



Fakultät für Physik der Technischen Universität München

Excellence Cluster Origin and Structure of the Universe

Search for neutrinoless double beta decay in ^{136}Xe with EXO-200

Wolfhart Feldmeier

Vollständiger Abdruck der von der Fakultät für Physik der Technischen Universität München zur Erlangung des akademischen Grades eines Doktors der Naturwissenschaften (Dr. rer. nat.)
genehmigten Dissertation.

Vorsitzender: Univ.-Prof. Dr. A. Ibarra
Prüfer der Dissertation: 1. Univ.-Prof. Dr. P. Fierlinger
2. Univ.-Prof. Dr. L. Oberauer

Die Dissertation wurde am 8.1.2015 bei der Technischen Universität München eingereicht und durch die Fakultät für Physik am 19.2.2015 angenommen.

Abstract

During the last two decades, great progress has been made in the field of neutrino physics. A large number of experiments have confirmed that neutrinos change their flavour during flight, which is generally accepted as a proof that neutrinos are massive. However the absolute values of neutrino masses, as well as the mechanisms that lead to massive neutrinos, are still unknown. Neutrinoless double beta decay is a hypothetical process that could probe many of the outstanding questions related to neutrino mass. The existence of this process would imply that neutrinos are their own anti-particles, so called Majorana particles, and the rate at which this process occurs allows to draw conclusions on the absolute values of neutrino masses. Furthermore, this process directly violates the conservation of lepton number, which could help in explaining the observed asymmetry between matter and anti-matter in the universe. EXO-200 is searching for neutrinoless double beta decay of ^{136}Xe with an expected sensitivity on its half-life of more than 10^{25} years.

This text describes the techniques by which EXO-200 deduces a limit on the half-life of this decay. The reconstruction of physics events in the detector data is presented, and a detailed description of the reconstructed data used for the analysis of neutrinoless double beta decay is given. An analysis which makes use of 2 years of reconstructed detector data is discussed, which yields a lower limit on the neutrinoless double beta decay of ^{136}Xe of 1.1×10^{25} yr. A metric that aims to give a better quantification of the spatial extent of detector events is introduced and used in a reanalysis of the data for improved background discrimination. Finally, a discussion of the results is given.

Contents

1. Introduction	1
2. Neutrino mass and double beta decay	5
2.1. Massive neutrinos	5
2.2. Neutrino oscillation	8
2.3. Neutrinoless double beta decay	10
3. The EXO-200 detector	17
3.1. Sensitivity considerations	17
3.2. Search for $0\nu\beta\beta$ in liquid ^{136}Xe	18
3.2.1. Radiation detection in liquid xenon	19
3.2.2. Background-free detection of double beta decay	21
3.3. The Time Projection Chamber	22
3.3.1. Charge detection	23
3.3.2. Scintillation detection	25
3.4. Data readout	26
3.5. Detector environment	28
3.5.1. Density of ^{136}Xe in the detector	28
3.5.2. Xenon recirculation and purification	29
3.5.3. Underground facility and muon shield	30
3.5.4. Calibration system	32
3.6. Detector Monte Carlo model	33
4. Event reconstruction	35
4.1. Signal finding	35
4.2. Signal parameter estimation	38
4.3. Signal reconstruction efficiencies	40
4.4. 3D event reconstruction	41
4.4.1. Signal bundling	41
4.4.2. Determination of Z position	42
4.4.3. Determination of X - Y -position	43
4.5. Event topology	43
4.6. Position reconstruction accuracy	44
4.7. 3D Reconstruction efficiencies	46
5. Continuous multiplicity description	49

6. Data processing and correction	57
6.1. Signal corrections	58
6.1.1. Electronegative impurity	58
6.1.2. Shielding grid inefficiency	60
6.1.3. Nonuniform light response	62
6.2. Energy calibration	62
6.2.1. Combination of charge and light	62
6.2.2. Calibration procedure	65
6.2.3. Energy resolution	66
7. Data selection	69
7.1. Detector live-time	69
7.2. Fiducial volume	70
7.3. Event selection criteria	72
8. Search for neutrinoless double beta decay	77
8.1. Maximum likelihood method	77
8.1.1. Binned likelihood fit	78
8.1.2. Parameter constraints and incorporation of efficiency uncertainties	79
8.1.3. Confidence intervals	79
8.2. Modelling of signal and backgrounds using the Monte Carlo simulation .	80
8.2.1. Background model	80
8.2.2. Agreement between simulation and data	83
8.3. Low background data fit	86
8.3.1. Parameter constraints	91
8.3.2. Fit results	92
8.4. MST metric as observable in the fit	96
9. Conclusions and outlook	101
A. Clustering parameters	103
B. MST metric for different diffusion parameters	107

1. Introduction

Although first postulated already over 80 years ago, and experiments as well as theory have made great progress in the last decades, the neutrino still has not revealed all of its secrets. A circumstance that is quite unique among the known elementary particles. The reason for this is that neutrinos interact very weakly with matter, which gave the force by which this interaction is mediated its name. Despite the fact that a huge amount of neutrinos are all around us - about 65 billion solar neutrinos pass through an area as small as a thumbnail every second [1] - only 40 interactions per year are expected to happen from these solar neutrinos within a human body¹.

Before being discovered experimentally, the neutrino was postulated by Wolfgang Pauli in an attempt to rescue the fundamental principle of conservation of energy and angular momentum. Unlike the discrete lines that were seen in alpha and gamma decays of radioactive Isotopes, the spectra of electrons in beta decay were discovered to be continuous. Pauli tried to explain that fact by introducing a new electrically neutral particle that was emitted together with the electron.

More than 20 years later in 1956, Cowan and Reines first detected the electron antineutrino via the use of beta capture [3]. They placed two water tanks with dissolved cadmium chloride next to a nuclear reactor and detected the scintillation light of both, the prompt annihilation gamma, and the coincident gamma emitted after neutron capture on ^{108}Cd . In 1962, Lederman, Schwartz and Steinberger showed that more than one flavor of neutrinos exists, by detecting the creation of charged particles in a muon neutrino beam [4]. After the tau lepton was discovered, it was generally believed that there would exist also a third flavor of neutrino, the corresponding tau neutrino. It was detected in 2000 by the DONUT collaboration [5].

In the 1970s, the Homestake Experiment [6] and several subsequent experiments detected a deficit in the rate of detected electron neutrinos compared to what was expected if the Standard Solar Model was correct. Some 20 years later, experiments like Super-Kamiokande [7] and the Sudbury Neutrino Observatory [8] showed, that neutrinos change flavor while travelling through space and therefore became invisible to those experiments that were only sensitive to electron neutrinos. This fact is known as Neutrino Oscillation and does exist because neutrinos have mass.

As of today the oscillation parameters are known to good precision and it is established that neutrinos have a small but non-zero mass. The absolute differences in mass between the three neutrino flavors are known, however some questions are still unanswered. In

¹ Solar neutrinos have low energy and thus their (measurable) interaction is dominated by neutral current scattering off electrons, whose cross section σ is approximately $1 \times 10^{-45} \text{ cm}^2$ [2]. An average human body weighs 70 kg and thus contains approximately $N = 2 \cdot 10^{28}$ electrons. With the solar neutrino flux density of $\Phi = 6.5 \times 10^{10} \text{ cm}^{-2}\text{s}^{-1}$, this gives $\Phi \cdot N \cdot \sigma \approx 40 \text{ yr}^{-1}$.

1. Introduction

the Standard Model of particle physics only vanishing neutrino masses seem natural and the question remains why neutrino masses are non-zero but small. While the absolute value of the mass differences is known, their sign as well as their absolute scale is not. Besides the fact that the mass of neutrinos is not known, also the origin of this mass is a mystery. Although neutrino masses could be added to the Standard Model in the same way as the masses of other fermions by introducing right-handed neutrinos, this seems somewhat implausible considering the fact that neutrino masses are so much smaller than those of the charged leptons and quarks (see figure 2.1), and it is believed that neutrinos acquire mass via a different mechanism than the other fermions. There exist appealing models, which explain the smallness of neutrino masses by introduction of very heavy right handed neutrinos, so called seesaw mechanisms. These mechanisms however require the neutrino to be its own anti-particle, a so called Majorana particle, a property that the neutrino might possibly have because it does not carry any charge.

Among other types of experiments that directly search for the mass of neutrinos, such as KATRIN [9] or cosmological surveys such as Planck [10], experiments that are searching for the hypothetical neutrinoless double beta decay can shed new light upon many of the open questions. The double beta decay of a nucleus without the emission of (anti-)neutrinos obviously violates the total lepton number by two units and requires the neutrino to be a Majorana particle. The existence of such a process would be a clear sign for new physics not described by the standard model of elementary particles. Furthermore, conclusions on the neutrino mass can be drawn from the rate at which this process occurs.

The chance to better understand the world on its smallest scales and to find out more about the properties of nature's most fundamental ingredients, the elementary particles, is motivation enough to search for the existence of this process. However the properties of this tiny particle have also important consequences on the largest imaginable length scales. Massive neutrinos might contribute significantly to the total energy density of the universe and therefore play an important role in its temporal evolution and formation of its large scale structures. Knowing the size of the neutrino masses therefore allows to better constrain other parameters in cosmological models, and hence, in summary, might result in a better understanding of the universe. The inherent lepton number violating nature of Majorana neutrinos, which could manifest itself in neutrinoless double beta decay, could also contribute to the solution of one of the most important unsolved problems in today's understanding of the universe: The apparent prevalence of matter over antimatter, which can not be explained by the current standard model of particle physics. Fukugita and Yanagida however showed in 1986, that the observed asymmetry of baryons in the universe could as well be generated by a lepton number violating process induced by the possible Majorana nature of neutrinos. The lepton number (L) excess created in this way could efficiently be transformed into a baryon number (B) excess via a $B - L$ conserving standard model process [11]. The observation of neutrinoless double beta decay hence would mark an important discovery with implications on a broad range of questions in fundamental physics.

The EXO programme is searching for this process in the isotope ^{136}Xe . Its current stage, EXO-200 is using 200 kg of enriched xenon in liquid phase and has made the first

observation of the two neutrino decay mode of ^{136}Xe [12, 13] and has placed a limit on the half-life of the neutrinoless mode [14, 15] of 1.1×10^{25} yr, a time almost 10^{15} times as long as the age of the universe. The hunt for such a rare process requires great control over radioactive backgrounds, which might otherwise completely overwhelm the tiny signal. Consequently, EXO-200 is built from ultra low radioactive materials and its design as a time projection chamber allows to better discriminate the residual backgrounds from the signal by their event topology.

This dissertation describes the analysis of data taken between September 2011 and September 2013. Much work has been put into the reconstruction of events from raw data and the restoration of the positions at which energy was deposited in the detector. After a theoretical motivation for the search for neutrinoless double beta decay in chapter 2 and a description of the experimental setup in chapter 3, this will be detailed in chapter 4. The processing of the reconstructed data and corrections that are applied to compensate detector nonlinearities are discussed in chapter 6 and the energy calibration procedure is outlined. Chapter 7 presents the data selection criteria, and the efficiencies and associated systematic errors are discussed for each cut that is applied to the data set. In chapter 8, the background model is presented and the limit on neutrinoless double beta decay of ^{136}Xe is extracted in a fit of the model to the data. Based on the work done for the position reconstruction, a finer graded metric for the spatial extent of detector events has been developed. This is presented in chapter 5, and a first analysis which makes use of this metric for improved background discrimination is shown in section 8.4. Finally, a conclusion and outlook is given in chapter 9.

2. Neutrino mass and double beta decay

In this chapter, the nature of massive neutrinos is discussed. First, the theoretical description as well as subtle differences between massive and mass-less as well as charged and neutral fermions are presented, followed by a discussion of the flavor mixing of massive neutrinos. Finally, the process of nuclear double beta decay is introduced, and the connection between the decay rate of the neutrinoless mode and neutrino masses is shown. This chapter is based largely on standard textbooks on the subject [16–18] and some detailed review papers (e.g. [19–22]).

2.1. Massive neutrinos

Particles with half-integer spin are described by the Dirac equation

$$(i\gamma^\mu\partial_\mu - m)\psi = 0, \quad (2.1)$$

where the particle is associated with the field ψ and the gamma matrices follow the anti-commutation relations

$$[\gamma^\mu, \gamma^\nu]_+ = 2g^{\mu\nu}. \quad (2.2)$$

The solution to this equation is generally an object with four independent complex components. Dirac had electrons in mind when he created this equation, and indeed for electrons and other charged massive fermions all four components are required for a correct description. At the time the Standard Model was built, neutrinos were considered mass-less. For mass-less fermions, the Dirac equation simplifies to

$$i\gamma^\mu\partial_\mu\chi = 0. \quad (2.3)$$

This is called the Weyl equation, whose solutions χ are called Weyl-fields and can be described by only two independent complex components (or 4 real components). A Weyl-field is a state with definite chirality, i.e. it is either left-handed or right-handed, but not a combination of both, in contrast to massive fermion fields which can be described as the linear combination of a left-chiral state and a right-chiral state. In 1957 Goldhaber et al. experimentally verified that only left-handed neutrinos take part in the weak interaction [23]. Subsequently, the weak interaction Lagrangian contains only left-handed neutrinos

$$\mathcal{L} = -\frac{g}{\sqrt{2}} \sum_l \bar{\nu}_{lL} \gamma^\mu l_L^- W_\mu^+ + h.c. - \frac{g}{2\cos\theta_W} \sum_l \bar{\nu}_{lL} \gamma^\mu \nu_{lL} Z_\mu^0, \quad (2.4)$$

2. Neutrino mass and double beta decay

and right-handed neutrinos are not part of the Standard Model. Therefore neutrinos in the Standard Model are left-chiral Weyl-fields.

However as a deficit in the number of electron neutrinos originating from the sun was observed [6], the question arose whether neutrinos might change flavor during flight, a property that would arise naturally if neutrinos have mass. And indeed over the past decade, several experiments asserted beyond doubt, that neutrinos change flavor during flight [7, 8, 24] and therefore have a small but nonzero mass. Consequently the description of neutrinos as Weyl-fermions must be abandoned. This changes the way how neutrinos can be described at a fundamental level and simultaneously gives birth to new questions.

Since neutrinos are massive, they cannot be described by equation 2.3. Instead the full Dirac equation must be used. However, for fields that carry no charge such as the neutrino, one can impose the following requirement

$$\psi = e^{i\alpha} \hat{\psi} \quad (2.5)$$

without violating charge conservation. Here

$$\hat{\psi} := \gamma_0 C \psi^* \quad (2.6)$$

is the charge conjugate field. Equation 2.5 says that the field and its charge conjugate are equivalent up to a complex phase, or in other words, that the particle is its own anti-particle. This is called the Majorana condition and effectively reduces the number of independent complex components of such fields to two. So the question that needs to be asked is whether the neutrino is a four component field which is described by the Dirac equation or whether it is a two component field, a so called Majorana field, which is subject to the additional requirement 2.5 and therefore is its own anti-particle. To date experiments searching for neutrinoless double beta decay are the only experiments that could answer this question. If the process exists, neutrinos must be Majorana particles.

Another question that comes up naturally is, why neutrino masses are so small. Direct mass searches using the beta decay of Tritium currently set an upper limit of $2.05 \text{ eV}/c^2$ (95 % CL) on the mass of the electron neutrino [25, 26]. Analyses of the cosmic microwave background data obtained from the Planck satellite give a somewhat model dependent upper limit on the summed neutrino mass of $0.23 \text{ eV}/c^2$ [10]. This means that neutrino masses are extremely small compared to the masses of all other fermions (see figure 2.1).

In the Standard Model, fermion masses are generated by coupling the fermion fields to the vacuum expectation value of the Higgs field after the symmetry is spontaneously broken:

$$- \mathcal{L}_{mass} = g \phi_0 \bar{\psi} \psi, \quad (2.7)$$

with the (constant) vacuum expectation value of the Higgs field ϕ_0 and a Yukawa coupling constant g which is to be determined experimentally. Neutrino masses could be added to the Standard Model in exactly the same way by adding right handed Dirac neutrinos to the model and choosing g adequately. However, to produce such small neutrino




Figure 2.1.: Masses of known fermions. The upper limit on neutrino masses is $\approx 1 \text{ eV}/c^2$, whereas the mass of the other fermions is at least 6 orders of magnitude higher. While the masses of the charged fermions also span across 6 orders of magnitudes, the mass range within one family is only about 1-2 orders of magnitude. Neutrinos however do not seem to fit in the scene.

masses, the coupling to the Higgs field must be much less than for other fermions. This seems unnatural and other theories that explain this discrepancy are generally preferred.

All fermion masses in the Standard Model are of the form 2.7 or

$$m_D \bar{\psi}_L \psi_R + h.c., \quad (2.8)$$

with the Dirac mass m_D . For Majorana fermions, one can add mass terms of the form

$$m_M \bar{\chi}_L \hat{\chi}_R + h.c., \quad (2.9)$$

with the Majorana mass m_M , which couples the field χ to its \mathcal{CPT} conjugate $\hat{\chi}$. Such terms do not conserve lepton number. However this conservation law is empirical, i.e. there is no fundamental symmetry that would require the lepton number to be conserved. Neutrinoless double beta decay obviously would violate conservation of the lepton number by two units. If one chooses the Majorana fields ν_L , ν_R and their \mathcal{CPT} conjugates as a basis, the most general mass term for one generation of neutrinos can be written in the form

$$-\mathcal{L}_{mass} = \frac{1}{2} \begin{pmatrix} \bar{\nu}_L & \bar{\tilde{\nu}}_L \end{pmatrix} \begin{pmatrix} m_M^L & m_D \\ m_D & m_M^R \end{pmatrix} \begin{pmatrix} \hat{\nu}_R \\ \nu_R \end{pmatrix}, \quad (2.10)$$

with Dirac mass m_D and Majorana masses m_M^L and m_M^R . If m_M^i are set to zero, this equation just reproduces 2.8 and the resulting neutrino is of Dirac type. In the general case, where m_D and m_M^i are non-zero, one can diagonalize the matrix in 2.10 and one

2. Neutrino mass and double beta decay

obtains two eigenstates which are of Majorana type, i.e. they obey 2.5. One can now assume that m_D is generated in the usual way via coupling of the neutrino field to the Higgs vacuum expectation value with coupling strength g on the *same* order as for the other fermions. If one further assumes $m_M^L = 0$ and m_M^R lies at a very high mass scale, i.e.

$$m_M^R \gg m_D, \quad (2.11)$$

then the matrix in 2.10 has the eigenvalues

$$m_1 \simeq \frac{m_D^2}{m_M^R}, \quad m_2 \simeq m_M^R. \quad (2.12)$$

This shows that by introducing a heavy Majorana neutrino, the mass of the other neutrino is scaled by a factor of m_D/m_M^R compared to the mass that is generated via coupling to the Higgs field. This mechanism is called Type I Seesaw [27]. Assuming that m_D is on the order of the τ lepton, then, to generate neutrino masses less than 1 eV, the scale for the heavy neutrino mass is

$$m_M^R \gtrsim 5 \times 10^9 \text{ GeV} \quad (2.13)$$

in this model. Various other Seesaw mechanisms are proposed, e.g. [28, 29]. They all have in common that they generate a small neutrino mass by introducing one or more new, heavier neutrinos.

2.2. Neutrino oscillation

Generally the neutrino states that appear in the weak interaction part of the Lagrangian 2.4 need not be the same as the states appearing in the mass term of the Lagrangian. And indeed the fact that neutrinos oscillate between flavor states shows that this is the case. The most general relation between weak flavor states $|\nu_l\rangle$ ($l = e, \mu, \tau$) and mass eigenstates $|\nu_\alpha\rangle$ ($\alpha = 1, 2, 3$) is

$$|\nu_l\rangle = \sum_i U_{li} |\nu_i\rangle, \quad (2.14)$$

where U is a unitary matrix, the so called Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix. It can be parameterized by three Euler rotation angles, one \mathcal{CP} violating phase that appears even for Dirac neutrinos and - in case neutrinos are Majorana fermions - two additional \mathcal{CP} violating phases:

$$\begin{aligned} U &= \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \\ &= \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \begin{pmatrix} e^{i\alpha_1/2} & 0 & 0 \\ 0 & e^{i\alpha_2/2} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \end{aligned} \quad (2.15)$$




Figure 2.2.: Ordering of neutrino masses. Left: Normal hierarchy where $m_1 < m_2 < m_3$. Right: Inverted hierarchy where $m_3 < m_1 < m_2$.

where s_{ij}, c_{ij} denote $\sin \theta_{ij}$ and $\cos \theta_{ij}$ respectively. θ_{12}, θ_{13} and θ_{23} are the three Euler angles, δ is the Dirac \mathcal{CP} violating phase and α_1 and α_2 are the two Majorana \mathcal{CP} violating phases. Since neutrinos are produced via the weak interaction 2.4, their initial state is a weak flavor eigenstate $|\nu_l\rangle$. The propagation through vacuum however is governed by the mass term of the Lagrangian, where the mass eigenstates $|\nu_\alpha\rangle$ appear and therefore:

$$|\nu_l(t)\rangle = \sum_i e^{-iE_\alpha t} U_{l\alpha} |\nu_\alpha\rangle. \quad (2.16)$$

Detection of the neutrinos then again happens via weak processes and so the amplitude for detecting a $\nu_{l'}$ from a source that creates ν_l is

$$\langle \nu_{l'} | \nu_l(t) \rangle = \sum_\alpha e^{-iE_\alpha t} U_{l\alpha} U_{l'\alpha}^*. \quad (2.17)$$

Since neutrino masses are tiny compared to typical energies, it is justified to work in the ultra-relativistic limit, in which the probability for finding a neutrino which was created

2. Neutrino mass and double beta decay

as ν_l in the state $\nu_{l'}$ after it traveled a distance $x = ct$ becomes

$$\begin{aligned}
P_{l \rightarrow l'} &= |\langle \nu_{l'} | \nu_l(t) \rangle|^2 \\
&= \delta_{ll'} - 4 \sum_{\alpha > \beta} \Re (U_{l\alpha}^* U_{l'\alpha} U_{l\beta} U_{l'\beta}^*) \sin^2 \left(\frac{\Delta m_{\alpha\beta}^2 x}{4E} \right) \\
&\quad + 2 \sum_{\alpha > \beta} \Im (U_{l\alpha}^* U_{l'\alpha} U_{l\beta} U_{l'\beta}^*) \sin \left(\frac{\Delta m_{\alpha\beta}^2 x}{2E} \right). \quad (2.18)
\end{aligned}$$

Without \mathcal{CP} violation, U is real and the second term in the last equation is zero. From 2.18 it is apparent that neutrino oscillation experiments are sensitive to the three

Parameter		Best fit	1 σ Range
Δm_{21}^2 [10^{-5} eV ²]		7.62	7.43 - 7.81
Δm_{31}^2 [10^{-3} eV ²]	normal	2.55	2.46 - 2.61
	inverted	2.43	2.37 - 2.50
$\sin^2 \theta_{12}$		0.320	0.303 - 0.336
$\sin^2 \theta_{23}$	normal	0.613	0.573 - 0.635 and 0.400 - 0.461
	inverted	0.600	0.569 - 0.626
$\sin^2 \theta_{13}$	normal	0.0246	0.0218 - 0.0275
	inverted	0.0250	0.0223 - 0.0276
δ	normal	0.80 π	
	inverted	-0.03 π	0 - 2 π

Table 2.1.: Neutrino mixing parameters obtained from global fit to oscillation experiments [30].

mixing angles θ_{ij} and the differences between the squares of the masses Δm_{ij} . Oscillation experiments have provided a great amount of information on neutrinos, most importantly they have proven that neutrinos have mass. Table 2.1 summarizes the parameters obtained from a global fit over oscillation parameters. Nevertheless, important questions remain that cannot be answered by neutrino oscillation experiments. The \mathcal{CP} violating Majorana phases enter U as a diagonal factor, which cancels in $P_{l \rightarrow l'}$. Therefore, oscillation experiments cannot provide any information about the possible Majorana nature of neutrinos. Since these experiments only measure the absolute differences of the squared neutrino masses, also the question whether neutrino masses are ordered in the so called Normal Hierarchy, where $m_1 < m_2 < m_3$ or in the Inverted Hierarchy, where $m_3 < m_1 < m_2$ (see figure 2.2) needs to be answered by different kinds of experiments.

