke and D. Kirkby. The RooFit toolkit for data modeling. *ArXiv Physics e-prints*, *arXiv:physics/0306116*, June 2003.
- [83] FFTW3. <http://www.fftw.org/>.
- [84] KamLAND-Zen Collaboration. Limit on Neutrinoless $\{\beta\}\{\beta\}$ Decay of Xe-136 from the First Phase of KamLAND-Zen and Comparison with the Positive Claim in Ge-76. *ArXiv e-prints*, November 2012.
- [85] C. Arnaboldi, et al. *Phys. Rev. C*, 78(3):035502, September 2008.
- [86] K.-H. Ackermann, et al. *European Physical Journal C*, 73(3):1–29, 2013.
- [87] <http://snoplus.phy.queensu.ca/News.html>.
- [88] R. Arnold, et al. *European Physical Journal C*, 70:927–943, December 2010.
- [89] A. Gando, et al. *arXiv:1211.3863*, 2012.