2.3. Neutrinoless double beta decay

Neutrinoless double beta decay is the best (if not only) candidate process that could answer the remaining questions. Double beta decay, under the emission of two electron




Figure 2.3.: Beta decay scheme for isobars with mass $A = 136$. For ^{136}Xe and ^{136}Ce , single beta decay is energetically forbidden. They can however decay to ^{136}Ba via double beta decay.

anti-neutrinos, is a standard model process, that is allowed for a large number of isotopes. Experimentally accessible are however only the decays of those isotopes, for which single beta decay is energetically impossible, because double beta decay is a second order process ($\propto G_F^2$) that shares its experimental signature with the single beta decay, which is a first order process ($\propto G_F$). This is the case for some isotopes with even number of nucleons. The semi-empirical Bethe-Weizsäcker formula

$$E_B = a_V A - a_S A^{2/3} - a_C \frac{Z^2}{A^{1/3}} - a_A \frac{(A - 2Z)^2}{A} - \delta(A, Z) \quad (2.19)$$

describes the nuclear binding energy of isotopes with mass number A and proton number Z . For isotopes with the same mass number A , the binding energy is quadratic in Z . The term $\delta(A, Z)$ describes the effect on the nuclear binding energy due to spin coupling of the nucleons. It is roughly given by

$$\delta(A, Z) \approx \begin{cases} -\frac{12 \text{ MeV}}{\sqrt{A}} & Z, N \text{ even (A even)} \\ 0 & A \text{ odd} \\ +\frac{12 \text{ MeV}}{\sqrt{A}} & Z, N \text{ odd (A even)} \end{cases}. \quad (2.20)$$

Thus for an even mass number A , one obtains two parabolas in Z , which are separated by approximately $24/\sqrt{A}$ MeV. For isobars with even A it can therefore happen that more than one beta-stable isobar exists. This is illustrated in figure 2.3 on the example of $A = 136$. While standard beta decay is energetically forbidden, it is possible for such

2. Neutrino mass and double beta decay

isobars to decay to the state with lowest energy via double beta decay ($2\nu\beta\beta$):

$$(Z, A) \rightarrow (Z + 2, A) + 2e^- + 2\bar{\nu}_e \quad (2.21)$$

or its analoga where either two positrons are emitted or one or two electrons are captured and one or zero positrons get emitted. This process is predicted by the Standard Model and shown in the left diagram of figure 2.4. It is of second order in the Fermi constant G_F and therefore highly suppressed. The rate of the decay is




Figure 2.4.: Feynman diagram of double beta decay. Left: The Standard Model allowed mode where two electrons and two electron anti-neutrinos are emitted. Right: The lepton number violating zero neutrino mode on the example of Majorana neutrino exchange. Regardless of the model by which the zero neutrino mode is created it can only occur if neutrinos are Majorana fermions.

$$[T_{1/2}^{2\nu}]^{-1} = G^{2\nu}(E, Z) |\mathcal{M}^{2\nu}|^2, \quad (2.22)$$

where $G^{2\nu}(E, Z)$ is the phase space factor and $\mathcal{M}^{2\nu}$ a nuclear matrix element which has to be calculated. Various models for the calculation of the nuclear matrix element exist, however they only agree up to a factor of 2 [31]. Standard double beta decay has been observed in several isotopes. Table 2.2 gives an overview over these isotopes and their half-lives. Although two neutrino double beta decay is a Standard Model process, its study is still valuable since it is an irreducible background in neutrinoless double beta decay experiments and the measurement of the processes' half-life provides valuable input to the models that calculate the nuclear matrix elements in 2.22 and those relevant for neutrinoless double beta decay.

If neutrinos are Majorana particles, the possibility exists that double beta decay can occur without the emission of (anti-)neutrinos:

$$(Z, A) \rightarrow (Z + 2, A) + 2e^-. \quad (2.23)$$

In fact, if the process is observed, neutrinos must be Majorana particles, independent of the mechanism by which it is caused. This is called the black box theorem [39]. A simple

Isotope	$2\nu\beta\beta$ half-life [10^{19} yr]	$0\nu\beta\beta$ half-life [10^{22} yr]	Experiment
^{48}Ca	$4.4^{+0.5}_{-0.4} \pm 0.4$	> 1.3	NEMO-3[32]
^{76}Ge	184^{+9+11}_{-8-6}	> 2100	GERDA[33, 34]
^{82}Se	$9.6 \pm 0.3 \pm 1.0$	> 36	NEMO-3[32]
^{96}Zr	$2.35 \pm 0.14 \pm 0.16$	> 0.92	NEMO-3[32]
^{100}Mo	$0.717 \pm 0.002 \pm 0.054$	> 110	NEMO-3[32]
^{116}Cd	$2.6 \pm 0.1^{+0.7}_{-0.4}$	> 7	Solotvina[35]
^{130}Te	$69 \pm 9 \pm 10$	> 280	NEMO-3[32], CUORICINO[36]
^{136}Xe	$216.5 \pm 1.6 \pm 5.9$	> 1100	EXO-200[13]
^{136}Xe	$238 \pm 2 \pm 14$	> 1900	KamLAND-Zen[37, 38]
^{150}Nd	$0.911^{+0.025}_{-0.022} \pm 0.063$	> 1.8	NEMO-3[32]

Table 2.2.: Summary of half-lives of various double beta decay isotopes and limits on the neutrinoless mode. If two errors are shown, the first is statistical, the second is systematic.

model that induces neutrinoless double beta decay is the exchange of a light Majorana neutrino as shown on the right side of figure 2.4. Here the neutrino is emitted at one vertex and absorbed at the other. This is clearly only possible if the exchanged neutrino is its own anti-particle. In this case, the complete energy of the decay is shared between the two electrons (and the small nuclear recoil) in contrast to the two neutrino mode, in which the energy is split between the electrons and neutrinos. Experiments that search for neutrinoless double beta decay therefore look for a small excess in the electron energy spectrum around the Q-value of the decay, as illustrated in figure 2.5. If the process is mediated by light Majorana neutrino exchange, the decay rate can be expressed as

$$[T_{1/2}^{0\nu}]^{-1} = G^{0\nu}(E, Z) \left| \frac{\langle m_{\beta\beta} \rangle}{m_e} \right|^2 |\mathcal{M}^{0\nu}|^2, \quad (2.24)$$

where the effective Majorana neutrino mass is given by

$$\langle m_{\beta\beta} \rangle = \left| \sum_i U_{ei}^2 m_i \right|. \quad (2.25)$$

The nuclear matrix element $\mathcal{M}^{0\nu}$ is a quantity that needs to be calculated, as in the two neutrino case. The decay rate is directly proportional to the effective Majorana mass. A measurement of the former hence would fix the absolute neutrino mass scale in the framework of light neutrino exchange. Assuming knowledge of the neutrino oscillation parameters, the effective Majorana mass can be related to the mass of the lightest neutrino within this model, and hence conclusions on the neutrino hierarchy can be drawn. Figure 2.6 shows this relation.

A large number of experiments have searched or are searching for neutrinoless double beta decay in various candidate isotopes. Table 2.3 gives an overview over prominent

2. Neutrino mass and double beta decay




Figure 2.5.: Summed energy of the electrons in double beta decay. The two neutrino mode (black dashed line) as well as the zero neutrino mode (full red line) are shown. In the left (right) plot the decay rate of the two neutrino mode is 100 (10^6) times higher than the zero neutrino mode. The right plot also zooms in around the region of the Q value of the decay.

isotopes and the best limits on the half-life of their neutrinoless decay modes. While a single experiment has the potential to discover neutrinoless double beta decay and hence to prove the Majorana nature of neutrinos, due to the large uncertainty in the nuclear matrix elements, measurements with several isotopes are necessary to provide information on the process by which neutrinoless double beta decay is mediated.




Figure 2.6.: Effective neutrino mass $m_{\beta\beta}$ vs mass of the lightest neutrino. Normal hierarchy is shown on the left, inverted hierarchy on the right. The bands are calculated from the parameters in table 2.1. The width of the bands is mostly due to the ignorance of the Majorana phases. Curves corresponding to different extreme values of Majorana phases and 1,2 and 3 sigma error bands are shown.

Isotope	Q [keV]	$G_{0\nu}^{(0)}$ [10^{-15} yr^{-1}]	IBM-2	QRPA	ISM
^{48}Ca	4272	24.81	1.98		0.54
^{76}Ge	2039	2.363	5.42	4.68	2.22
^{82}Se	2995	10.16	4.37	4.17	2.11
^{96}Zr	3350	20.58	2.53	1.34	
^{100}Mo	3034	15.92	3.73	3.53	
^{116}Cd	2814	16.70	2.78	2.93	
^{130}Te	2527	14.22	4.03	3.38	2.04
^{136}Xe	2458	14.58	3.33	2.22	1.70
^{150}Nd	3371	63.03	2.32		

Table 2.3.: Q -values, phase space factors $G_{0\nu}^{(0)}$ [40] and nuclear matrix elements [31] for different double beta decay isotopes.

3. The EXO-200 detector

After the theoretical motivation to search for neutrinoless double beta decay presented in the last chapter, this chapter will describe how such a search can be performed in the candidate isotope of ^{136}Xe . The first section lays out the general strategy and difficulties of such a rare event search. Due to its unique properties, ^{136}Xe is one of the most promising among the neutrinoless double beta decay candidate isotopes. This will be motivated in section 3.2. Finally, the EXO-200 detector, a liquid xenon time projection chamber searching for neutrinoless double beta decay, will be described.

3.1. Sensitivity considerations

Assuming a constant decay rate, the number of radioactive decays of an element within an observation time t follow a Poisson distribution. In double beta decay experiments, the observation times are much smaller than typical isotope half-lives ($T_{1/2}$), and therefore the constant rate assumption is fulfilled.

The expected number of observed signal events N_s then scales with observation time as

$$N_s = \epsilon \frac{aM}{m_a} \frac{\ln 2}{T_{1/2}} t, \quad (3.1)$$

where ϵ is the efficiency for detecting an event, a is the isotopic abundance, M is the total mass of the decaying material, and m_a the mass of a single atom. It makes sense to define the experimental sensitivity as the median value of the lower limit on the half-life the experiment will report if no signal is observed. For a simple counting experiment, N_s follows a Poisson distribution from which an upper limit on the true number of signal events can be constructed, given that no signal events are observed. The upper limit on N_s can then be translated into a lower limit on $T_{1/2}$ using equation 3.1. In an ideal background-free experiment the sensitivity on $T_{1/2}$ is then proportional to

$$S(T_{1/2}) \propto \frac{\epsilon a M}{m_a} t. \quad (3.2)$$

No experiment however is background-free, so if one assumes a background b (in units of counts per energy per exposure) that is constant over the energy region of interest and constant in time, then the expected number of observed background events in the region of interest is proportional to

$$N_b \propto \epsilon b \Gamma M t, \quad (3.3)$$

3. The EXO-200 detector

where the energy resolution Γ of the experiment defines width of the region of interest. The experiment then observes

$$N_{tot} = N_s + N_b \quad (3.4)$$

events and therefore the uncertainty on the true number of signal events N_s must be calculated from the Poisson distribution of N_{tot} , whose uncertainty ΔN_{tot} will be dominated by the uncertainty on the number of background events N_b , because $\Delta N_{Poisson} = \sqrt{N}$. This weakens the upper limit one would extract if the experiment were background free, and the sensitivity of a realistic experiment instead goes like

$$S(T_{1/2}) \propto \frac{\epsilon a M}{\Delta N_{tot} m_a} t \propto \frac{\epsilon a}{m_a} \sqrt{\frac{Mt}{b\Gamma}}. \quad (3.5)$$

A competitive experiment therefore needs to maximize exposure Mt and at the same time minimize the radioactive background b to the smallest possible level. The better (smaller) the energy resolution of the detector, the less background events pollute the signal region and dilute the signal sensitivity. The candidate isotope must either be available with high natural abundance or reasonably well enrichable and the detector should be designed for high detection efficiency. EXO-200 has been designed to meet these requirements.

3.2. Search for $0\nu\beta\beta$ in liquid ^{136}Xe

There is no single best isotope that is optimal in all properties important for an experimental search for neutrinoless double beta decay. Even if there was one - due to the large computational uncertainties in the nuclear matrix elements, searches in several isotopes are necessary to finely constrain the effective neutrino mass. ^{136}Xe nevertheless has many desirable qualities which make it one of the first choices among the 35 candidate isotopes that undergo double beta decay.

^{136}Xe has a relatively large Q value of 2457.8 keV [41]. This has two advantages. First, according to Sargent's Rule, the decay rate is proportional to Q^5 , and so isotopes with larger Q values are expected to have lower half lives. Second, many background gamma lines have lower energies and therefore do not contribute events to the region of interest. Relevant exceptions to this are ^{208}Tl which emits a gamma at 2615 keV [42], ^{214}Bi which emits a gamma at 2448 keV [43], and ^{60}Co whose summed gamma peak is at 2505.7 keV [43]. The former is part of the thorium decay chain, the latter appears in the uranium chain. The radionuclides produced by cosmogenic activation of xenon are short lived, which requires only short storage times in an underground facility before it can be used in a low background experiment. The relatively high density of liquid xenon of 3 g/cm³ at 167 K [44] together with its high atomic number provides good self shielding and the fact that the source- and detector material are the same, reduces the amount of other potentially radioactive materials needed to construct the detector. Furthermore, as a noble gas, xenon is easy to purify of any chemically reactive contaminants and once enriched can be reused in later experiments. This is particularly useful, because to

explore effective Majorana masses down to the inverted hierarchy, ton-scale experiments will be required. Because the detector material is liquid and non-reactive, it should be relatively easy to scale liquid xenon experiments to high target masses.

3.2.1. Radiation detection in liquid xenon




Figure 3.1.: Mean free path of photons [45] and stopping distance of electrons [46] in liquid xenon. The left (red) axis is for photons, the right (blue) axis for electrons.

Liquid xenon has the property to produce charge carriers (electrons and ionized xenon atoms) as well as scintillation light when exposed to nuclear radiation. The scintillation light is prompt compared to the electron drift velocity in liquid xenon for typical electric fields, and the electron diffusion is on the order of millimeters for typical fields and drift times [47]. One can model the electronic structure of liquid xenon as a band structure with a gap between the valence and conduction band of 9.22 eV [48], which makes it a good isolator. The range of electrons and gamma radiation is shown in figure 3.1. Electrons have far less reach than photons, whose dominating interaction process is Compton Scattering for energies around the ^{136}Xe Q-value. Energy depositions from radioactive gamma backgrounds will therefore show a different topology than double beta decay events, which can be used to separate these background events from signal events in a suitable detector. These properties make liquid xenon an ideal material for a time projection chamber (TPC).

High energy radiation loses energy in matter by ionizing or exciting the atoms in the material. Scintillation light is emitted when the excited states decay to their ground state. A fraction of the created electron-ion pairs recombine and emit scintillation light as well. The remaining charge carriers can be drifted to electrodes where their voltage pulse is detected.

Ionization

The energy loss of an ionizing particle in liquid xenon can be expressed as

$$E_0 = N_i E_i + N_{ex} E_{ex} + N_i \epsilon \quad (3.6)$$

where N_i, E_i are the number of ionized atoms and average energy required to ionize an atom, N_{ex}, E_{ex} are the number of excited atoms and the average energy required to excite an atom, and ϵ is the average kinetic energy of ionized electrons. Hence the average energy to create an electron ion pair is

$$W := E_0/N_i = E_i + \frac{N_{ex}}{N_i} E_{ex} + \epsilon \quad (3.7)$$

which has been measured to 15.6(3) eV in liquid xenon [49]. Comparing this to the band gap value of 9.22 eV yields that $\approx 60\%$ of the radiation energy goes into ionization and $\approx 40\%$ into excitation. The electrons created in the ionization process can be collected on an electrode by applying an electric field. Not all of these will arrive there however. Some fraction of the electron ion pairs will recombine and produce scintillation light. In addition, a certain amount of electrons can attach to electronegative impurities to form ions during the drift through the liquid xenon. Due to their larger mass, the ions drift much slower and therefore escape from the main charge signal peak at the anode. The main contribution comes from oxygen or oxygen molecules such as water and nitrogen because of their strong electronegativity and natural presence in air. The number of surviving electrons can be modeled by


$$\frac{dN_S}{dt} = -k_S N_e N_S \quad (3.8)$$

where k_S is the capture rate for the impurity S and N_e, N_S are the concentrations of electrons and impurities respectively. In addition to the type of impurity, k_S also depends on the electric drift field. The number of drift electrons thus decreases exponentially with a lifetime of

$$\tau_e = \frac{1}{k_S N_S} \quad (3.9)$$

Scintillation

Scintillation light in liquid xenon is produced by the decay of excited dimers, Xe_2^* which are produced directly by the incident radiation and indirectly by the recombination of xenon ions Xe^+ with ionized electrons. The first process is



The recombination process is



and the excited xenon atom decays as in the first process. The produced scintillation light peaks at 177.6 nm and has two different decay rates due to the deexcitation of singlet and triplet states of the excited dimer. The former has a decay time of 2.2 ns and the latter has a decay time of 27 ns [50].

Energy resolution

The energy resolution of a detector is directly related to the number of quanta, i.e. electron ion pairs and photons in this case, that are generated per energy of the incident radiation. If ionization and excitation were Poisson processes, the uncertainty in the number of created quanta would just be the square root of the number itself. The succeeding processes that create single quanta from one incident radiation particle are not independent however, and so the uncertainty in the number of created quanta is instead

$$\Delta N = \sqrt{FN} \quad (3.12)$$

where F is called the Fano factor, which has been calculated to be 0.059 for liquid xenon in the model that gives the largest value [51]. The energy resolution can then be expressed as

$$\frac{\sigma}{E} = \sqrt{\frac{FW}{E}} \quad (3.13)$$

which would result in an energy resolution of 0.06 % for liquid ^{136}Xe at the Q-value. However no liquid xenon experiment to date has ever come close to such good resolutions. The reasons for this remain unclear.

A fraction of the electron ion pairs created by ionizing radiation recombine according to equation 3.11. This fraction depends on the electric drift field and on the density of the ionized matter. The higher the electric field in the detector, the faster the electrons drift to the anode and the ions to the cathode and the less time remains for the recombination process to happen. Similarly, a higher ionization density leads to a higher recombination rate, because a free electron is more likely to find a nearby ion. The latter can be used to identify different types of radiation. Alpha particles lose their energy within a very short distance compared to electrons. The electron ion cloud created by alpha particles is thus much denser than those created by electrons (and those electrons created by photon scattering). As a consequence the light to charge ratio is much higher for alphas and can be used to reject such events.

Because the recombination process reduces the number of free charges and increases the number of photons, the amount of detected charge and scintillation is not independent but highly anti-correlated. Using a combination of the amount of charge and light as an energy estimator thus will result in better energy resolution compared to detectors that are only sensitive to one channel.

3.2.2. Background-free detection of double beta decay

Via double beta decay, a ^{136}Xe nucleus transitions to doubly ionized $^{136}\text{Ba}^{++}$, as shown in figure 2.3. If one would succeed in detecting the barium ion by any means, this would be

a definite evidence that the energy deposited in this event must have come from double beta decay of ^{136}Xe , since no other realistic process could create doubly positively charged barium ions. If the ions can be detected with acceptable efficiency, the only background for a ^{136}Xe neutrinoless double beta decay experiment would thus be the barium ions created in the two neutrino decay mode. Since liquid xenon detectors usually apply an electric drift field to collect the charges created in the ionization process, the barium ions created by double beta decay are continuously swept towards the cathode and thus do not pollute the area within the detector at which they should be detected.

Because of this attractive feature, great efforts are made (e.g. [52–54]) to detect the barium ion within a double beta decay detector scenario. The desire is to be able to exploit such a barium tagging technique in the next generation of xenon double beta decay experiments, but progress is made in small steps. During the course of this thesis thus, some amount of time was dedicated to the extraction of ionized barium atoms from liquid xenon.

3.3. The Time Projection Chamber




Figure 3.2.: Schematic of the EXO-200 TPC. Radiation in the liquid xenon creates free electrons and prompt scintillation light. The light is collected by arrays of avalanche photo diodes on both ends of the TPC. Charge drifts along the homogeneous electric field and passes the induction (V) plane before it is deposited on the anode (U) plane. From the induction signals in the V plane and the charge deposition signals on the U plane, the event position perpendicular to the drift direction can be reconstructed. The position along the drift direction is determined by the product of the time difference between charge and light signals and the drift velocity.




Figure 3.3.: TPC during insertion into the cryostat (right). The six supporting legs that contain signal cables and the xenon supply lines as well as the high voltage feedthrough in the middle of the cylinder can be seen.

The main part of the experiment is a double Time Projection Chamber (TPC) filled with liquid xenon enriched to $80.672\% \pm 0.014\%$ ^{136}Xe , which serves as both the source and detector material. It is split into two almost identical parts by a common cathode. See figure 3.2 for an illustration of the concept. The detector is surrounded by a cylindrical vessel made out of ultra-low radioactivity copper. The vessel's diameter is 45.8 cm. Like all other materials used to construct the detector, this was chosen after an extensive study of radioactive impurities in candidate materials [55]. The walls of the vessel are only 1.37 mm thick in order to place as little radioactive material near the detector region as possible, while still providing enough pressure stability for the vessel. Besides the copper used for the support structure and walls, the predominant other materials are bronze for conductors and acrylic, polytetrafluorethylene (PTFE, Teflon) and polyamide for dielectrics. The vessel is supported by six legs that position it inside a doubly walled cryostat (see figure 3.3). In addition to that, the legs contain the signal cables for the wires and the supply lines for xenon recirculation. Construction of the TPC was done below a concrete overburden to keep cosmic activation of the materials at a minimum level.

3.3.1. Charge detection

The TPC is divided into two drift regions by a central cathode grid which is set to a voltage of -8 kV (see figure 3.4). The anodes on each side of the TPC consist of two

3. The EXO-200 detector




Figure 3.4.: Photograph of one half of the TPC. The cathode grid can be seen on top. Field shaping rings made out of copper surround the cylindrical Teflon reflector. They are connected by a resistor chain (not visible on the photograph).




Figure 3.5.: Photograph of one anode plane. The copper support structure holding the two wire planes is shown. The electronic readout connection (yellow band cable) for the wires on the right side of the picture can be seen. In the background is the APD platter (without the APDs filled in).

wire planes, a collection plane and an induction plane. Cathode and induction planes are separated by 198 mm. The spaces in between are referred to as the main drift regions. The collection planes are kept at virtual ground and are used to collect the drifting electrons. The induction planes which are placed 6 mm in front of each collection plane serve two purposes. They detect induction signals produced by the electrons that drift towards the collection plane and shield the collection plane from fields created by charges in the main drift region. In order to provide full transparency for the drift electrons, the voltage of the induction planes is set to -780 V which provides an approximately two times higher electric field between the collection and induction plane than in the main drift region. Each of the 4 wire planes consist of 114 wires with a pitch of 3 mm which are mounted on a hexagonal support structure. The wires were constructed by photo etching a thin phosphor bronze plate. They are of roughly square shape with a side length of $125\text{ }\mu\text{m}$. A readout channel comprises three wires each, which are physically connected. The wires on the induction plane (V-wires) are at an angle of 60° to the ones on the collection plane (U-wires). This allows to reconstruct the position of charge arriving in this plane. The U (V) coordinate is defined as the direction perpendicular to the U-wires (V-wires). Figure 3.5 shows a part of one of the wire planes and their mounting on the support structure. A very uniform electric drift field is necessary to provide a uniform detector response. Accurate position reconstruction requires the electrons to drift in straight lines and with uniform drift velocity perpendicular to the anode and the magnitude of the field needs to be constant throughout the detector to yield a uniform charge to light ratio as discussed in section 6.2.3. To provide this, a set of field shaping rings is aligned along the detector axis surrounding the drift region as shown in picture 3.4. The field shaping rings are electrically connected via a chain of resistors sitting on an acrylic block. Each of them has a resistance of $900\text{ M}\Omega$.

3.3.2. Scintillation detection

The 178 nm scintillation light is detected by 234 large area avalanche photo diodes (LAAPDs or APDs) on both ends of the TPC. The devices are 1.3 mm thick and their diameter varies between 19.6 mm and 21.1 mm . Due to their manufacturing process and small size they contain much less intrinsic radioactive contaminations than photo multiplier tubes. This advantage is mitigated by the fact that they typically yield higher electronic noise and less gain. To reduce the amount of possibly radioactive mass, the APDs are mounted without their standard housing, which also increases the packing density and the light efficiency, because the front window is missing. The electrical insulation is instead provided by the liquid xenon which is a good isolator. A platter on both ends of the TPC holds the APDs in place (see figure 3.6) and provides a voltage of -1400 V to reverse-bias the APDs. In addition, this voltage prevents the electrons from drifting any further and ensures that all charges end up on the anode plane. The platters are coated with gold on the back side to provide good electric contact with the diodes. The APDs are wired up in groups of 7 (with some exceptions in which a group only contains 5 or 6). Because the voltages required to provide the same gain on each individual device varies about 2% [56], similar devices are grouped together and trim



Figure 3.6.: Back (non-active) side of the light collection plane. The APDs are partially filled in and in the lower half the cabling is already installed.

voltages are applied via the signal cables to bring the groups to a common gain level. Overall, 48% of the platter surface are photosensitive. The non-photosensitive parts of the platter and detector boundaries are designed to provide good reflectivity of the ultraviolet scintillation light. For this purpose, the front of the APD platters is coated with aluminum and magnesium fluoride. The drift volume is surrounded by a Teflon cylinder jacket as visible in figure 3.4, which provides high reflectivity for the scintillation light of liquid xenon [57].

3.4. Data readout

Figure 3.7 shows a schematic of the detector's readout electronics. All signals are guided through flex cables of ≈ 1 m length to the front end electronics which are located outside the cryostat at room temperature. This reduces the amount of possibly radioactive material near the detector and allows easy maintenance access or exchange of the electronics. Part of the cabling for one APD plane can be seen in figure 3.6. The flexible flat signal cables are made in layers. The signal trace layer is confined by two solid copper / Kapton layers on either side. These have the purpose of electrical shielding and provide the bias voltages for the V-planes and APD-planes with low resistance (< 50 m Ω). The signal traces are made by photo-etching polyamide coated copper strips and have a resistance of ≈ 1 Ω . In the U- and V-wire cables every other trace does not carry a signal but is used for electrical shielding. The signal processing is the same for all three types of signals. They only differ by the shaping times which are listed in table 3.1. After amplification and feeding through two differentiators and integrators, the signals




Figure 3.7.: Schematic of the read out electronics. The front end electronics are housed outside the cryostat at room temperature. The signals are guided through flex cables of ≈ 1 m length and then amplified, shaped and digitized. The digitized signals are fed into a trigger module (TEM) which instructs the DAQ to record signals.

are digitized by a 12 bit Analog to Digital Converter (ADC) to continuous waveforms with a sampling frequency of 1 MHz. The digitized signals are fed into a Trigger Event Module (TEM) via optical fibers of ≈ 20 m length. If any individual waveform or the summed waveforms exceed the trigger threshold, the TEM constructs an event which consists of the waveforms of all channels 1024 μ s before and after the trigger time. The trigger threshold is very permissive and set to prevent storing large amounts of data on which only pure noise is present. In addition to the physics trigger, there is also a forced trigger which constructs an event every 10 seconds. These forced triggers are used to monitor the DAQ and to calculate the detector live time. An event consists of 76 U-wire, 76 V-wire and 74 APD waveforms each with 2048 samples centered around the trigger time. The average event rate is ≈ 0.6 Hz. The DAQ writes all events to removable disks, which are periodically shipped to the SLAC national laboratory, where the data undergoes offline processing.

Signal Type	Shaping Type and Time [μ s]				
	D (Preamp)	D	I	D	I
APD	300	10	3	10	3
U	60	40	1.5	40	1.5
V	60	10	3	10	3

Table 3.1.: Signal shaping stages in the order in which they are wired. 'D' denotes a differentiation stage, 'I' and integration stage.

3.5. Detector environment

The TPC is surrounded by a heat transfer fluid to keep the temperature stable and uniform around the detector walls. The TPC and the heat transfer fluid HFE-7000¹ are contained in a double walled cryostat made from 2.7 cm thick low radioactive copper plates. The inner wall of the cryostat is wrapped in an aluminized polyester foil and the space between inner and outer wall is pumped to high vacuum for thermal insulation. Besides its temperature stabilizing properties, the HFE with its density of 1.8 g/cm³ at 170 K also acts as an additional shield for radioactivity. The outer cryostat is surrounded by a lead wall of 25 cm on all sides. The supply systems and electronic connections enter the cryostat at the front. A second lead wall was constructed at this side to prevent gamma radiation to enter the detector within line of sight of the supply lines. The cryostat, its surroundings, the xenon and HFE supply systems are housed in a class 100 clean room, which is flanked by muon veto panels on four sides. A schematic of the detector's environment is shown in figure 3.9.

3.5.1. Density of ¹³⁶Xe in the detector

To translate a measured number of double beta decay events into the corresponding lifetime of ¹³⁶Xe, a precise knowledge of the density (atoms per volume) of this isotope inside the active detector region is necessary. The density of the xenon at the operating temperature of 166.6 K corresponds to a density of 3.03 g/cm³ (see [58, 59]). The following systematic effects are considered:

- Knowledge of the isotopic fraction of ¹³⁶Xe in the detector liquid.
- Different isotopic fraction in the liquid phase compared to the gas phase at which the fraction is measured
- Uncertainties in the xenon density due to temperature variations / uncertainties.
- Dilution of the xenon from natural xenon adsorbed in the plastic components of the detector during the preceding natural xenon test run

¹3MTMNovecTMEngineered Fluid HFE-7000

The enrichment of the xenon was done by ultracentrifugation. The fraction of ^{136}Xe atoms has been measured from a gas sample taken during detector filling using dynamic dual-inlet mass spectroscopy to be 80.672(14) % with a remaining abundance of ^{134}Xe of $19.098\% \pm 0.014\%$ and other lighter isotopes (less than 0.25 %).

The xenon is in gas phase during the measurement in the mass spectrometer, but in liquid phase during the course of the experiment. The differences in vapor pressures between ^{136}Xe and ^{130}Xe at the boiling point have been measured to be 0.01 % [60]. The difference in vapor pressures is reflected linearly as a difference in substance concentration. Considering that the mass difference between isotopes is smaller in EXO-200, since the main residual isotope is ^{134}Xe , the 0.01 % difference is a conservative estimate of this effect.

Uncertainties in the measurement of the xenon temperature are also directly translated to uncertainties in the knowledge of its density. The variation of xenon density with temperature at the operating point has been estimated to $-0.008 \text{ g/cm}^3/\text{K}$ from a combination of the measurements in [58, 59]. The temperature is continually monitored by thermocouples mounted on the cryostat which houses the cooling liquid. These were initially calibrated by a reference thermistor whose uncertainty at the liquid xenon temperature is 0.2 K. The temperature variation measured by the thermocouples during the data taking period is 0.15 K. Combining these two errors quadratically yields a 0.06 % uncertainty on the xenon density in the detector.

The plastic parts inside the can absorb some amount of xenon. This effect was observed as a pressure decrease during the initial filling of the detector. Because the detector was filled with natural xenon before it was exchanged by enriched xenon, it is possible that some of the natural xenon still contained in the plastic parts can dilute the isotopic fraction. This effect is conservatively estimated to be less than 0.04 %.

The combination of above effects results in a total uncertainty on the number of xenon atoms per unit volume of 0.26 %, which is a direct contribution to the systematic error in the half-life measurement of ^{136}Xe decay.

3.5.2. Xenon recirculation and purification

The walls of the TPC are only 1.37 mm thick. Therefore its necessary to keep the pressure difference between the xenon in the TPC and the HFE surrounding it small. A sophisticated system that handles feeding and bleeding, as well as recirculating the liquid xenon in and out of the TPC, ensures this. A simplified schematic of the system is shown in figure 3.8. EXO-200 uses $\approx 200 \text{ kg}$ of enriched ^{136}Xe . Approximately 175 kg (58 l) are in liquid phase in the TPC, of which 110 kg are in the sensitive detector region which is bounded by the wire planes and the Teflon cylinder. The remaining $\approx 25 \text{ kg}$ are in the xenon system and in gas bottles. A slow controls program feeds or bleeds xenon into the system as required to keep the pressure difference between the inside and outside of the TPC at a constant small level. In addition to regulating the pressure difference, the system continually circulates the xenon through the TPC. The xenon which leaves the TPC is heated and transferred into gas phase. A magnetically coupled piston pump [61] pumps the xenon gas through two hot zirconium purifiers which remove

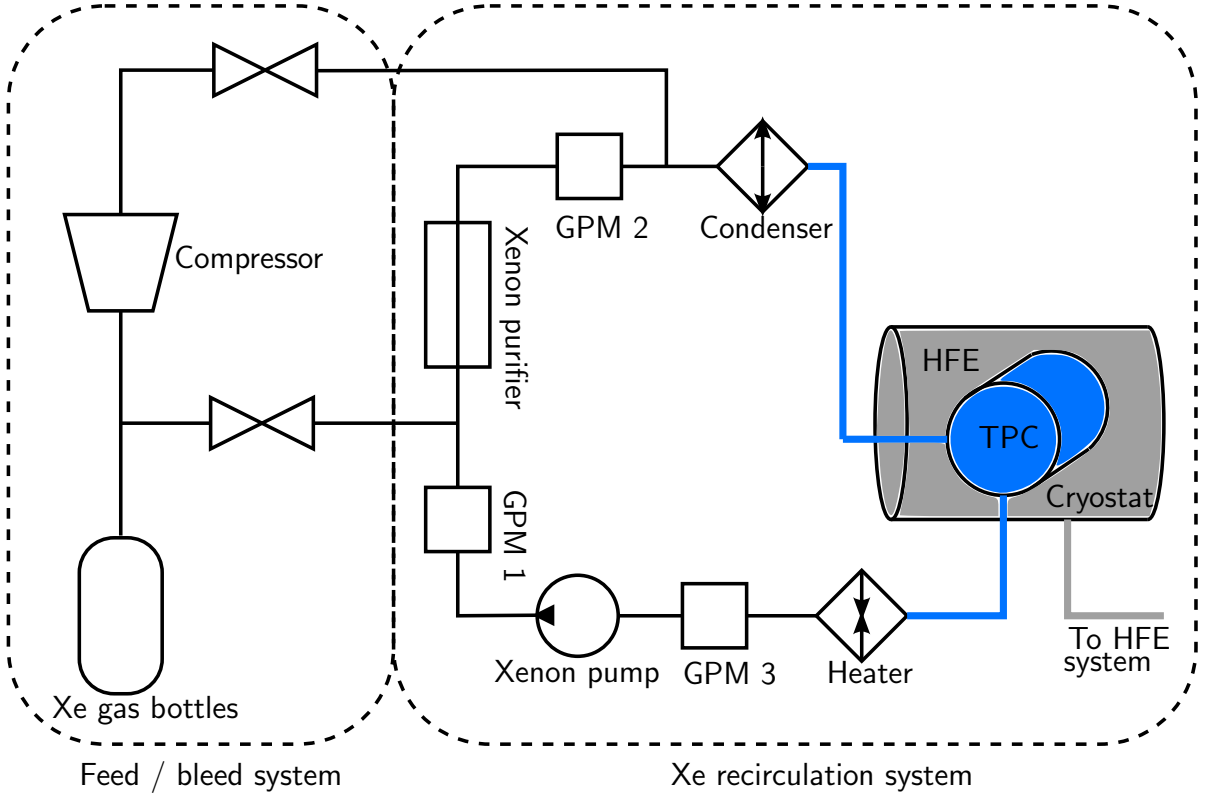


Figure 3.8.: Simplified schematic of the xenon system. The liquid xenon is recirculated continuously during normal operation. After leaving the cryostat, the xenon is heated and transferred into gas phase. A purification system removes electronegative impurities. The purified xenon is liquefied and re-transferred into the TPC. Three Gas Purity Monitors (GPMs) monitor the xenon purity at different locations in the recirculation loop.

electronegative impurities from the gas. The xenon is cooled down and condensed again before it is transferred back into the cryostat. Within the recirculation loop, three gas purity monitors (GPMs) [62] measure the level of cleanness of the gas. The largest source of impurities are the gas bottles in which the xenon is stored. Therefore the feed system enters the xenon loop right before the gas purifiers.

3.5.3. Underground facility and muon shield

The experiment is located at the Waste Isolation Pilot Plant (WIPP), a salt mine ≈ 655 m underground near Carlsbad, New Mexico. The primary purpose of the site is the permanent disposal of transuranic waste, however a small part of the mine is also used for low radioactive background experiments. The solid rock overburden provides a muon shielding of ≈ 1600 m water equivalent. The radioactive contamination has been determined to $\lesssim 1$ ppm in the salt. The air flow through the mine is high compared to usual sites, and so the radon content in the air of 7 Bq/m^3 is not significantly higher




Figure 3.9.: Schematic of the detector surroundings. The TPC is contained within a bath of heat transfer fluid (HFE), which is maintained at 167 K in a double walled cryostat. The cryostat is surrounded by at least 25 cm of lead on all sides and housed in a class 100 clean room. On four of the six sidewalls, muon veto panels are mounted.

3. The EXO-200 detector

than in surface air.

The vertical muon flux at WIPP has been measured to be $4.0 \text{ Hz/cm}^2/\text{sr}$. Muons that transverse the TPC ionize a lot of xenon atoms and their tracks in the TPC make them easily identifiable as muons. Muons that do not transverse the TPC but miss the sensitive detector region only by small amounts however might still create radioactive backgrounds via spallation neutrons. EXO-200 is therefore surrounded by muon veto panels which are attached to four sides of the clean room in which the TPC is located. The muon veto consist of a total of 30 plastic scintillator panels which are attached to eight photo multiplier tubes each. The efficiency of the veto to detect a muon which traverses the TPC has been measured to be $96.0(5) \%$.

3.5.4. Calibration system




Figure 3.10.: The calibration sources are little beads which are housed in cylindrical steel capsules as shown on the left. The sources can be placed at various positions around the detector via a cable in a copper guide tube which is shown on the right.

isotope	gamma lines (keV)	weak source activity (kBq)	strong source activity (kBq)	half-life (yr)
^{60}Co	1173, 1332	0.530	7.060	5.27
^{137}Cs	637	2.820	13.14	30.2
^{226}Ra	many, see [63]	0.257	-	1600
^{228}Th	2614	1.417	34.04	1.91

Table 3.2.: Activities of EXO-200 calibration sources. The activities for all sources except ^{226}Ra are as of September 1st, 2009 as specified by the manufacturer.

EXO-200 uses four radio isotopes to calibrate the response of the detector. Each of the isotopes is available as a weak and as a strong source as listed in table 3.2. The sources consist of small spheres housed in a cylindrical capsule. The capsule sits at the end of a bead chain which is used to push the source through a guide tubing made from copper. The copper guide tube embraces the detector as shown in figure 3.10, and allows the calibration sources to be positioned at various places around the TPC. Extended calibration campaigns using all sources are done several times per year to provide the basic energy calibration of the detector. A routine calibration using the ^{228}Th source is done a couple of times per week for ≈ 2 hours to track any possible time variations of the energy response and to monitor the electron lifetime.

3.6. Detector Monte Carlo model




Figure 3.11.: Model of the time projection chamber implemented in the simulation.

EXO-200 makes use of a full detector Monte Carlo simulation, built partially using the GEANT4 toolkit [64], to model the detector response to physics processes such as radioactive decays within the detector and surrounding materials. The software consists of a first part including a three dimensional model of the detector and relevant physics

3. The EXO-200 detector

processes and a second part which is capable of simulating the detector's electronic signal response to the processes simulated in the first part.

The first part implements a detailed three dimensional GEANT4 model and material description of the TPC and its surroundings, including the HFE liquid, the copper cryostat and the lead shielding. Figure 3.11 shows the model of the TPC that is implemented within the simulation. Default GEANT4 physics processes are used to simulate the interaction of radiation with the detector and surrounding materials. To store the output of the first step, the sensitive region of the detector is binned into voxels of 0.15 mm side length. If any interaction happens within the detector's sensitive region, the resulting amount of energy deposited into the detector material and its position voxel are recorded as so called pixelated charge deposits (PCDs).

The second part simulates the detector's response to charge deposits in the sensitive region. The amount of ionization created per unit energy is treated as constant and uncorrelated to scintillation. It is based on a numerical simulation of the detector's electrostatic properties using MAXWELL [65] and the Shockley-Ramo theorem [66, 67], which allows to calculate the signal amplitudes given a weighting potential and the charge drift path. The MAXWELL simulation uses a two dimensional model of the wire geometry, in which the V- and U-wires are assumed to be parallel and of infinite length for simplicity. It provides a two dimensional model of the electric field and weighting potential inside the detector. To generate the signals on the detector's 76 wire channels, the charge contained in the PCDs is drifted along this drift field, and the voltage induced on each wire channel at each time step is calculated using the weighting potential.

The treatment of light signals is somewhat simpler. Since the APDs possess short intrinsic rise times of 10 ns to 100 ns, their signals are simulated as step functions before shaping. The amplitude of light pulses is calculated using a energy and position dependent response function, which was obtained from a Monte Carlo study of the TPC's optical properties.

All channels are shaped with their respective electronics transfer function and random white noise with an amplitude approximately the size of real noise is added. The second part outputs waveforms in the same format as real data, which allows simulated data to be treated in the same way as physics data.

4. Event reconstruction

As presented in section 3.4, the data output of the detector is organized in so called events, where each event is created after a trigger in either the charge or the light channels. The following chapter describes the reconstruction of the amplitude, time and position of all energy deposits contained in the event from the raw data.

Each trigger creates an event, which consists of the waveforms of all 226 channels, digitized with a sampling rate of 1 MHz. Each waveform contains 2048 samples which are centered around the trigger time. These waveforms are stored in a binary format by the DAQ on removable disks and shipped to the SLAC national laboratory for processing. At SLAC, the data is converted into a ROOT data format which is the base for further processing. The data processing will be discussed in more detail in section 6. The event reconstruction process is done in three steps. The first step (section 4.1) searches for signals on individual waveforms and provides a first guess for the timing, amplitude and number of signals. In the second step (section 4.2), these guesses are refined by fitting signal templates to the raw waveforms. This provides accurate information on the signal timing and amplitude. The third step (section 4.4) combines all signals that were found in the previous steps to provide information about the event positions and its topology. Since simulated data is available in the same format as detector data, it is treated in exactly the same way by the reconstruction process.

4.1. Signal finding

The reconstruction process starts by searching for signals on each individual wire channel and on the summed waveforms of all APD channels on one plane. The shape of the signals on U- and APD-channels is almost fully determined by the readout electronics in the signal shaping stage described in section 3.4. The scintillation flash is converted to a voltage jump in the APDs, which is acting as a step signal input to the signal shaping electronics. The voltage rise in the U-channels is almost instantaneous too, because the U-wire plane is electrostatically shielded from the drifting charge by the V-wire plane and the drift time between the V-wire plane and the U-wire plane as well as the extent of the electron cloud is small compared to the shaping times used for U-wire channels. The shape of signals on V-wire channels is determined by both, the shaping electronics as well as the drift of the electron cloud. Figure 4.1 shows typical example waveforms for the three channel types.

Since the expected shapes of the signals are known or can be modeled, a matched filter can be applied to the waveforms. The matched filter correlates the input waveforms $w(t)$ with a template waveform of the expected signal shape $s(t)$, i.e. the output of the

4. Event reconstruction




Figure 4.1.: Example waveforms for the three channel types. Top left: U-wire signal. Top right: V-wire signal. Bottom left: APD signal. Each waveform consists of 2048 samples.

matched filter is defined as:

$$o(t) := (s \star w)(t) = \int_{-\infty}^{\infty} s^*(\tau)w(t + \tau)d\tau \quad (4.1)$$

If the input waveform $w(t)$ is equal to the expected signal $s(t)$, then $o(t)$ is just the autocorellation function of the input waveform. For real waveforms this is only approximately the case. If $w(t)$ is has the form of the expected signal $s(t)$ with added Gaussian white noise, then the matched filter is the filter with optimum signal to noise ratio. In practice, equation 4.1 is used in the Fourier domain:

$$\mathcal{F}(o) = \mathcal{F}(s \star w) = \mathcal{F}(s)^* \cdot \mathcal{F}(w) \quad (4.2)$$

where the cross correlation is just a simple multiplication. Hence the matched filter weights the spectrum of the input waveform with the spectrum of the expected signal.




Figure 4.2.: The signal finding steps for a U-wire waveform (containing two consecutive signals), zoomed-in around the signal region. The raw waveform is shown top left. The waveform is transformed using a matched filter (top right) which is good at deciding whether a signal is present. Waveforms that exceeded the matched filter threshold are unshaped and reshaped digitally with short differentiation times (lower figures). This allows to detect signals following each other in short time.

In reality, the assumption of white Gaussian noise on the waveforms does not hold. APD waveforms are therefore pre-whitened by dividing by a measured noise power spectrum when the matched filter is applied. The noise power spectrum of U-wire and V-wire waveforms was not measured after an electronics change and hence cannot be used to pre-whiten their waveforms, however the signal to noise ratio for these is already much better than for APD waveforms.

Signal templates for the matched filter are generated by applying the transfer functions corresponding to the shaping electronics (table 3.1) to estimates of the detector signal output. The detector output of the APDs is modeled as a simple step function. The output for U- and V-wires is modeled by a simulated charge drift in a two dimensional electrostatic detector model.

After the waveforms have run through the matched filter, an estimate for their noise amplitude is calculated. This is done by calculating the mean absolute deviation (MAD) from the baseline of the waveforms. The MAD is evaluated in an iterative process that

consecutively removes parts of the waveforms that exceed a value of 3 times the MAD of the previous iteration. Each peak on the filtered waveform that exceeds a value of 5 times (4 times) the MAD for wires (APDs) is marked as a signal. The amplitude of the signal is guessed from the magnitude of the raw waveform at the time of the peak in the filtered waveform. The top right panel in figure 4.2 shows this process.

While the matched filter is good at deciding whether a signal is present or not, its ability to resolve signals that pile up is very limited. This limitation is not a problem for APD signals, where pileup of two or more signals is very unlikely with an average signal rate of ≈ 0.6 Hz and a pulse duration on the order of $100\ \mu\text{s}$. A gamma particle however can scatter multiple times in the detector, and while this creates only one scintillation flash (within time resolution of the detector), separate electron clouds will result. The time needed for electrons to drift from the center of the detector to the anode is on the order of $100\ \mu\text{s}$. This is on the same order as the pulse duration on U-wire channels, and therefore distinct charge clouds arriving on the same U-wire channel won't be resolved by the matched filter. For this reason, U-wire waveforms on which a signal was found, are digitally unshaped by convolving the waveforms with the inverse transfer function of the hardware signal shapers (table 3.1) and reshaped with a $2\ \mu\text{s}$ short triangular filter. Peaks in this reshaped waveforms are then identified as signals in the same way as done with the matched filter. An example of this is shown in the lower panels of figure 4.2

4.2. Signal parameter estimation




Figure 4.3.: U-wire waveform and fit by a signal template (red). The signal amplitude and the signal time are floating in the fit.

After first guesses on the timing and amplitudes of signals are obtained, a second step tries to refine these parameters. This is done by performing a χ^2 fit of signal templates to the raw waveforms, where the time and amplitude of the signals is allowed to float freely. The signal templates are the same as those used in the matched filter. Figure 4.3 shows

an example fit of a U-wire signal. The fit can handle an arbitrary number of signals on a waveform and is done in an iterative way, in which signals can be combined or removed entirely. If the best fit results in two signals with very similar time and large amplitude uncertainty, the signals are combined into one and the fit is repeated. Likewise, signals that fit to very low amplitude or high amplitude uncertainty are removed and the fit is repeated.

Charge drifting onto a U-wire can sometimes induce a signal on a neighbouring wire which might mistakenly be reconstructed as a charge deposition signal. These induction signals however will have a relatively small amplitude and a shape that differs significantly from the shape of charge deposition signals. This fact is used to tag induction signals. The decision on whether a signal is treated as a charge deposition or an induction signal is based on four metrics:

- The baseline subtracted integral of unshaped induction signals is much smaller than the corresponding value for deposition signals
- Induction signals can only occur if there's a large deposition signal on a neighbouring channel
- Induction signals have sharper rise- and fall off times than deposition signals.
- The fit of a template induction signal gives a better χ^2 value than the default deposition template fit.

Signals that were identified as induction signals are removed from the subsequent analysis.

Wire gain correction

The individual electronics components used to shape and digitize the wire signals can have slightly different gain, resulting in a channel-wise variation of the detector response to a defined amount of charge. To correct for this variation, two different methods are used to measure the electronics gain of each wire channel. The electronics has a capacitor for each readout channel whose capacitance has been precisely measured. The capacitors thus allow to inject a very well known amount of charge onto each channel. From the resulting signal amplitude, the channel-specific gain can be derived. This method is used to gain-correct V-wire signals. A second method makes use of electron positron pair production events provided from a strong ^{228}Th calibration source. The 2615 keV photon from this source can produce an electron positron pair with a combined energy of 1593 keV. These events are used, because electrons (and likewise positrons) lose their energy within a very small volume, and thus usually most of the charge ends up on one U-wire channel. The channel response to this known energy is used to derive a channel-specific gain for U-wire signals.




Figure 4.4.: Reconstruction efficiencies for the three signal types. U-wire efficiencies are estimated using simulated pulses on real data noise-only waveforms. The efficiencies of V-wire and scintillation signals are determined from calibration source data using U-wire signals as energy reference.

4.3. Signal reconstruction efficiencies

Figure 4.4 shows the reconstruction efficiencies of all three types of signals. Since the U-wires have a lower threshold than V-wire and scintillation signals, the amplitude of U-wire signals can be used as an energy reference in the determination of the other two reconstruction efficiencies. Data from a calibration source run, where the source was placed in the cathode plane has been used to extract the V-wire and scintillation efficiencies in figure 4.4. For the U-wire reconstruction efficiency itself however no energy reference exists, and so simulated pulses on top of real detector noise waveforms are used instead to estimate their reconstruction efficiency. This is expected to give an accurate result, since the pulse shape for U-wire signals is known to be modeled very well. The response to scintillation light is not necessarily uniform throughout the detector because it is not fully surrounded by APDs. Placing the calibration sources behind the anode plane however results in a very similar scintillation reconstruction efficiency curve, which allows the conclusion that the detector is fully efficient above ≈ 600 keV.

4.4. 3D event reconstruction

After all reconstructed signals have been gain corrected, they are combined in a process called clustering to retrieve the event's 3D position. The clustering process is done in three separate steps, first bundling nearby signals of the same sort (e.g. U-signals with other U-signals), then determining the z position of charge deposits by associating U-signal bundles with scintillation signals, and finally connecting U-bundles and V-bundles to form charge clusters with full 3D position information.

4.4.1. Signal bundling

Before signals of different type are connected to reconstruct the position of an energy deposition inside the detector, signals of the same type that likely originate from the same energy deposition are lumped together in bundles. As discussed in section 3.2.1, the extent of the charge cloud created by alpha or beta radiation is on the order of several millimeters. In contrast, a gamma particle that Compton scatters, produces several such charge clouds within the detector volume. To reconstruct these charge clouds, U-wire signals are bundled together if they occur on adjacent channels within a time span of $3.5\text{ }\mu\text{s}$ from the largest signal in the bundle. The time of the bundle is defined as the amplitude weighted sum over all signals in the bundle and its U-coordinate is an amplitude weighted average of the wire locations. Typically, one or two channels are comprised within a bundle, in agreement with the discussion above. In less than 5% of the cases, a bundle spans across more than two channels.

A similar treatment is done with the V-wire signals from the induction plane. However the structure of the signals that are created when charge drifts by this plane is different from the structure of U-wire signals. Even very localized charge clouds usually induce a signal in several of the channels that are near the drift path. In addition, signals on channels further away from the drift path are systematically reconstructed at earlier times due to electrostatic shielding effects from other wires. To take this into account, V-wire signals on adjacent channels are bundled according to the following relationship:

$$|t_i - 2.97\text{ }\mu\text{s}/\text{chan} \cdot \Delta V - t_0| \leq 4.5\text{ }\mu\text{s}, \quad (4.3)$$

where t_i are the times of the individual signals, t_0 is the time of the largest signal in the bundle, and ΔV is the absolute channel number difference between the individual signal and the signal with the largest amplitude. In contrast to U-signal bundles, t_0 also defines the time of the V-signal bundle. The V coordinate of the bundle is the amplitude weighted average of the wire locations.

Scintillation signals are reconstructed on the sum waveform of each of the two TPC halves and hence one scintillation flash will result in at most two scintillation signals from opposite TPC halves. These signals will be bundled together if they occur within $6\text{ }\mu\text{s}$. This time is defined by the sum of the integration time in the APD signal shaping electronics (table 3.1).

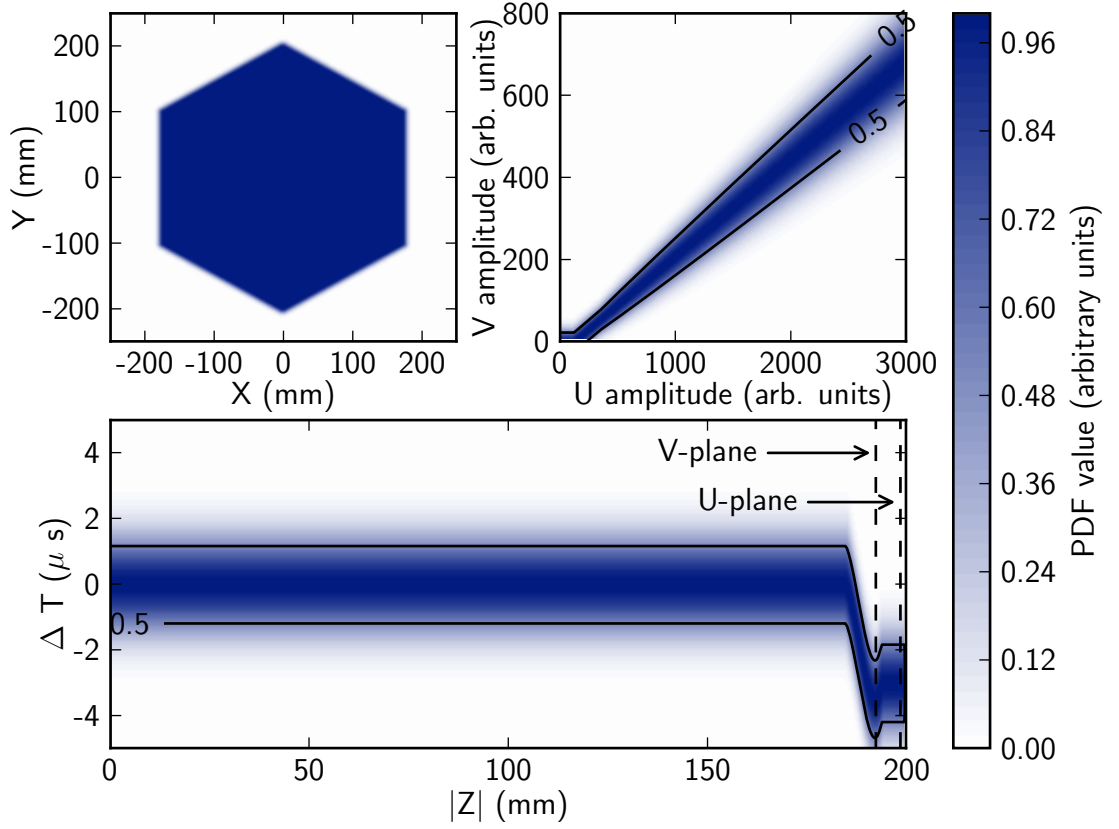


Figure 4.5.: The three PDFs used by clustering. Upper left: The position PDF which heavily penalizes positions outside the wire hexagon. Upper right: The signal amplitude PDF makes use of the correlation between U-signal and V-signal bundle amplitudes. Bottom: The PDF weighs the time difference of U-signal and V-signal bundles.

4.4.2. Determination of Z position

The second clustering step aims to reconstruct the Z position of the energy depositions in the detector. This is done by connecting U-signal bundles with scintillation signal bundles. Since the scintillation is prompt, U-signal bundles are only connected to those scintillation bundles that occur between $3\text{ }\mu\text{s}$ after, or the maximum possible charge drift time plus $3\text{ }\mu\text{s}$ before the time of the U-signal bundle. The $3\text{ }\mu\text{s}$ allowance on each side again corresponds to the integration time for scintillation signals. From these scintillation signal bundles, the U-signal bundle is connected to the one with the minimum absolute time difference from the U-signal bundle. The Z coordinate is then calculated as the product of this time difference and the known drift velocity of $1.71\text{ }\mu\text{s}$.

4.4.3. Determination of X - Y -position

In the last step, the X - Y -position of the energy deposition is determined. This is done by finding the most likely combination of U- and V-signal bundles that were created in the first step. The likelihood is determined by three probability density functions (PDFs) that correlate summed signal amplitudes, the times of the signal bundles and their resulting X - Y -position.

The first describes the correlation of the summed signal amplitudes between U-signal and V-signal bundles. The second models the coincidence between V-wire signals and U-wire signals. The difference in the arrival times of U-wire bundles and V-wire bundles is independent of the event location in the detector, except for events near the anode planes. This effect is addressed to the V-wire signal reconstruction process, which results in slight time offsets for energy depositions near the anode planes and is modeled in the PDF. Because of the hexagonal shape of the wire frames, not all U-wires overlap with every V-wire. The third PDF hence ensures that only such combinations of U-signal and V-signal bundles are allowed, which result in physically possible positions. All three PDFs are based on studies done with calibration source and low background data, which are detailed in appendix A. Figure 4.5 gives a visualization of the PDFs.

The likelihood of a connection between a U- and V-signal bundle is defined as the product of the three PDFs. The negative logarithm of this likelihood is used as a cost function for this specific connection. A matching algorithm tests all possible connections between U- and V-signal bundles, including whether multiple bundles of one type correspond to one bundle of the other type, and returns the one set of combinations whose sum of costs divided by the number of connections is lowest. For each connection in this set, a so called charge cluster is formed and linked with the scintillation bundle that was linked to the U-signal bundle in the previous step. Each of these charge clusters thus have full 3D position information. Bundles for which no connection was found are converted into charge clusters with partially or fully missing position information. Since the U- and V-wires are at an angle of 60° , the U- and V- coordinates of charge clusters are transformed into a rectangular coordinate system in the following way:

$$\begin{aligned} X &= \begin{cases} V - U & \text{in TPC 1} \\ U - V & \text{in TPC 2} \end{cases} \\ Y &= \frac{U + V}{\sqrt{3}} \end{aligned} \quad (4.4)$$

4.5. Event topology

Reconstruction of the 3D position of charge clusters allows to use this information in background rejection. Apart from two neutrino double beta decay, the main background signals in EXO-200 stem from gamma radiation. For energies on the order of the double beta decay Q-value, Compton scattering is the dominant interaction process of gamma radiation with the detector material, and their mean free path is on the order of several centimeters (figure 3.1). In contrast, the reach of the electrons from double beta decay

is on the order of several millimeters, which is on the same order as the 9 mm spacing between wire channels. Therefore, gamma radiation is expected to normally result in more than one charge cluster per event, whereas beta radiation events usually contain only one. We use this fact to label each event as either single site (SS) or multi site (MS). The definition of these two categories is:

- An event is labeled as single site, if:
 - It contains exactly one charge cluster within drift time of a scintillation cluster.
 - The charge cluster contains no more than two U-wire signals on different channels.
 - The definition of a charge cluster implies that these signals must stem from adjacent channels.
- All other events with at least one charge cluster are labeled as multi site.

With this definition, more than 90 % of double beta decay events are categorized as single site, whereas approximately two third of gamma background events are multi site.

4.6. Position reconstruction accuracy

The position resolution is determined by the wire pitch in the X - Y -plane and by electronics shaping times, digitizer sample rate and drift velocity in Z direction. Since charge clusters consist of U-wire and V-wire bundles, whose position is the amplitude weighted average over the wire locations on which a signal occurred, the position resolution can be better than the 9 mm wire pitch. The position reconstruction accuracy can be tested to a certain extent using low background data as shown in figure 4.6.

To estimate the resolution in Z direction, the fact that the cathode grid is a weak emitter of alpha radiation is used. Alpha events from the cathode are ideal for this study, since their penetration length is small and the cathode grid is very thin and thus such events are a source of charge at a known physical location in Z direction. The top plots in figure 4.6 show the distribution of these events in Z for both TPC halves. The width of these distributions, 0.73 mm corresponds to the Z position reconstruction uncertainty for a single charge cluster. If the drift velocity and the distance between cathode and anode plane were exactly the same in both TPCs, the distributions should be centered around 0. However they are slightly offset. This could be explained if the cathode were off-centered by ≈ 0.5 mm, which is within the expected tolerance after the vessel was cooled down to cryogenic temperatures.

For an estimate of the uncertainty in X and Y direction, a simple model can be used, which assumes uniform event distributions throughout the detector. All charge depositions within one wire channel pitch are mapped to the same coordinate. The standard deviation of the uniform distribution then is a measure of the position reconstruction uncertainty in the U and V coordinate. For a channel pitch of 9 mm this corresponds

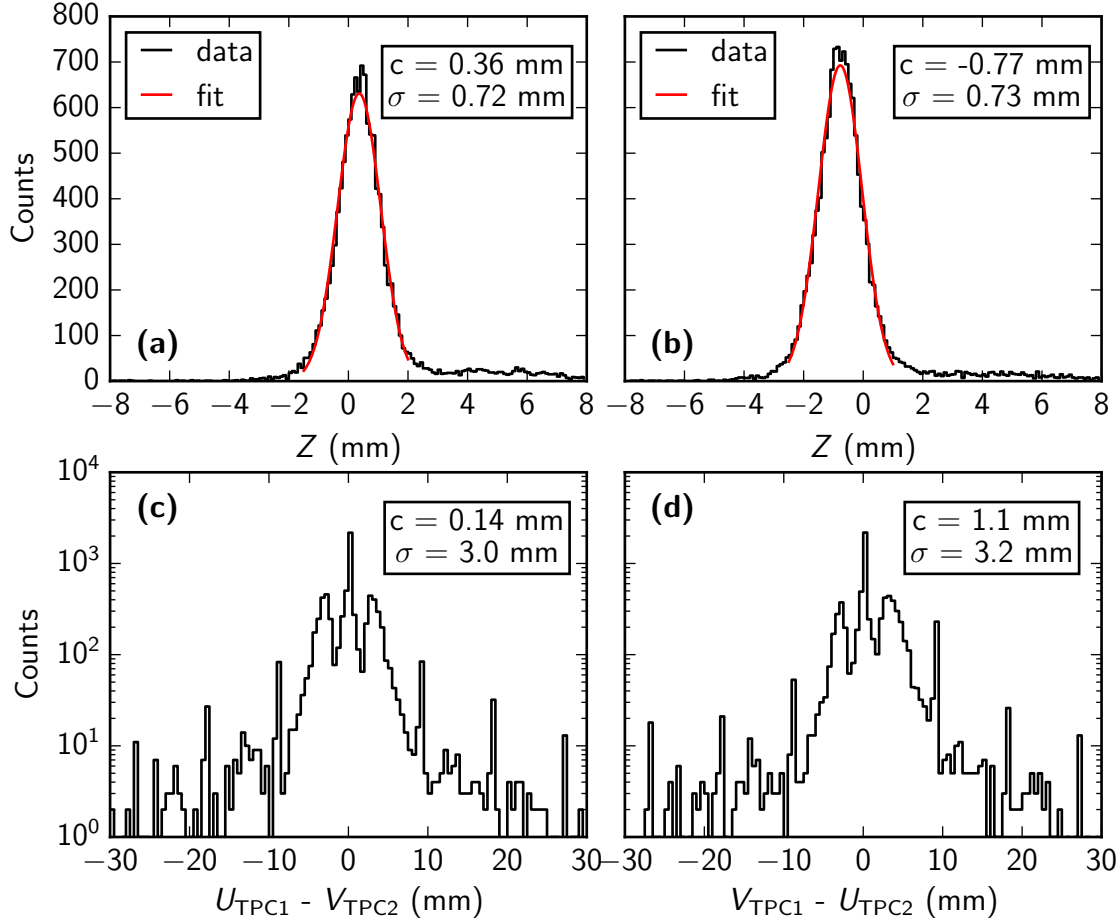


Figure 4.6.: Position reconstruction accuracy from low background data. The mean values (standard deviations) of the distribution are labeled as c (σ). **(a)**: Reconstructed Z position of alpha decays originating from the cathode grid which drifted in TPC 1. **(b)**: The same for TPC 2. **(c)** and **(d)**: Difference of U and V coordinate of charge depositions that were split by the cathode grid and ended up in both TPCs. Since the U -wires in one TPC are parallel to the V -wires in the opposite TPC, the U coordinate in one TPC should be the same as the V coordinate in the other TPC for such events. The satellite peaks at ± 4.5 mm are from 2- V -wire events.

to a position error in U and V of 2.6 mm. Since the transformation to X, Y coordinates involves the sum of U and V , the uncertainties in this coordinate system require an additional factor of $\sqrt{2}$. Because the bundle coordinates are amplitude weighted averages, the true position reconstruction uncertainties are expected to be smaller than in this simple model. A study using the detector Monte Carlo simulation yielded a standard deviation of 2.4 mm in U direction, 1.2 mm in V direction and 2.6 mm in the composite X coordinate from the true simulated position. To cross-check these results in real data, charge depositions that were split by the TPC's central cathode grid are used. Since the cathode grid is at the voltage maximum, charge depositions at the center of the TPC are split and the fractions drift to opposite anode planes. Yet they are known to originate from the same position and hence can be used to test the position reconstruction across the two TPC halves. The U -wires of one TPC half are parallel to the V -wires of the opposite half, so the U coordinate of such events in one TPC should have the same value as the V coordinate in the opposite TPC. The bottom plots of figure 4.6 show the distributions of the difference of these coordinates. The distributions have a standard deviation of ≈ 3 mm, consistent with the result from the Monte Carlo simulation, considering that this measurement is a measurement of a sum of U and V . It can be seen that only in rare cases the two coordinates differ by at least one channel pitch (9 mm). The distribution of $V_{\text{TPC1}} - U_{\text{TPC2}}$ is off-centered by ≈ 1 mm, which might indicate a slight parallel displacement between the V -wires in TPC 1 and the U -wires in TPC 2, which if true, would not contribute to the position reconstruction uncertainty.

4.7. 3D Reconstruction efficiencies

In order to obtain the full 3D position information of a charge cluster, all three types of its signals have to be reconstructed. Since the scintillation reconstruction has the highest energy threshold, it dominates the position reconstruction efficiency of single site clusters, as can be seen in figure 4.7 a. Multi site events usually share the scintillation signal, and so for charge clusters in a multi site event, the position reconstruction efficiency is a mixture of the V signal reconstruction efficiency of the individual cluster and the scintillation reconstruction efficiency of the sum of all clusters. For this reason, the position reconstruction efficiency for individual multi site charge clusters can have a lower energy threshold than the one for single site clusters. The position reconstruction efficiency of charge clusters in a multi site event is shown in figure 4.7 b. For the subsequent analysis, the efficiencies for individual clusters are less interesting than the efficiency for having an event in which all charge clusters have fully reconstructed position information. For single site events, this is just the same as the corresponding individual cluster efficiency. For multi site events, it is the product of the individual cluster efficiencies, and therefore depends on the structure of these events. Figure 4.7 c shows the efficiency of having a multi site event in which all charge clusters have full position reconstruction. It can be seen, that this efficiency depends on the type of calibration source used for the study. The general trend downward is explained by the fact, that higher energy events usually create more charge clusters, which enter in the efficiency product.

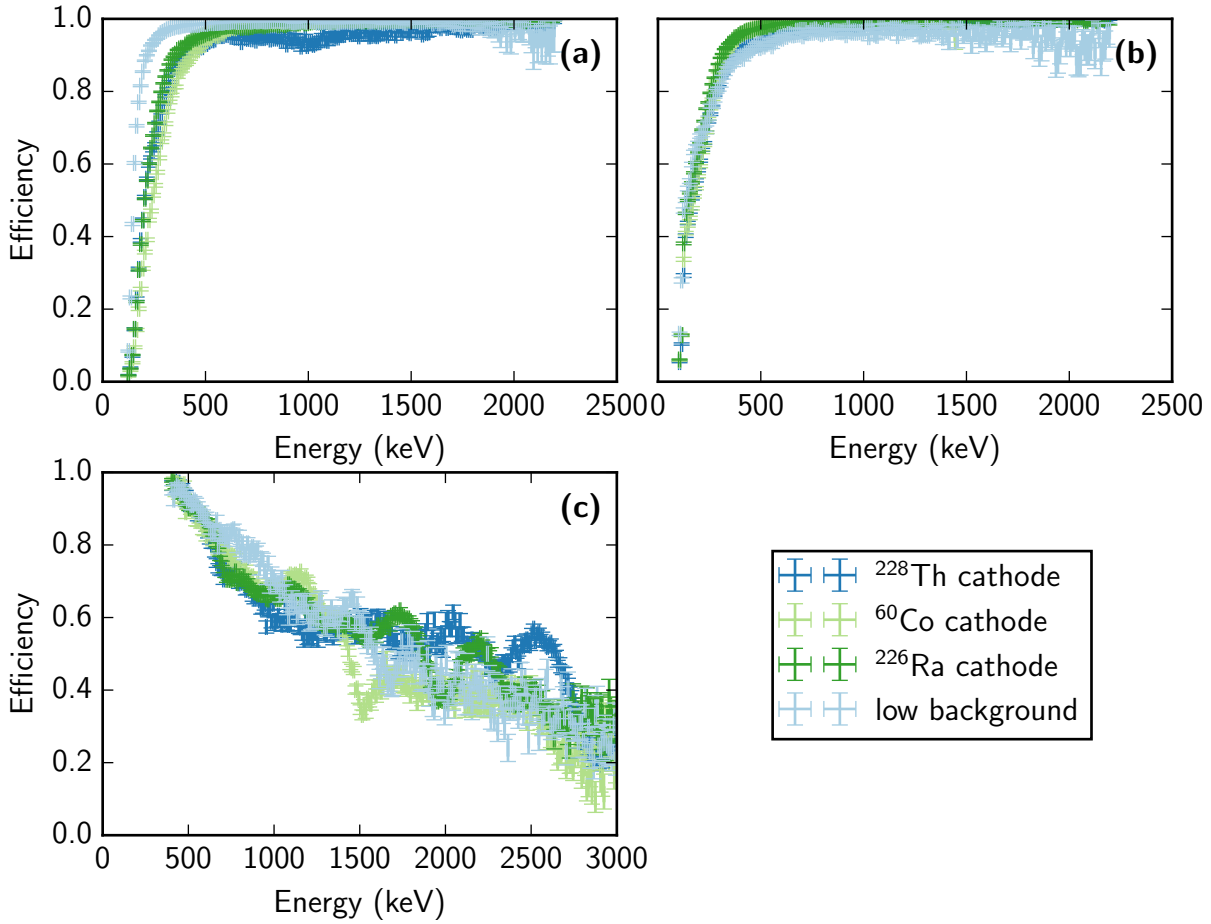


Figure 4.7.: Single-cluster 3D reconstruction efficiency, defined as the number of charge clusters with full 3D reconstruction divided by the total number of charge clusters. In **(a)** for single site events and in **(b)** for multi site events. **(c)** shows the probability that *all* charge clusters in a multi site event have full 3D position reconstruction vs the summed charge cluster energies.

5. Continuous multiplicity description

The main analysis presented in this text makes use of the event multiplicity to gain discrimination power between gamma backgrounds and beta decay signals. For this purpose a binary distinction in single site and multi site events is chosen, which was defined in section 4.5. However there is motivation for a finer graduation of the event multiplicity. ^{60}Co for example is an important background in the experiment and is also used as a calibration source. This nuclide undergoes beta decay to ^{60}Ni with the emission of two prompt gammas. The energies of the two gammas are 1173.2 keV and 1332.5 keV, with the lifetime of the intermediate state being only 0.7 ps. Because of their coincidence, the summed energy of both gammas can appear in the detector, which is less than 50 keV away from the $0\nu\beta\beta$ Q-value and about 100 keV from the ^{208}Tl gamma line in the ^{232}Th background. Since a sum peak event consists of two gammas, its multiplicity or number of energy depositions in the detector is higher than that for a single gamma event, which allows it to be discriminated from other backgrounds or beta decay events. On the low end of the multiplicity spectrum, background discrimination might also profit from a finer graduation. Currently charge depositions ending up on one or two neighbouring wire channels are categorized as single site. Due to Compton scattering however, charge depositions from gamma radiation might systematically end up more often on two neighbouring wire channels than on one compared to charge from beta radiation, a fact that the current binary multiplicity can not exploit.

Various multiplicity metrics, have been considered with the focus on discrimination power between single and multiple gamma events [68]. One metric that gave particular good discrimination between these two types of events is the sum of edge weights of the minimum spanning tree (MST) of the reconstructed charge clusters in the event. A MST is a concept in graph theory. A graph is a set of vertices, in this case the charge deposits in the detector, and edges which connect the vertices. A spanning tree of a graph is a subgraph of that graph that contains no loops and connects all vertices of the graph. If the edges contain weights, one can find a minimum spanning tree, which is a spanning tree whose sum of edge weights is less than or equal to the sum of edge weights of all other spanning trees of the graph. For exact definitions see for example [69]. Figure 5.1 illustrates a graph with weighted edges and its minimum spanning tree.

The edge weights were defined as the euclidean distance between charge clusters scaled by an energy dependent factor, chosen to give the best discrimination power between single- and multiple-gamma events. However such an additional energy dependent factor turns out to yield a too complicated metric to give reasonable agreement between simulation and actual data. Instead just the euclidean distance between charge depositions is used as edge weights. This simplifies the metric, while still retaining good discrimination power between single and multiple gamma events as will be shown below. The metric

5. Continuous multiplicity description




Figure 5.1.: Example of a minimum spanning tree. The figure shows a graph with weighted edges. Vertices are drawn as black dots, edges are shown as dashed black lines or full red lines. The spanning tree of this graph with the lowest sum of edge weights is marked by the red lines.

is a natural description of the event's spatial extent, and its dimension is a length. A distinction between one-wire charge depositions and two-wire charge depositions is possible, if not just the centers of each charge cluster (as defined in section 4.4) are used as vertices of the MST, but each individual U-wire signal within a charge cluster. The U coordinates of the vertices are the U coordinates of the corresponding channel, and the V coordinate is inherited from the signal's parent charge cluster. This vertex definition gives the MST metric a resolution on the single channel level while enabling it also to reflect differences between single, few, and many photon scatters. Figure 5.2(a) shows the energy spectrum of the ^{60}Co calibration source located in the cathode plane versus the described MST metric. While the single gamma full absorption peaks at 1173.2 keV and 1332.5 keV are mostly contained within MST values of less than ≈ 100 mm, combinations of two gammas like the Compton area above the full absorption peaks and the sum peak lie mostly above this value. This could be exploited by supplementing the single site and multi site categories with a third category of events whose MST metric is larger than a certain threshold. Figure 5.2(b) illustrates this on the example of the ^{60}Co calibration source. Cutting events with an MST metric larger than 110 mm retains most of the single gamma events, while reducing the part of the spectrum originating from two gammas by approximately a factor of 10.

Most of its information is however used if the MST metric is not used to define a third multiplicity category, but if it is treated as a continuous variable in the fit instead, superseding the binary classification into single site and multi site. An essential requirement that needs to be fulfilled though, is the agreement between simulation and data on the event distribution of the MST metric, so that the event distribution functions in the MST metric produced by the simulation correctly predict the MST metric distribution in data. With the current simulation package however, the agreement is rather bad, as




Figure 5.2.: The top plot shows the distribution of the MST metric versus total event energy for the ^{60}Co calibration source. The lower plot compares the energy spectrum of all events to that with an MST metric of less than 100 mm. Most of the single gamma events at or below the two single gamma full absorption peaks are contained within low values of the MST metric, while two gamma events above the 1332.5 keV peak are systematically at higher values of the MST metric.




Figure 5.3.: Agreement between simulation and data for the ^{228}Th calibration source in the cathode plane. The top plot shows the full range of the MST metric. The bottom plot shows the distributions at low values of the metric. Data is shown in black, the distribution obtained from the default simulation package is shown in blue. The bad modelling of the MST distribution at low values by the simulation is due to the lack of simulated charge carrier diffusion. Adding this feature to the simulation improves the agreement with data drastically as shown by the red curve.

can be seen in figure 5.3, which compares the MST distributions between a simulation (blue) of the thorium calibration source in the cathode plane and corresponding calibration data (black). The most significant disagreement happens for MST values at 0 mm and for values around 9 mm. The first case are events that produce only one charge deposition signal on the U-wires and therefore have only one vertex in the MST. Since the MST metric is defined as the sum of edge lengths, its value is zero in case of only one vertex. The second case is from events with two coincident charge deposition signals on neighbouring wire channels. In this case there are two vertices whose distance is the width of one channel, which is 9 mm. The simulation overpredicts the one wire case by approximately the same amount it underpredicts the two wire case.

The main reason for this disagreement is, that the diffusion during the drift of the charge cloud towards the anode is not modeled at all in the default simulation. Instead, charge depositions are treated as point like objects during the drift simulation and signal generation. Ionization electron diffusion in liquid xenon depends on the strength of the electric field and is typically much stronger in the plane transverse to the drift direction than the diffusion in the drift direction. Measurements of the transverse diffusion coefficients exist for slightly higher fields than the one used in EXO-200 [47]. An extrapolation of these measurements to the field used in EXO-200 predicts a transverse diffusion coefficient on the order of $100 \text{ cm}^2/\text{s}$, which corresponds to a diffusion length of $\approx 1 \text{ mm}$ over the full drift length. This can cause the charge cloud of a single interaction point to end up on more than one U-wire, whose channel spacing is 9 mm.

To examine the effects of charge carrier diffusion in the TPC, the simulation package has been extended with a diffusion model. The simulated charge depositions, which were assumed point-like in the default version, are split into a fixed number of sub-depositions with the total amount of charge spread evenly among them. The sub-depositions each independently undergo a random walk in the plane perpendicular to the drift direction in addition to their drift movement. The number of sub-depositions into which the original deposition is split can be set for the simulation, and already a quite small number of 10 sub-depositions seems to be sufficient to model the effects that determine the distribution of the MST metric at low values. To estimate the effective transverse diffusion coefficient for EXO-200, the calibration sources were simulated with different transverse diffusion coefficients, ranging from $60 \text{ cm}^2/\text{s}$ to $100 \text{ cm}^2/\text{s}$. A comparison of these simulations to data from the ^{228}Th calibration source is shown in the appendix B. A transverse diffusion coefficient of $80 \text{ cm}^2/\text{s}$ gives reasonable agreement between simulation and data in the MST metric. The red curves in figure 5.3 show the distribution of the MST metric for a simulation of the ^{228}Th calibration source with this diffusion coefficient. Clearly, the incorporation of charge carrier diffusion into the simulation eliminates most of the discrepancies between simulation and data.

The MST metric is sensitive to the different relative fractions of events that end up on one U-wire compared to two neighbouring wires. This adds additional beta / gamma discrimination power that the default SS / MS multiplicity definition is insensitive to. The same is true for the fall off of the MST metric at large values. The MST distribution of beta decays falls off very steeply, since fast electrons deposit their energy within a very small volume except for the rather rare cases of Bremsstrahlung. Gamma events are

5. Continuous multiplicity description

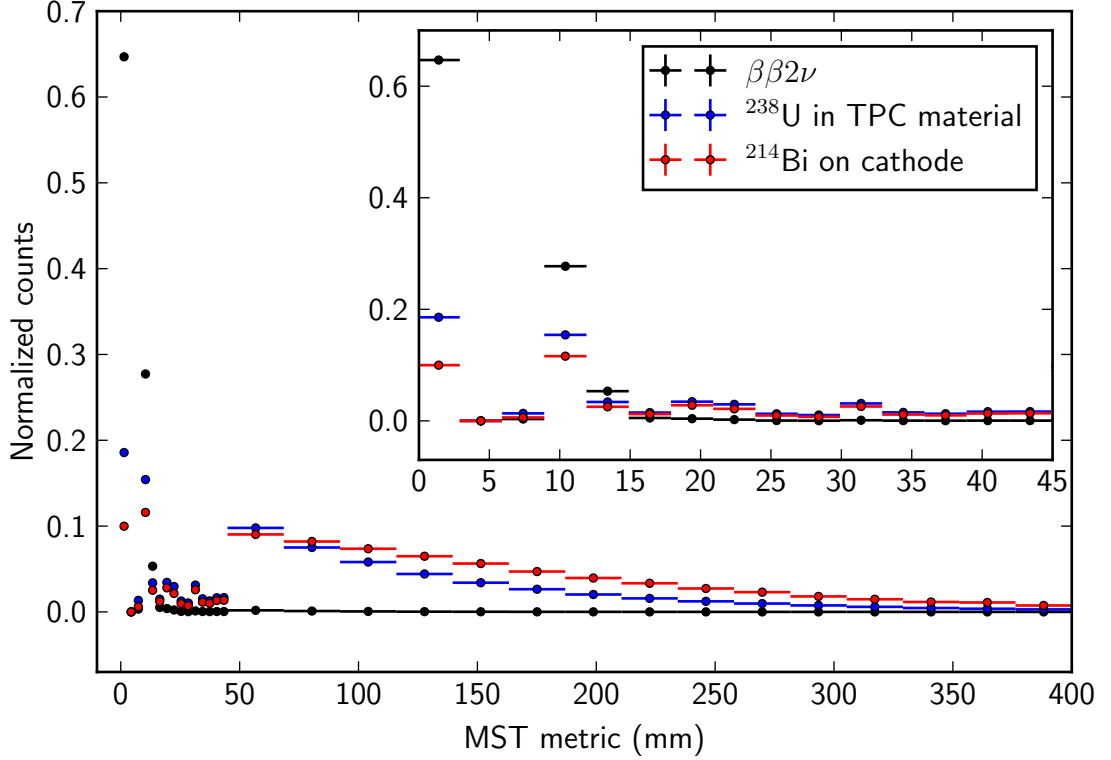


Figure 5.4.: Comparison of three different background PDFs projected onto the MST metric dimension. The binning is a section-wise uniform binning, with a bin width of 3 mm for values of MST metric < 45 mm and a bin width of 23.7 mm for greater values. The inset shows a zoom into low values of the MST metric.

expected to fall off much slower than beta decay events and coincident gamma events systematically occupy very large MST values. To make use of this information in a fit of various background and signal components to the data taken by the detector, probability density functions (PDFs) describing the distribution of the MST metric for each background component must be created using the simulation package. The process of building these PDFs will be detailed in section 8.2.1. They are based on histograms in the corresponding observable, for example energy or in this case the MST metric, of simulations of various backgrounds. The finer the resolution of the PDFs, i.e. the finer the binning of the underlying histograms, the more data must be simulated. To keep the amount of data that needs to be simulated at a maintainable level, an efficient binning must be chosen. The PDFs should be detailed at low values of the MST metric to reflect the features of the discrete anode wire spacing. For larger MST values it is sufficient to model the general trend and detailed resolution is unnecessary. This is reached by choosing a fine grained bin width of 3 mm for MST values of less than 45 mm, and 23.7 mm for larger values. Figure 5.4 compares the MST PDFs of three different background components whose behaviour is in agreement with above discussion. The one-wire to two neighbouring wire ratio is higher for double beta decay events than for

the gamma backgrounds of the ^{238}U and ^{222}Rn chains (the latter represented by ^{214}Bi). Additionally the double beta decay PDF falls off rapidly for MST values greater than 10 mm in contrast to the gamma backgrounds.

The analysis presented in the following chapters is based on the binary single site / multi site multiplicity categories. In section 8.4, the fit of the background model to the low background data is repeated with the single site / multi site categorization replaced by the MST metric and the associated gain in sensitivity to neutrinoless double beta decay is discussed.

6. Data processing and correction

The data used in this analysis was taken between 22 September 2011 and 1 September 2013, which is a period of 710 days. Some fraction of this time is covered by source calibration runs and periods at which the electronegative purity in the detector was low. After subtraction of these periods, the amount of low background data used for the analysis corresponds to a total live time of 477.6 days.

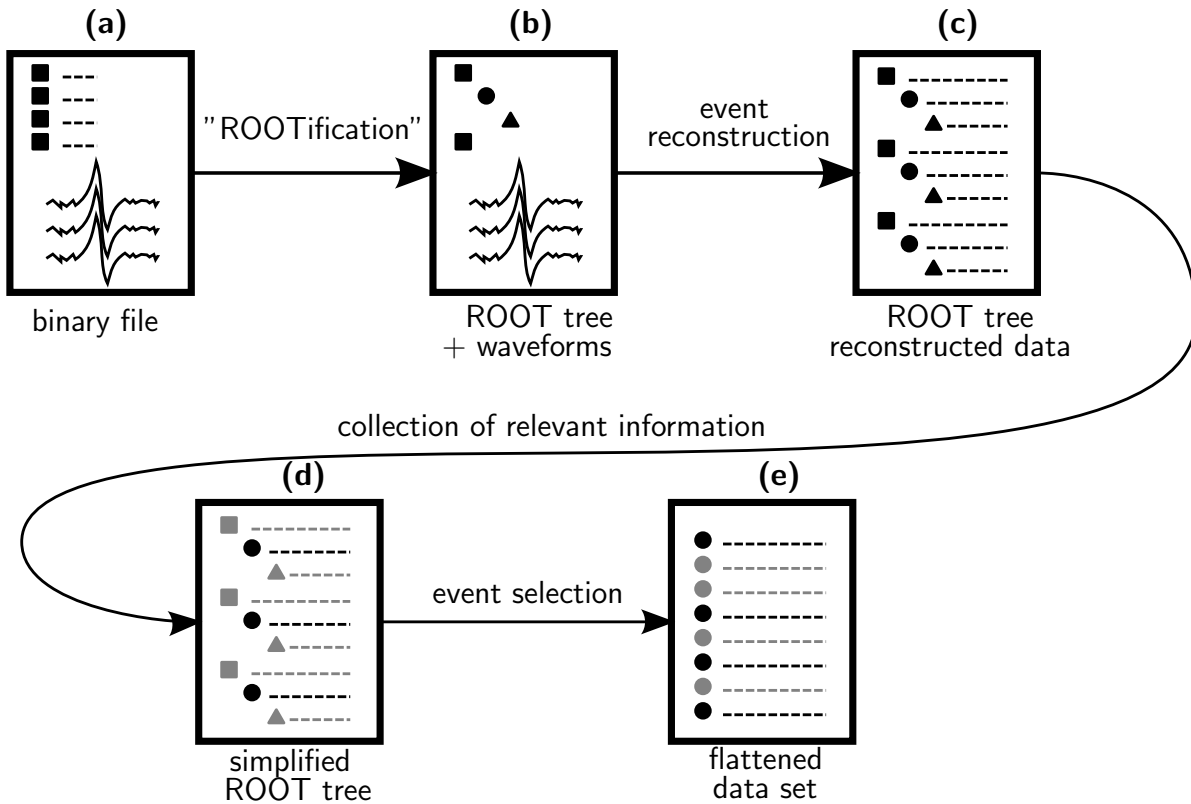


Figure 6.1.: Schematic of data processing. **(a)**: The DAQ records waveforms and trigger information to a binary format. **(b)**: The data is converted to a ROOT tree format which is empty except for waveforms. **(c)**: The reconstruction process fills the data structure. Waveforms are dropped. **(d)**: Data is converted to a lighter ROOT tree. Information not relevant to high level analysis is dropped. **(e)**: Event selection cuts are applied and the data set is reduced to contain only the data that is used in the fit.

All data is processed according to the scheme shown in figure 6.1. The processing happens in several steps and the output of each step is stored for possible future repro-

cessing. This has the advantage that only upstream processes need to be redone if one of the processes was changed. The raw waveforms and trigger information are provided by the DAQ in a custom binary format. In a first processing step, this data is converted to a ROOT [70] tree format which contains the waveforms and an empty data structure that can hold all possible information on the event.

This data structure is filled in the reconstruction process which was discussed in chapter 4. The data structure is hierarchical and for example allows bidirectional relations between all signals that were found in an event and higher level objects such as charge clusters and their reconstructed position. In the same step, the raw signal amplitudes are corrected in order to obtain an optimal energy resolution, because the response to an energy deposit is not completely uniform throughout the detector. These corrections are described in section 6.1, except for the correction of different wire gains, which was discussed in section 4.2. The waveforms are dropped at this stage to keep the file sizes at an acceptable level.

In the next step, the data is reduced to a lighter format containing only the information relevant for higher level analysis, while still keeping all of the reconstructed events. Single signal amplitudes are dropped at this stage and instead the calibration function is applied which translates the charge and light amplitudes to an estimate of the total energy in the event. The calibration procedure is explained in section 6.2.

Event selection cuts are applied in the final step, which produces the data sets used in the fit. Only the total energy, the standoff distance and the time of the event are saved, and the data set is split into a single site and a multi site part. The event selection criteria are detailed in chapter 7.

6.1. Signal corrections

6.1.1. Electronegative impurity

As discussed in section 3.2.1, the ionization electrons may scatter on electronegative impurities on their way to the anodes, which leads to a loss of electrons in the charge cloud following equation 3.8. The result is an exponential decrease of free electrons with drift time:

$$N(t) = N_0 \exp(-t/\tau_e), \quad (6.1)$$

where N_0 is the initial number of electrons created in the ionization process, and τ_e is the electron lifetime. This effect introduces a Z dependent energy response to the detector. Charge clusters near the cathode that need to cover a larger drift distance, yield smaller ionization signals per unit energy than those near the anodes. For this reason the electron lifetime is monitored periodically using the ^{228}Th calibration source, and the data is corrected for the expected attenuation. To measure the electron lifetime, the calibration source is placed at the cathode and the ionization yield from a known energy is measured in different parts of the detector. For this purpose, each TPC half is divided into 16 equal drift time bins, and a Gaussian plus error function model is fit to the full absorption peak of the source's ^{208}Tl line. The central value of the peak and the




Figure 6.2.: Measurement of electron lifetime. Each TPC half is divided into 16 drift time bins and the amplitude of the ^{228}Th source's full absorption peak is plotted for each bin. The data is fitted by equation 6.1 for both TPCs. Figure adapted from [13].

6. Data processing and correction

mean drift time in each bin are plotted as shown in figure 6.2 and fit by equation 6.1 to extract the electron lifetime for both TPCs.

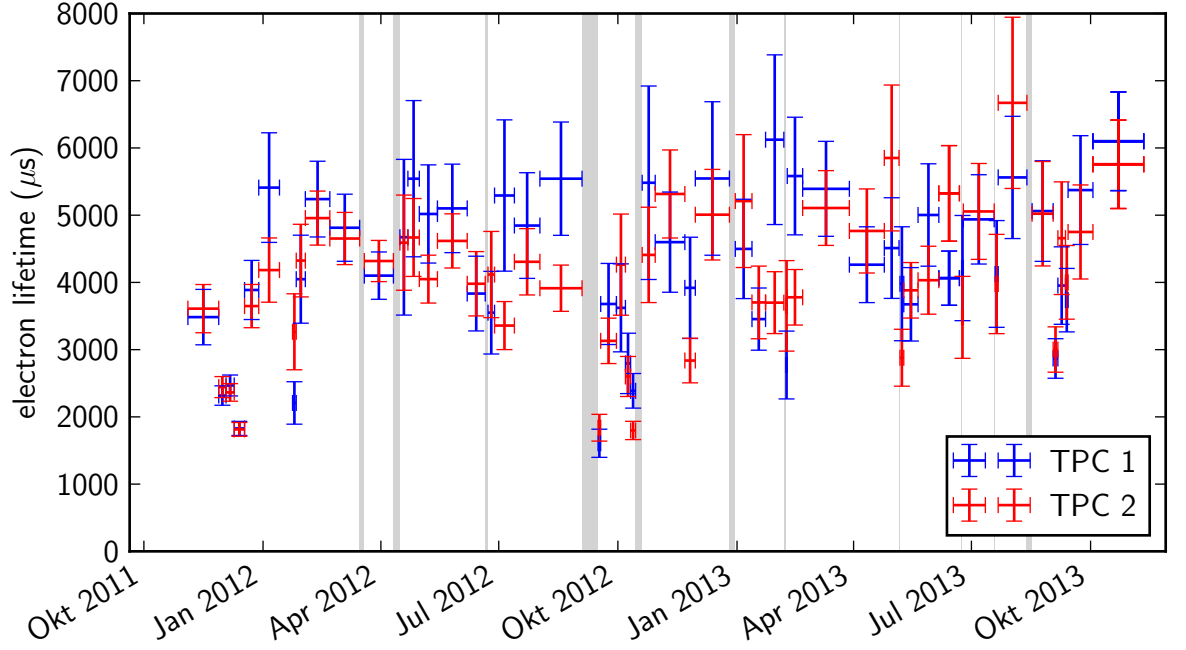


Figure 6.3.: Electron lifetime between September 2011 and November 2013. The shaded regions mark periods of bad purity, coincident with interruptions in the xenon recirculation or gas feed events. The shaded regions are not used in the analysis. Figure adapted from [71].

Although the xenon purification system described in section 3.5.2 works quite well, interruptions in the xenon circulation due to power outages or maintenances and gas feed events cause the electron lifetime to vary over time. For this reason the electron lifetime is measured several times per week with the method described above. The time variation of the lifetime for the whole data set is shown in figure 6.3. These values are used to correct the amplitudes of ionization signals in the data by a factor $\exp(t/\tau_e)$. For lifetimes $> 1000 \mu\text{s}$, the maximum correction is less than 12 %.

6.1.2. Shielding grid inefficiency

# U-signals	p_0	p_1 (mm)	p_2 (mm ⁻¹)
1	0.041	5.84	$4.16 \cdot 10^{-6}$
2	0.079	6.64	$2.11 \cdot 10^{-6}$
more	0.041	5.84	$4.16 \cdot 10^{-6}$

Table 6.1.: The parameters of equation 6.2, used to correct for the inefficiency of the shielding grid.




Figure 6.4.: Determination of the shielding grid's inefficiency using simulation. The plots show the reconstructed sum of U-signal amplitudes divided by the true simulated energy versus the Z position of the energy deposit for events with one U-signal **(a)** and two neighbouring U-signals **(b)**. The red curves show the correction model (equation 6.2) times a normalization constant.

The secondary purpose of the V-wire plane is to shield the U-wires electrostatically from the charge in the drift volume. Since the V-wires do not form a perfect conduction plane, the charge of the slowly drifting ions creates a small dependence of the U-wire signal amplitude on the Z position of the charge deposit (e.g. [72]). We model the Z dependence of the U-wire pulse amplitudes by the following empirical formula

$$A_{\text{meas}}(Z) = A_0 \cdot \frac{1 + p_2 Z}{1 + p_0 \exp\left(\frac{|Z| - Z_{\text{max}}}{p_1}\right)}, \quad (6.2)$$

where A_{meas} is the measured amplitude, A_0 is the amplitude one would obtain with a perfect shielding grid, Z_{max} is the maximum drift distance (192.5 mm) and p_0 , p_1 and p_2 are empirical parameters. We use the detector Monte Carlo simulation to estimate the parameters of this equation. The simulation uses electric fields and weighting potentials of the detector, calculated using Maxwell [65], to generate the pulses of U-wire signals. The reconstructed pulse amplitudes are then compared to the true simulated energy as in figure 6.4, and the parameters p_0 , p_1 and p_2 are extracted by fitting equation 6.2 (times a normalization constant) to the data. The resulting parameters, together with the inverse of equation 6.2 are used to correct the data for this effect. It was found, that the inefficiency behaves slightly different for charge clusters that deposit their charge on two neighbouring U-wires. We therefore treat this case separately. The parameters used for the signal amplitude correction are listed in table 6.1.

6.1.3. Nonuniform light response

Due to the good uniformity of the electric field in the detector, the amplitude of charge pulses of a certain energy deposit depend only on the Z position and the dependencies can be modeled quite well. The situation is different for light signals. A scintillation flash can be reflected several times, for example by the Teflon cylinder, or can be shaded by the cathode grid, before it hits an APD. Furthermore, scintillation signals from different locations in the detector will shine brighter onto different APDs, depending on distance and angle. This, in combination with the fact that the APDs vary substantially in gain and noise level, makes it harder to correct the light signal in order to achieve a uniform energy estimate. In addition, the APD signals suffer from noise introduced by the front end electronics, which is correlated across different channels. A substantial effort has been undertaken to develop an optimal energy estimator from the light signals under these conditions in [73], which will be described briefly in the following.

During the light signal reconstruction process, discussed in chapter 4, the summed waveforms over all APD channels in one plane are used to find scintillation signals. Fitting a pulse template to the summed waveform to extract the signal amplitude, and using this amplitude as an energy estimator like in the case for charge signals however would lead to sub-optimal energy resolution. Because the noise on APD channels is correlated, and a light flash at one position preferentially creates a large signal on a few channels and less to no signal on the rest of the channels, summing up the waveforms of all channels actually dilutes the signal to noise ratio. On the other hand the pure random (uncorrelated) part of the noise improves with the square root of the number of summed channels. The method described in chapter 4 of [73] makes use of the detailed knowledge of the correlated noise on each channel, the expected signal shape and the light response on each channel to a scintillation signal at a certain time and position in the detector to derive an optimal energy estimator from the APD waveforms in each event. The knowledge of the expected signal shape is used to weight more heavily those frequency components on the waveforms that give larger signal to noise ratio. The knowledge of the light response of each single APD channel to a scintillation event at a certain position in the detector is used to weigh more heavily those channels which are expected to give a large signal due to their proximity to the interaction point. And finally the knowledge of the noise correlations across channels can provide a better estimate of the noise on one channel than an estimate on only the individual channel can provide, and those channels with higher signal to noise ratios are given larger weights. The method uses all this information to calculate an optimal energy estimate that is linear in the waveform samples and yields a uniform energy response throughout the detector.

6.2. Energy calibration

6.2.1. Combination of charge and light

For a given particle energy, the charge and light yield in a liquid xenon detector are anticorrelated [74, 75]. This means that the energy resolution in a single of the two




Figure 6.5.: Combination of charge and light signal. The figure shows the calibrated energy estimates obtained from light signals vs the one obtained from charge signals for data from the ^{228}Th calibration source. The gamma lines of the source are visible as diagonal islands, which demonstrate the anticorrelation between charge and light in the detector. A measurement of the anticorrelation parameter makes it possible to form a combined energy estimate (diagonal axis) with improved energy resolution. Projections of this data on the three axes are shown in figure 6.6.




Figure 6.6.: Scintillation only, charge only, and combined energy spectrum of the ^{228}Th calibration source. The plots show projections of the data in figure 6.5 on the three axes. It is apparent that the charge channel yields better energy resolution than the light channel. The best resolution is achieved by a combination of the charge and light signals.

channels is smeared out. If however both channels are measured, they can be combined to give an energy estimate with improved resolution. This can be seen in figure 6.5, which shows the energy estimated using the light signal versus the energy estimated using the charge signal of the ^{228}Th calibration source. It can be seen that the gamma lines in the source's spectrum are tilted by a certain angle towards the axes of the plot. This means that the energy resolution can be improved by using an appropriate combination of the charge and light amplitude, as indicated by the third (diagonal) axis in the plot. The energy estimate we use is a linear combination of the corrected charge and light amplitudes

$$C = \cos(\theta(t)) A_C + \sin(\theta(t)) A_S, \quad (6.3)$$

where A_C and A_S are the completely corrected charge and light signal amplitudes respectively that were discussed in the previous section, and the parameter θ is defined as the one that gives best energy resolution for the 2615 keV peak of the ^{228}Th source. θ is determined weekly by minimizing the width of the 2615 keV peak of the calibration source and therefore labeled as time dependent. Note that C is an uncalibrated magnitude, and the second degree of freedom in the linear combination of A_C and A_S is

absorbed in the calibration. The combination yields an energy resolution at the source's 2615 keV peak of 1.4 %, which is more than a factor of 2 better than the energy resolution of the charge only channel. This is demonstrated in figure 6.6 which shows a projection of the data in figure 6.5 onto the charge only, light only, and combined energy axes. The 2615 keV peak is much sharper in the combined energy estimate, and some lower energy features become visible, which are washed out in the charge- and light-only signals.

6.2.2. Calibration procedure

To calibrate the combined amplitude C , several photopeaks of different calibration sources are used. Full scale calibration campaigns involving the cesium, cobalt and thorium sources are done several times a year to provide the overall scale of the calibration and the ^{228}Th source is deployed a couple of times per week to track any time dependencies. The positions of the calibration photopeaks C_{Th} , C_{Co}^1 , C_{Co}^2 and C_{Cs} are measured by fitting a Gaussian plus error function model to the spectra. Examination of the data within a calibration campaign yields that, while the absolute values of C_i show some significant time variation, the ratios between the amplitudes remain quite stable. Our approach thus expresses possible nonlinearities of the energy scale S in terms of amplitudes relative to the amplitude of the ^{228}Th source's 2615 keV peak, which is measured on a weekly basis:

$$x := \frac{C}{C_{Th}(t)} \quad (6.4)$$

$$S(C) = S(x(C)) = p_2 x^2 + p_1 x + p_0. \quad (6.5)$$

This method allows us to have constant parameters p_i and model the time dependence of the energy scale by the time dependence of the ^{228}Th photopeak. The parameters p_i are determined by fitting $S(x)$ to a plot of the measured photopeaks x_{Th} , x_{Co}^1 , x_{Co}^2 and x_{Cs} versus their expected values.

Applying the the Gaussian plus error function model to extract the position of the photopeak to Monte Carlo simulations of the calibration source, shows that the extracted positions have a small bias relative to the true energy of the simulated gamma particle. This bias is due to the incomplete modeling of the photopeak and Compton shoulder by a Gaussian plus an error function. The biases are comparable for all sources and we use the one estimated from a simulation of the ^{228}Th source, $E_{\text{bias}} = 2.9 \text{ keV}$ (6.1 keV) for SS (MS) to correct for this fact.

The complete calibration function is

$$E = (E_{Th} - E_{\text{bias}}) \cdot S(x) = (E_{Th} - E_{\text{bias}}) \cdot (x^2 p_2 + x p_1 + p_0) \quad (6.6)$$

$$x = \frac{C}{C_{Th}(t)}, \quad (6.7)$$

where E_{Th} is the true energy of the ^{228}Th source's full absorption peak ($E_{Th} = 2615 \text{ keV}$). E_{bias} , p_0 , p_1 and p_2 are time independent calibration parameters as discussed above, and $C_{Th}(t)$ is the time dependent parameter, which represents the magnitude of the 2615 keV

peak in units of the combined amplitude C . The calibration is done independently for single site and multi site data.

6.2.3. Energy resolution




Figure 6.7.: The energy resolution of the detector. The top panel shows the absolute energy resolution ΔE in keV, the bottom plot shows the relative energy resolution $\Delta E/E$ in percent. 1σ confidence regions are shown as bands around the resolution curves. Figure adapted from [76].

As discussed in chapter 3, good energy resolution is a requirement to reach high sensitivity to the signature of neutrinoless double beta decay. In addition, a precise knowledge of the detector's energy resolution is necessary to construct the various spectral distribution functions of backgrounds as seen by the detector. We assume that the detector's response function follows the form of a Gaussian whose width, the energy resolution σ , depends on the amount energy deposited in the detector. We use a semi empirical model for the energy dependence of the resolution, following the arguments in [77]:

$$\sigma^2(E) = p_0^2 E + p_1^2 + p_2^2 E^2 \quad (6.8)$$

The coefficient of the energy independent term p_1 can be interpreted as the basic electronic noise level of the detector readout system. The coefficient p_0 can be addressed to statistical fluctuations in the number of quanta created in the ionization and scintillation process. As discussed in section 3.2.1, these processes are approximately of poissonian nature, and the variance due to these fluctuations thus should be proportional to the deposited energy and poses a fundamental limit on the achievable energy resolution. The contribution of the p_2 term to the energy resolution depends linearly on the deposited energy. Various effects such as drifting gains or other inhomogeneities in the detector's energy response in space and time can contribute to this term.

The energy resolution shows a slight variation in time during the analysis period. However we determine average parameters for p_0 , p_1 and p_2 to construct the spectral distribution functions used in the analysis of low background data and include the time variations as systematic uncertainties. The whole set of calibration data from the cobalt and thorium sources is used to determine average values and covariances of the three energy resolution parameters for the analysis period. The parameters are extracted by convoluting the resolution model of equation 6.8 with exact energy spectra obtained from Monte Carlo simulations of the two sources and fitting the resulting spectra to the calibration source data. The fit weights the calibration data periods by the fraction of low background data taken during that period and treats single site and multiple site data and parameters independently. The variances determined in this way represent the statistical uncertainties on the six parameters (three for single site and multi site each). The variation due to the time dependence of the parameters is estimated by individually extracting parameters for the 7 calibration campaigns within the analysis period. The resulting parameters are listed in table 6.2 and the resolution curves are shown in figure 6.7.

	p_0 ($\sqrt{\text{keV}}$)	p_1 (keV)	p_2
SS	0.628	20.8	0.0011
MS	0.602	25.8	0.0040

Table 6.2.: Average energy resolution parameters (equation 6.8) determined by fits to all available cobalt and thorium calibration source data.

Since the summed energy of the two electrons emitted in neutrinoless double beta decay is a discrete value, it is suitable to define a region of interest (ROI) around the decay's Q-value whose size depends on the energy resolution at the Q-value. We define the $1\text{-}\sigma$ ($2\text{-}\sigma$) ROI, which corresponds to the energy interval that extends ± 1 (± 2) times the energy resolution around the Q-value. Their values are shown in table 6.3.

ROI	energy range (keV)
1- σ	[2416.6, 2491.6]
2- σ	[2379.0, 2529.2]

Table 6.3.: The definition of the energy region of interest used in this analysis.

7. Data selection

Not all of the low background data that was recorded is actually used in the analysis. Chunks of data that are obviously bad or events that are definitely not signal-like are not allowed to contribute to the final measured energy spectrum. Since such event selection cuts affect the signal efficiency and detector live-time, it is necessary to study their effects in detail, which is the purpose of this chapter.

The selection of low background data that is used in the analysis is defined by two mechanisms. In a first step, broader selection criteria are applied that remove chunks of data due to bad environmental conditions. These selection criteria define the actual live-time of the detector, which is less than the low background data taking time of 511.6d. In a second step, data within the live-time of the detector are selected based on event-by-event selection criteria. Events that are clearly background related, noise, or originate from less well modeled regions of the detector for example are not used in the analysis. These selection criteria or cuts, while not contributing to detector dead-time, affect the $0\nu\beta\beta$ signal efficiency and thus it is important to study the effect and associated uncertainty of each.

7.1. Detector live-time

The detector live-time differs from the low-background run time due to removal of data according to the following criteria:

Muon veto Data recorded between 1 ms before and 25 ms after a trigger of the muon veto system is cut from the data set. This cut aims to reduce the amount of background introduced by cosmogenic activation of the detector material or the HFE fluid. Many of the possible cosmogenic states are short lived. Analysis of the cut data for example yielded a neutron capture time in the hydrogen of the HFE fluid of $740(120)\mu\text{s}$. In addition, if a muon went through the TPC and was reconstructed as such by the reconstruction process, we cut 60 s of the data following the muon event. The main source of background from a muon traversing the TPC despite the muon itself is the capture of spallation neutrons on ^{136}Xe , which produces the beta-unstable isotope ^{137}Xe . This xenon isotope undergoes beta decay with a half-life of 3.8 min [78]. The 1 min cut has been chosen as a compromise between vetoing these decays and losing as little live-time as possible. Together, these two cuts reduce the live-time by approximately 5 %.

Poor data Data that was taken during periods of poor environmental circumstances, such as mining activities at the site, experimenters in the clean room module that

7. Data selection

contains the detector, or times at which the free electron lifetime was less than $1000\mu\text{s}$ are removed from the data set.

Run boundaries Data 60s after the start and 1second before the end of a low background data taking run are cut. This ensures that the 60s muon veto cut is applied, even if the muon event happened right before the start of the run and was therefore missed.

The periods which pass these criteria define the live-time of the detector. It is calculated by counting the number of forced triggers recorded by the DAQ, which are requested every 10s. For this analysis period, the total detector live-time is $477.60(1)\text{d}$. The amount of data removed due to the different vetoes discussed above is summarized in table 7.1.

	time (days)	%
vetoed time	34.0	6.6
environmental	8.9	1.7
TPC muon	21.9	4.3
panel muon	3.1	0.6
live time	477.6	93.4
total run time	511.6	100

Table 7.1.: Summary of detector run time and vetoes used in this analysis.

7.2. Fiducial volume

To convert a measured event rate into a half-life, the number of decaying atoms must be known within reasonable uncertainty. The detector volume on which the analysis is based is called the fiducial volume. Since EXO-200 is able to reconstruct the position of charge clusters, we can select arbitrary volumes inside the detector as the fiducial volume. The most obvious choice would be to use the whole sensitive region of the detector in order to maximize the target mass. However it is also important to keep systematic effects as low as possible, and so a volume is chosen that is well understood and properly modeled in simulation. A hexagonal right prism in each of the two TPC halves is chosen as the fiducial volume, because the wire planes have the form of a regular hexagon. For this choice, four of the volumes surface boundaries are parallel to the anode wires, which allows to express these boundaries in terms of the natural detector coordinates U and V . The two boundaries which are not parallel to the anode wires are expressed in terms of the derived coordinate X . Each wire plane consists of 38 channels. The size of the hexagon is chosen so that the fiducial volume encompasses all but the volume covered by the outer wires. This choice is made because the Monte Carlo simulation of the electron drift models the wires as an infinite repetition of infinitely long wires. Possible edge effects due to field lines bending in towards the outer wires are therefore not considered




Figure 7.1.: Projections of the fiducial volume onto the XY and XZ planes. Fiducial volume boundaries are shown as dashed lines, detector boundaries are shown as solid lines. The fiducial volume is bounded in the XY plane by the intersection of a hexagon with 162 mm and the cylindrical Teflon reflector walls, which have an inner radius of 183 mm.

in the simulation, but such effects might distort the actual volume that is covered by the wires, especially on the outermost ones. 3D simulations of the drift field suggest that this is indeed the case. Therefore the fiducial volume is restricted to the area covered by all wires except the outermost, which is a hexagon with an apothegm of 162 mm. Since this hexagon would extend slightly outside the Teflon reflector walls at its corners, it is cropped by a circle of 183 mm radius, which corresponds to the inner radius of the Teflon reflector wall.

The volumes are bounded in Z direction by the condition $10 \text{ mm} < |Z| < 182 \text{ mm}$, i.e. 10 mm in each direction from the cathode and 10 mm from the V-wire plane. This choice has been made because the cathode and anode grids emit a higher radioactive background than the surrounding liquid xenon as can be seen in figure 7.2(a). Since the contamination is mainly beta like as shown in figure 7.2(b), 10 mm of liquid xenon are enough to shield that background completely. Projections of the fiducial volume are shown in figure 7.1. This volume, multiplied by the liquid xenon density of 3.03 g/cm^3 corresponds to a ^{136}Xe mass of 76.5 kg, which multiplied by the live time gives a total exposure of 100 kg yr.

The uncertainty on the size of the fiducial volume is based on the reconstruction uncertainties for the four coordinates that define the fiducial volume as discussed in section 4.6, and on the difference on the event rate produced by calibration sources and the predicted event rate estimated using Monte Carlo simulations of the calibration




Figure 7.2.: **(a)** shows the Z distribution of events in the detector. It is apparent that the cathode ($Z = 0$ mm) and anode ($|Z| = 192$ mm) grids have a higher radioactive contamination than the liquid xenon in the bulk of the detector. **(b)** shows the energy spectra of events within 10 mm from the cathode (blue) and anode (red). Both are beta like.

sources. The uncertainties in the coordinates were found to be 2.4 mm in U , 1.2 mm in V , 2.6 mm in X , and 0.42 mm in Z . This results in an uncertainty on the chosen volume of 1.7%.

As mentioned above, the simulation models the outer parts of the detector less well. To take this into account, we compare the difference in event rates between calibration source data and simulations of the calibration source for different sizes of the hexagonal volume. Error-weighted averages of this number over different calibration sources and source positions are calculated for the chosen fiducial volume and the smaller volume used in [13]. A 1.7% difference is found between the two volumes. Conservatively, full correlation of this uncertainty due to detector edge effects with the basic fiducial volume uncertainty (also 1.7% as discussed above) is assumed. The total systematic uncertainty related to the choice of fiducial volume is thus $1.7\% + 1.7\% = 3.4\%$.

7.3. Event selection criteria

Several criteria must be fulfilled for an event to make it into the final data set. For example all positions in the event must be fully reconstructed. This reduces some systematic uncertainties but also reduces the signal efficiency. The following criteria must be fulfilled for each event:

Forced triggers As mentioned in section 3.4, the DAQ system records an event every




Figure 7.3.: The low background data set (black) after all cuts have been applied. The total scintillation amplitude vs the total (fully corrected) ionization amplitude per event is shown. Each event is represented by one dot in the plot. The events shown in red fail the light-to-charge ratio cut and are not included in the low background data set. Two island consisting of alpha decay events are visible in the upper left.

ten seconds, regardless of whether a signal that exceeded the trigger threshold is present or not. These forced triggers are used to measure the detector live time. Although we record the waveforms for these triggers, we cut events that lie within the $\pm 1024 \mu\text{s}$ frame of a forced trigger. The probability for a real decay to coincide with such a forced trigger is very small and reduces the signal efficiency by $\approx 1 \times 10^{-3} \%$.

Noise Single events that are obviously noise, as tagged by a simple algorithm before reconstruction, are neither reconstructed nor added to the data set. Such events occur with an average rate of $\approx 11 \text{ mHz}$. Comparing this to the average event rate of 75 mHz leads to an estimated random coincidence of a noise event with a real physics event in $9 \times 10^{-3} \%$ of the cases. In addition, a hand scan of events which were tagged as noise found no real signal on the investigated waveforms which leads to a conservative $6 \times 10^{-2} \%$ systematic error on the estimated signal efficiency due to this cut.

Event-to-event coincidence This cut aims to reduce backgrounds from fast decays in a chain. If two or more events occur within one second, both events are removed from the data set. A main target are is the bismuth polonium coincidence in the ^{222}Rn decay chain. The intermediate state in the decay $^{214}\text{Bi} \xrightarrow{\beta} ^{214}\text{Po} \xrightarrow{\alpha} ^{210}\text{Pb}$ has a half-life of $164.3\text{ }\mu\text{s}$ and therefore should be always covered by this cut if properly reconstructed. Random coincidences that happen within the 1 s cut are removed as well, which reduces the signal efficiency. The efficiency due to this cut has been calculated to $93.1(2)\%$.

Light-to-charge ratio Following the discussion of section 6.2.3, alpha radiation is expected to yield a higher scintillation to ionization ratio in the detector. This is used to effectively remove alpha decay events from the data set. Figure 7.3 shows the final data set and the cut on the light-to-charge ratio. The cut also removes another bismuth polonium coincidence. The decay $^{212}\text{Bi} \xrightarrow{\beta} ^{212}\text{Po} \xrightarrow{\alpha} ^{208}\text{Pb}$ is part of the ^{220}Rn decay chain and the intermediate state has a half-life of only 300 ns, which means that both decays end up in one event. The high light to charge ratio of the alpha decay however still makes the event identifiable as background. The impact on the $0\nu\beta\beta$ signal efficiency is estimated to be negligible.

One scintillation cluster To remove data from coincident decays that end up on the same event, we require that the event consists of only one scintillation cluster. In addition, we require that the scintillation cluster is reconstructed at least $120\text{ }\mu\text{s}$ before the end of the waveform, which allows all charges, which have a maximum possible drift time of $120\text{ }\mu\text{s}$, to be fully reconstructed. The rate at which the former requirement is applied mistakenly due to misreconstruction of noise as an additional scintillation cluster has been estimated by simulating pulses of known amplitude on top of noise traces. In 0.1% of the cases, the reconstruction process found an additional signal. Another study of this effect on the events in the low background data set that pass all other cuts but this one yielded that in $\approx 0.7\%$ of the cases a second signal was reconstructed although only one was really present.

Energy The total energy in the event must lie within the range of 980 keV to 9800 keV , where the trigger and signal reconstruction are known to be 100% efficient. The efficiency loss for a $0\nu\beta\beta$ signal due to this cut is estimated to be negligible.

Event positions The fiducial volume was discussed above in section 7.2. An event must be fully contained within the fiducial volume to be allowed into the data set. This is the case if all of the event's charge clusters lie within the fiducial volume. To ensure this is the case, one must add the additional requirement that all charge clusters have fully reconstructed 3D position information. The 3D reconstruction efficiencies for a single charge cluster have been discussed in section 4.7. The efficiency for an event however depend on its charge cluster composition and therefore cannot be directly tested using the gamma calibration sources, since they are more likely to create multiple charge clusters due to Compton scattering. As no beta emitting calibration source with an energy near the double beta decay Q-value

Source	Signal efficiency (%)	Error (%)
Forced trigger	99.99	-
Noise	100	< 0.06
1 s event coincidence	93.1	0.2
More than one scintillation cluster	100	< 0.7
No full position reconstruction	90.9	7.8
Fiducial volume	-	3.4
Light-to-charge ratio	100	0.15
Total	84.6	8.6

Table 7.2.: Summary of signal efficiencies and associated systematic errors due to the event selection cuts.

is available, the 3D reconstruction efficiency for $0\nu\beta\beta$ events is estimated using the Monte Carlo simulation. The efficiency is calculated as the fraction of events removed by the full 3D reconstruction requirement after all other analysis cuts discussed above have been applied, resulting in a value of 90.9 %. Since the Monte Carlo model is not perfect, an uncertainty must be assigned to this estimate. This is done by comparing the energy dependent 3D reconstruction efficiency curve for $2\nu\beta\beta$ as predicted by simulation to the one of background subtracted low background data. The difference between the two efficiencies calculated in this way are fairly energy independent and it's therefore justified and believed to be conservative to use the maximum deviation of the two efficiency curves (7.8 %) as an estimate of the efficiency's uncertainty.

After the application of all cuts, the total low background data set comprises 62764 events (50456 single site, 12308 multi site). The hit in efficiency and the associated error for each of the above cuts is summarized in table 7.2.

8. Search for neutrinoless double beta decay

The analysis of the $0\nu\beta\beta$ half-life is based on a binned maximum-likelihood fit on the data described in the previous chapter. The data is split between the single site and multi site categories and is simultaneously fit by a set of probability density functions (PDFs) that describe the population of different background sources as well as the $0\nu\beta\beta$ signal. Three observables are used in the fit:

Multiplicity This observable describes the event multiplicity which is related to the number of charge depositions in an event and has only two states, either single site (SS) or multi site (MS). An exact definition of the event multiplicity has been given in section 4.5.

Energy The total energy contained in the event with fully applied corrections and calibration as discussed in chapter 6.

Standoff distance This observable aims to describe the location of the charge deposits inside the detector. It is defined as the minimum distance of a charge deposit to a detector material other than liquid xenon or the cathode grid. External backgrounds should generally be distributed around small standoff distances due to the good shielding properties of liquid xenon, whereas double beta decay events are evenly distributed within the liquid xenon, which leads to a higher average standoff distance. For multi site events, the standoff distance is defined as the minimal standoff distance of all charge clusters in the event.

Each event that meets the selection criteria discussed in the last chapter is boiled down to these observables, and PDFs in these observables are built for several background sources and locations. The fit extracts the number of events addressed to each background and signal component as well as their single site fractions. A confidence interval for the number of signal events in the data set is extracted using the profile likelihood method and converted into a confidence interval for the $0\nu\beta\beta$ half-life of ^{136}Xe .

8.1. Maximum likelihood method

One method to determine which parameters of a given model best fit a set of observed data is called the maximum likelihood method. The essential quantity used in this

method is called the *likelihood* of a set of parameters $\boldsymbol{\theta}$, given the observed set of outcomes \mathbf{x}_i and assuming a model $f(\boldsymbol{\theta}, \mathbf{x})$, which is the probability density function (PDF) of the observables \mathbf{x} for a specific set of parameters $\boldsymbol{\theta}$. The likelihood is defined as

$$\mathcal{L}(\boldsymbol{\theta}|\mathbf{x}_i) = \prod_i f(\boldsymbol{\theta}, \mathbf{x}_i) \quad (8.1)$$

which is a function of only the parameters $\boldsymbol{\theta}$. Finding the set of parameters $\hat{\boldsymbol{\theta}}$, which best describes the observed distribution of the \mathbf{x}_i is then a matter of maximizing their likelihood:

$$\mathcal{L}(\hat{\boldsymbol{\theta}}) = \max_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}) \quad (8.2)$$

In some special cases this can be done analytically. In most practical applications however, this is done numerically, and since $\mathcal{L} \leq 1$ by definition, calculating \mathcal{L} usually means multiplying lots of small numbers which can easily exceed the numeric floating point limits of common computers. Therefore in practice one usually minimizes the negative logarithm of the likelihood (NLL)

$$-\ln \mathcal{L}(\boldsymbol{\theta}|\mathbf{x}_i) = -\sum_i \ln f(\boldsymbol{\theta}, \mathbf{x}_i), \quad (8.3)$$

which is equivalent to maximizing \mathcal{L} , because the logarithm is a monotonically increasing function. The minimization in this analysis is done using the widely used package MINUIT [79].

8.1.1. Binned likelihood fit

For large data sets, computing the NLL in each minimization step can become quite tedious, since the sum in equation 8.3 runs over all points in the data set. A way around this is to partition the data set into bins and do a binned maximum likelihood fit instead. Although this is accompanied by some loss of information which depends on the bin size, now the sum in equation 8.3 runs over all bins instead, which reduces the computing time drastically for large data sets. The number of entries in a bin j , k_j follows a Poisson distribution with the expected number of entries μ_j being the integral of the PDF over bin j , times the total number of events n addressed to the PDF. The negative log likelihood can then be written as

$$-\ln \mathcal{L} = \sum_j \mu_j - k_j \ln \mu_j + \text{const} \quad (8.4)$$

with $\mu_j = n \int_{\text{bin } j} f(\boldsymbol{\theta}, \mathbf{x}) d\mathbf{x}$

The constant in equation 8.4 only depends on the number of observed entries in each bin and can therefore be dropped when minimizing $-\ln \mathcal{L}$.

8.1.2. Parameter constraints and incorporation of efficiency uncertainties

The maximum likelihood method makes it easy to put constraints on parameters that have been determined externally. The likelihood is simply multiplied by the probability density function of the parameter that was determined in a different study. If for example a set of parameters $\boldsymbol{\theta}$ are known to be normally distributed around central values of $\boldsymbol{\theta}_0$ with covariances $\boldsymbol{\Sigma}$, then one can constrain the likelihood in the following way:

$$-\ln \mathcal{L}_{\text{constr}} = -\ln \mathcal{L} - \ln \mathcal{C} = -\ln \mathcal{L} + \frac{1}{2} (\boldsymbol{\theta} - \boldsymbol{\theta}_0)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{\theta} - \boldsymbol{\theta}_0). \quad (8.5)$$

The normalization term of the multivariate Gaussian in equation 8.5 was again dropped because it is a constant independent of the parameters $\boldsymbol{\theta}$.

This method also allows for an easy way to incorporate systematic uncertainties of signal efficiencies into the likelihood function. The number of events addressed to each background or signal component is simply scaled by a common normalization factor which is centered around 1 and constrained by a Gaussian whose width corresponds to the uncertainty of the signal efficiency.

8.1.3. Confidence intervals

Maximizing the likelihood yields the set of parameters that fit the data best within the given model. The likelihood is however not only useful for point estimation but can also be used to build a test statistic which allows to construct confidence intervals for these parameters. Suppose there is one or more parameter of interest π (the number of $0\nu\beta\beta$ events for example) and a set of additional nuisance parameters $\boldsymbol{\theta}$ (such as the number of events from different backgrounds). Then one can construct a likelihood-ratio test statistic which eliminates the nuisance parameters, called the profile likelihood (see for example [80])

$$\lambda(\pi|\mathbf{x}_i) = \frac{\sup [\mathcal{L}(\pi, \boldsymbol{\theta}, \mathbf{x}_i); \boldsymbol{\theta}]}{\sup [\mathcal{L}(\pi', \boldsymbol{\theta}, \mathbf{x}_i); \boldsymbol{\theta}, \pi']} \quad , \quad (8.6)$$

where the supremum in the denominator runs over the full parameter space $(\pi, \boldsymbol{\theta})$ and the one in the numerator runs over the nuisance parameters only. The profile likelihood λ therefore is only a function of the parameter of interest π . The distribution of λ can be found via a Monte Carlo study. However under very general conditions, the distribution of λ can be approximated by

$$\lambda(\pi) \sim \chi_{\text{dim}(\pi)}^2 \quad (8.7)$$

where $\chi_{\text{dim}(\pi)}^2$ is the chi square distribution with degrees of freedom equal to the number of parameters of interest. This relation is called Wilks' theorem [81] and will be applied in this analysis to find confidence intervals for the $0\nu\beta\beta$ half-life.

8.2. Modelling of signal and backgrounds using the Monte Carlo simulation

The Monte Carlo package developed for a full detector simulation was described in section 3.6. This package is used to build PDFs for a set of backgrounds from various places within and around the detector. The simulated events are processed by the same reconstruction algorithm as the detector data, for the energy of the simulated event however, the true energy is used. This ensures that the simulation is subject to similar signal efficiencies as the data from the detector. Residual differences in signal efficiency originating for example from the differences in electronic noise and incompleteness of the Monte Carlo model will be discussed at the end of this section. For each background component, a single site and a multi site PDF are created by condensing the simulated events to the two remaining observables standoff distance and energy. The energy distribution is smeared according to the measured energy resolution function discussed in section 6.2.3.

8.2.1. Background model

The following background PDFs are built for the fit:

- Backgrounds from within the liquid xenon: $2\nu\beta\beta$, ^{135}Xe , ^{137}Xe , ^{222}Rn .
- Backgrounds from the detector copper material: ^{238}U , ^{232}Th , ^{40}K , ^{60}Co , ^{65}Zn , ^{54}Mn .
- Backgrounds from radon in the air gap between the cryostat and the lead shield wall.
- Backgrounds from ^{232}Th in the inner cryostat wall.
- Backgrounds from captures of cosmogenic and radioactivity induced neutrons on the detector and cryostat copper, as well as on the xenon and HFE fluids.

The two double beta decay modes are simulated using the Fermi function proposed in [82].

Both xenon background isotopes are short-lived beta emitters which can be produced by cosmogenic activation of the main isotope and the major residual isotope (^{134}Xe) respectively.

For the PDFs of ^{238}U and ^{232}Th , the entire decay chains were simulated with the assumption of secular equilibrium between the single states. For ^{222}Rn , the chain is simulated down to ^{210}Pb .

The radon chain in the liquid xenon needs special treatment because alpha decays are usually vetoed due to their high light to charge ratio. While the final data set should contain almost no alpha decay events due to the cut efficiency of almost 100 %, simulated events would not be vetoed, because the simulation is not capable of reproducing the




Figure 8.1.: Two of the PDFs used in the fit. The top two are the SS and MS PDF respectively for ^{232}Th in the TPC material, the bottom two are SS and MS PDFs for $2\nu\beta\beta$. Projections onto the energy and standoff distance axes are shown in each case. All scales are linear.

correct light to charge ratio for alpha events. Therefore the simulated spectrum of ^{222}Rn in the liquid xenon would be overpopulated by these events, if the whole decay chain were simulated. For this reason the ^{222}Rn decay chain is split into relevant components that contain beta and gamma decays.

Due to the short $164\text{ }\mu\text{s}$ half-life of ^{214}Po which decays under the emission of an alpha particle, the preceding ^{214}Bi beta decay is usually also vetoed by the light to charge ratio cut. However since the decay products are ionized, they accumulate on the surface of the cathode grid. Therefore in approximately 50 % of the cases, the alpha particle emitted by ^{214}Po is directed towards the cathode and hence not seen by the detector. This means that the preceding ^{214}Bi decay will eventually not be vetoed and ^{214}Bi beta decays on the cathode surface must be treated by a separate PDF.

The neutron capture backgrounds are not entirely simulated using the GEANT4 software package. Instead a multi-step approach is done. Muon induced neutrons are simulated by using the observed muon flux at the site (measured in [83]) as an input to a FLUKA [84, 85] based simulation which contains a simple model of the detector. The FLUKA simulation takes care of the propagation of the muons as well as spallation, propagation and capture of neutrons within the detector materials. It outputs capture gammas which in turn are used as input to the default GEANT4 based simulation. Radioactivity induced neutrons from the surrounding salt are first generated using the SOURCES4A code [86] and then put into the FLUKA simulation.

Model completeness

The background model contains all sources within the detector materials that are expected to yield a significant contribution to the energy spectrum. However it must be checked whether more remote backgrounds for example from the cryostat walls or the HFE cooling liquid need to be included in the model. A study has been done whether the addition of the ^{238}U and ^{232}Th chains or ^{60}Co in the HFE or copper cryostat is a reasonable extension of the model. With the ^{222}Rn in the air gap PDF, the background model as described above already contains a remote component that is part of the uranium chain. The expected number of events assigned to an additional remote uranium component in either the HFE or the cryostat walls by the fit was found to be highly anticorrelated with the number of events the fit assigns to radon in the air gap. While the best fit is reached with a non-zero contribution of both, with a slight preference to the radon in the air gap PDF, the hypothesis of either of the two components being zero is consistent within 1σ . This means that the fit is not able to discriminate between the exact locations of remote uranium components and the radon in the air gap PDF will act as a collecting point for remote uranium backgrounds. The effect of using the radon in the air gap PDF as a placeholder for all remote uranium backgrounds on the ROI has been checked. Substituting this PDF by other remote uranium PDFs changes the expected number of counts in the ROI by 2.7 % at most. The situation is completely analogous with remote ^{232}Th components. The fit already contains a remote thorium component at the inner cryostat wall and an additional remote component does not add any information to the fit. The situation with ^{60}Co is also similar. Since the guide tube

for the calibration sources is made out of different copper than the TPC parts themselves, significant additional remote cobalt backgrounds are expected to come primarily from there. Exchanging the source of the ^{60}Co background between the TPC copper and the calibration source guide tube changes the expected number of events in the ROI by 5.0 %.

8.2.2. Agreement between simulation and data

In order to obtain reliable results from the fit, it is essential that the simulation, which is used to produce the PDFs, can accurately reproduce the background spectra and their single site / multi site distribution. While this can not directly be tested for all backgrounds believed to be in the detector, the calibration sources and simulations thereof provide a powerful tool to cross-check the agreement between the simulation package and the real detector. The large scale calibration campaigns during which several of the calibration sources are deployed within the timescale of a week are used for this purpose.

Figure 8.2 shows the spectrum of the ^{226}Ra calibration source taken during a calibration campaign in June 2013 and the spectrum obtained using the simulation of the source at the same location in the cathode plane. The qualitative agreement between simulation and data is very good. The largest disagreement is observed at low standoff distances (near the detector walls). This trend is common for all available calibration sources as can be seen in figure 8.4, which shows the residual disagreement between simulated and measured spectra for all sources. While the reason for this disagreement near the detector boundaries is not entirely known, its effect on the signal is quantified further below and accounted for in the analysis.

Since the distinction between single site events and multi site events carries much information about whether an event is signal-like or background-like, it is also important that the simulation correctly predicts the fraction at which an event of a certain source is single site over the full energy spectrum. The calibration source campaigns are again used to validate the predictions from the simulation. Figure 8.3(a) shows the single site fractions of three calibration sources. While they all share an overall downward trend with increasing energy, specific features are expected to be different due to different event topologies. The full absorption gamma peak for example is generally more likely to be single site than the Compton shoulder right next to it. Figure 8.3(b) shows the agreement in the single site fraction between the same data and a simulation of the sources. Such deviation curves have been calculated for all available calibration campaigns and sources and we take the weighted average of the maximum deviation of each curve as a conservative estimate of the single site fraction disagreement between simulation and data. The weighting is by the inverse of the error and the live time of the calibration source run. The systematic error on the single site fraction calculated in this way is 9.6 %.

8. Search for neutrinoless double beta decay

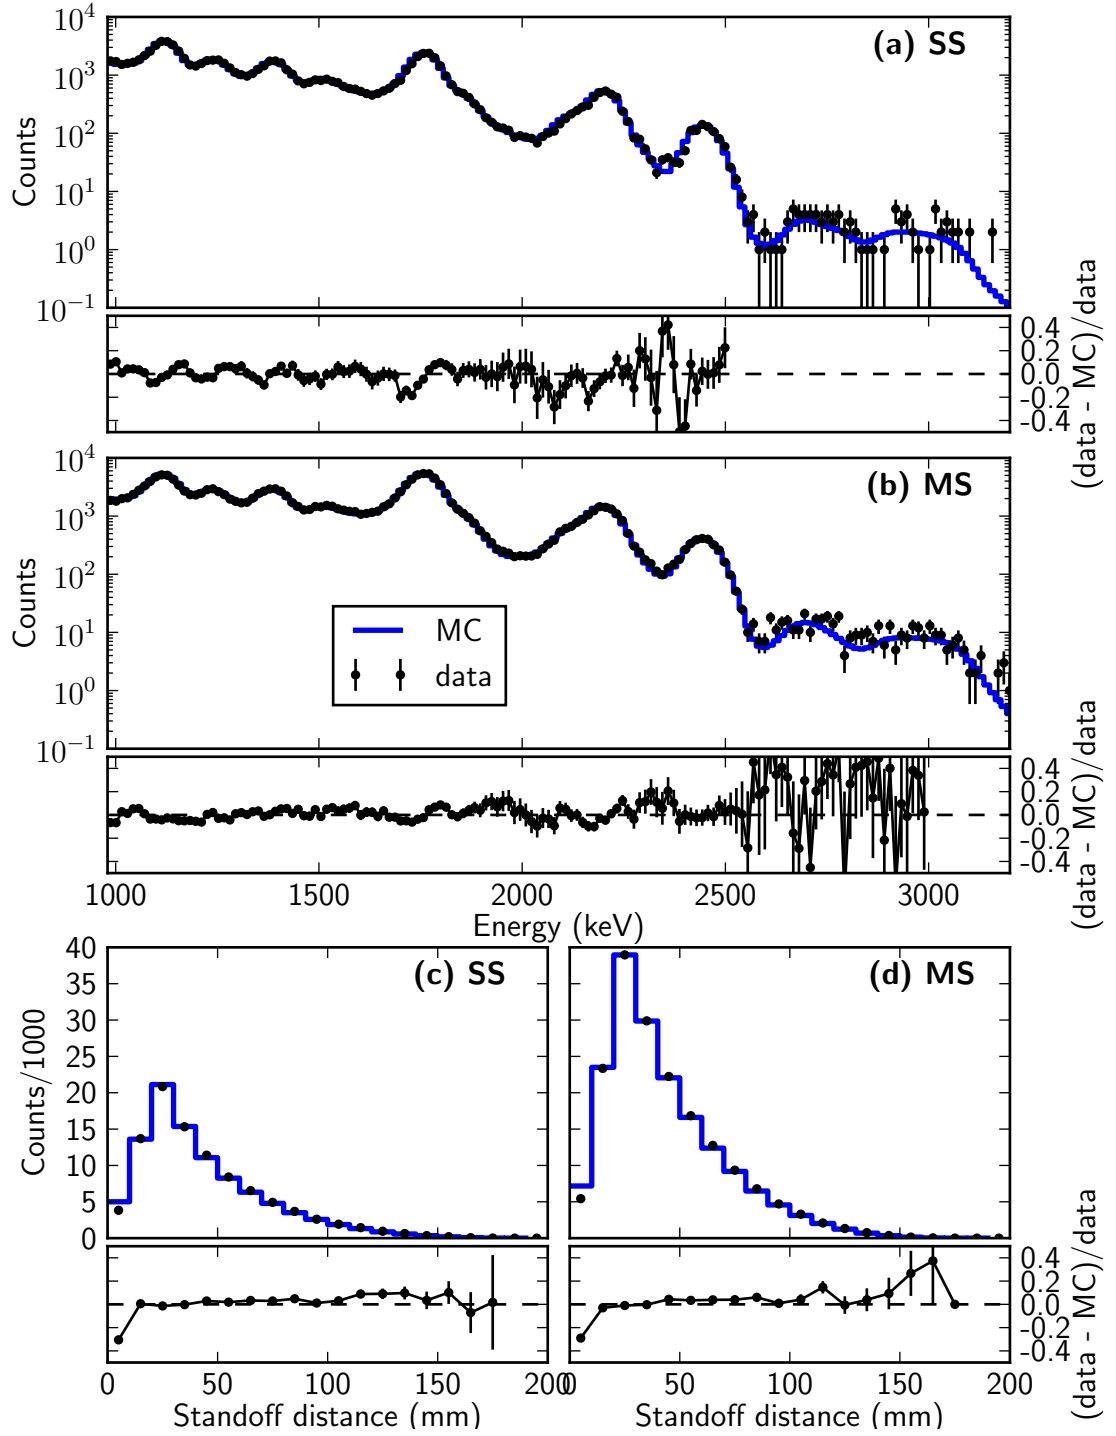


Figure 8.2.: Agreement between spectra produced using the simulation package (blue) and calibration data (black) for data obtained from the ^{226}Ra calibration source which was located in the cathode plane. Plot (a) and (b) show the SS and MS energy spectra respectively, plot (c) and (d) show the SS and MS standoff distance distributions. Overall a very good agreement is observed, with the exception of the first standoff distance bin (0 mm to 10 mm)




Figure 8.3.: Event multiplicity of calibration sources. (a) shows a plot of the single site fraction $SS/(SS + MS)$ of the ^{228}Th , ^{226}Ra and ^{60}Co calibration sources versus energy. All three calibration sources show a similar energy dependence. (b) shows the agreement between the single site fraction predicted by the simulation and the one observed in data versus energy. Data taken from [87].

Effect of PDF shape uncertainties

A monte carlo study has been performed to estimate the effect of the residual shape disagreement between the PDFs and data on the expected number of counts in the region of interest (ROI). The basic idea is to observe the disagreement between the number of events in the ROI as reported by the fit and the true number of events in the ROI due to the uncertainties in the shape of the background PDFs. To accomplish this, a large number of toy data sets are generated from PDFs on which the shape disagreement is divided out (called unskewed PDFs in the following). These toy data sets are then fit with the original PDFs and the fractional difference in the estimated number of events in the ROI obtained in this way compared to what is obtained without dividing out the shape disagreement is used as a measure of the impact of the shape disagreement on

the counts in the ROI. The PDFs are unskewed by dividing out the calibration source spectral residuals shown in figure 8.4. Uranium chain and the ^{214}Bi PDFs are divided by the radium source residual curves (figure 8.4 black), ^{60}Co background PDFs are divided by the cobalt source residual curves (figure 8.4 red) and the remaining PDFs, with the exception of neutron capture PDFs, are divided by the thorium source residual curves. The neutron capture PDFs are not unskewed because no similar calibration source is available and their energy range spectrum extends beyond the energies provided by the calibration sources. Due to their small contribution to the overall background rate their contribution to the total PDF shape uncertainty systematic is expected to be negligible. Figure 8.5 shows a comparison of an unskewed PDF to its original version on the example of the PDF for the ^{238}U background in the detector copper. The toy data sets are generated from the unskewed PDFs, with each PDF contributing to the toy data set according to its fractional contribution to the low background data as estimated by a fit of the PDFs to a subset of the data used in this analysis. These toy data sets are fit with the original PDFs and the estimated counts in the ROI are compared to what would be obtained without unskewing the toy data sets. The results are shown in figure 8.6. The disagreement in the spectral shapes between the simulation and data thus lead to a 9.2% bias in the estimated number of events in the ROI. The study is also done with unskewing the energy or standoff distance dimension only. It is apparent that the shape disagreement in the standoff distance is the main contributor to the effect, most likely due to the large disagreement in the first (0 mm to 10 mm) bin.

8.3. Low background data fit

Information on the number of $0\nu\beta\beta$ events is extracted by a binned maximum likelihood fit of the set of background and signal PDFs on the data set, as introduced in section 8.1.1. The likelihood that is minimized is slightly more complicated than in equation 8.4. The data set is partitioned in a single site and a multi site part, but both parts are simultaneously fit to extract information on the total number of specific background and signal events. In addition, some parameter constraints are added to the likelihood, and normalization variables are introduced to incorporate systematic effects that affect the overall efficiency as well as specifically the efficiency in the signal region. The expectation value for the number of counts in bin j of the single site (multi site) data set, $\mu_j^{\text{SS (MS)}}$, is the product of an overall normalization parameter N and the sum over the individual expected number of single site (multi site) events contributed to bin j from the background or signal component i . The latter in turn are the products of the estimated total number of events from component i , n_i and the relative fraction (1 - the relative fraction) of single site events in component i , m_i . With the same conventions




Figure 8.4.: Residual difference of the calibration source spectra between simulation and data for the ^{228}Th , ^{226}Ra and ^{60}Co calibration sources used to unskew the PDFs (see text). Plot (a) and (b) show the differences in energy spectra for SS and MS respectively, (c) and (d) show the differences in standoff distance distributions. Data from [87].

8. Search for neutrinoless double beta decay




Figure 8.5.: Example of an original (black) and and unskewed PDF (red) on the example of the ^{238}U in the vessel multi site PDF. The unskewing was done by dividing the PDF by the ^{226}Ra residual curve of figure 8.4.

as used in equation 8.4 and before, the likelihood then takes the following form:

$$-\ln \mathcal{L} = \sum_j \mu_j^{\text{SS}} + \mu_j^{\text{MS}} - (k_j^{\text{SS}} \ln \mu_j^{\text{SS}} + k_j^{\text{MS}} \ln \mu_j^{\text{MS}}) - \ln \mathcal{C} \quad (8.8)$$

$$\text{with } \mu_j^{\text{SS}} = N \int_{\text{bin } j} \left[\sum_i n_i m_i f_i^{\text{SS}}(\mathbf{x}) + n_{\text{neutron}} \sum_l p_l m_l f_l^{\text{SS}}(\mathbf{x}) \right] d\mathbf{x}$$

$$\mu_j^{\text{MS}} = N \int_{\text{bin } j} \left[\sum_i n_i (1 - m_i) f_i^{\text{MS}}(\mathbf{x}) + n_{\text{neutron}} \sum_l p_l (1 - m_l) f_l^{\text{MS}}(\mathbf{x}) \right] d\mathbf{x},$$

with the constraints labeled as \mathcal{C} . The index j runs over all bins and the index i labels a specific background or signal component except the neutron components which are labeled by the index l . The second term in μ_j is the contribution from neutron components, where the contribution of a single component l is proportional to the product




Figure 8.6.: The bias introduced into the expected number of events in the ROI due to the disagreement of the energy and standoff distance distributions in simulation and data. The plot shows the biases if only the disagreement in energy distributions (blue), in standoff distance distributions (red), or in both (black) is taken into account.

of the total neutron capture counts n_{neutron} times the relative fraction p_l of this component to the total neutron capture counts. The common normalization parameter N is constrained to be centered around 1 with a standard deviation corresponding to the uncertainty in the signal efficiency.

To account for an efficiency that is specific to the signal region, the single site PDFs $f_i^{\text{SS}}(\mathbf{x})$ are split into a part outside and a part inside the ROI:

$$f_i^{\text{SS}}(\mathbf{x}) = r_i(1 - \alpha)f_i^{\text{SS inner}}(\mathbf{x}) + [1 - r_i(1 - \alpha)]f_i^{\text{SS outer}}(\mathbf{x}), \quad (8.9)$$

where r_i is the relative integral of the PDF inside the ROI to outside the ROI:

$$r_i = \frac{\int f_i^{\text{SS inner}}(\mathbf{x}) d\mathbf{x}}{\int f_i^{\text{SS outer}}(\mathbf{x}) d\mathbf{x}} \quad (8.10)$$

This allows the normalization of the PDFs within the ROI to float independent from the normalization of the complete PDF. The common parameter α is constrained to be centered around 0 with a standard deviation corresponding to the signal efficiency specific to the region of interest. Table 8.1 gives a summary of all parameters that are floating in the fit.

8. Search for neutrinoless double beta decay

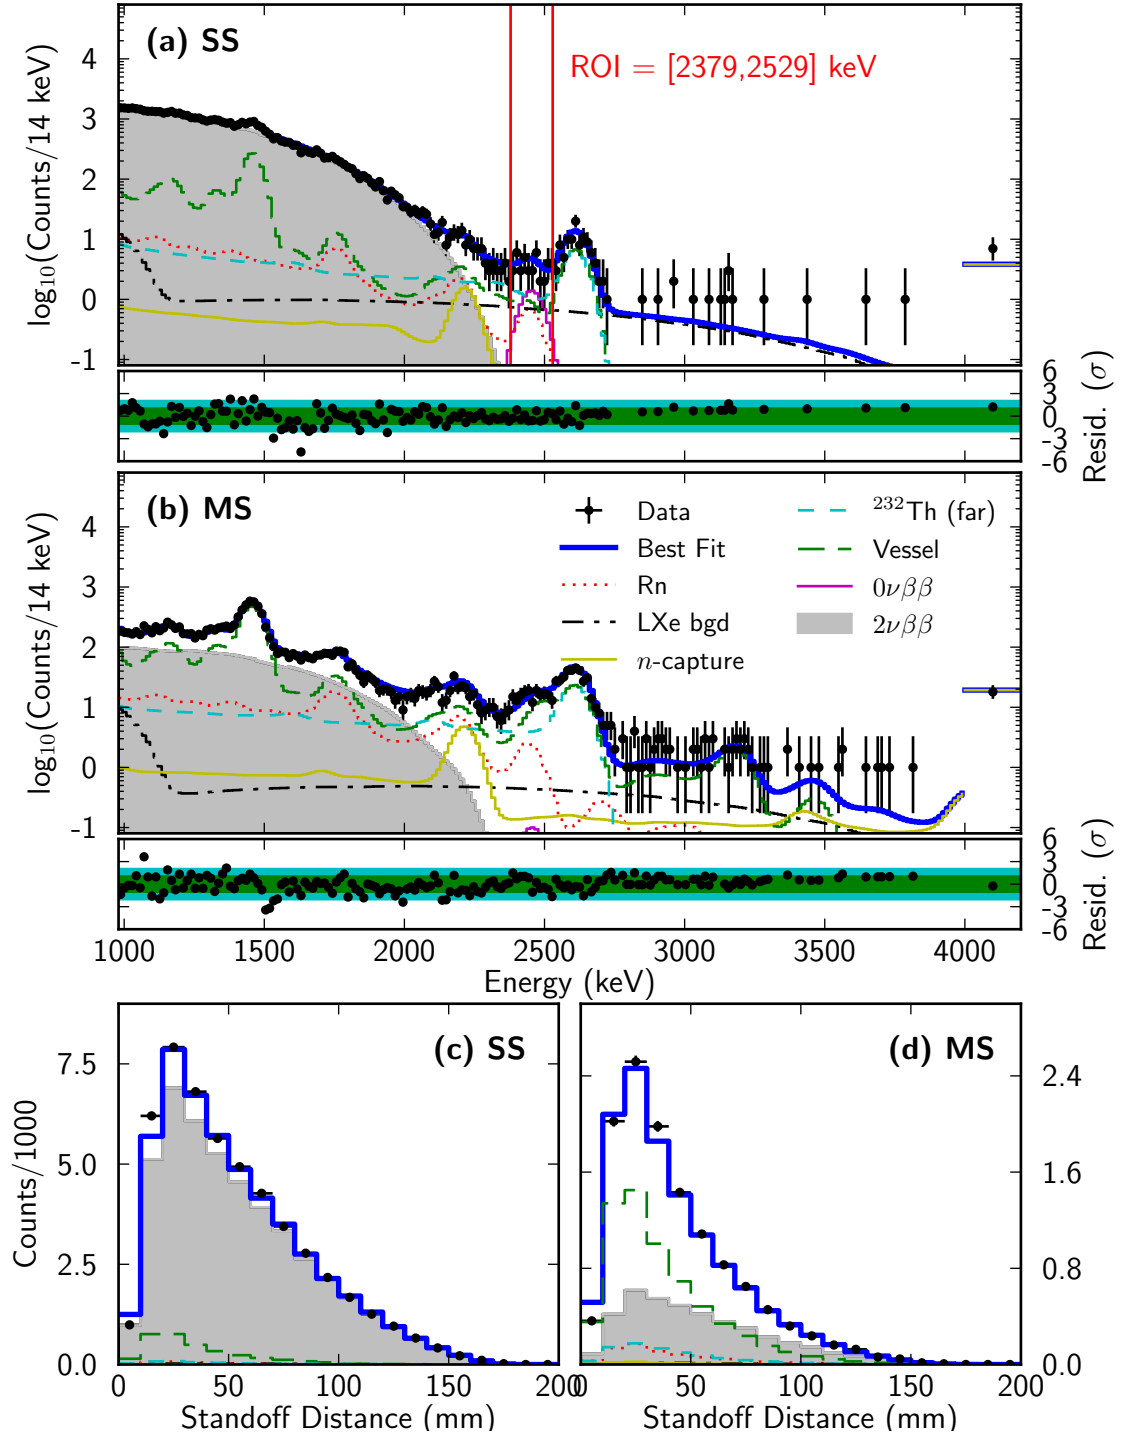


Figure 8.7.: Fit of the PDFs to the low background data set. Projections onto the energy axis are shown for SS events (a) and MS events (b), projections on the standoff distance axis are shown in (c) and (d). Several similar PDFs are grouped together for better visibility. The lower panels of the plots show residuals between data and best fit, normalized to the Poisson error of the data points. The green (blue) shaded regions in these panels indicate $\pm 1\sigma$ ($\pm 2\sigma$) deviations. The last bin in the energy spectrum summarizes events and PDFs for energies between 4000 and 9800 keV. Figure adapted from [15].

parameter	description	constrained
n_i	total counts of component i	no, except for ^{222}Rn
n_{neutron}	total neutron capture counts	no
p_l	fractional contribution of individual neutron capture PDF to n_{neutron}	yes
m_i	single site fraction of component i	yes, fixed for neutron capture
N	overall normalization factor shared by all components	yes
α	ROI specific normalization factor shared by all components (that extend to the ROI)	yes

Table 8.1.: Complete list of parameters that are floating in the fit.

8.3.1. Parameter constraints

Of the parameters listed in table 8.1, only the expected number of events for each background or signal component are floating freely in the fit. The rest of the parameters is constrained using multivariate Gaussian constraints, whose central values and (co-)variances have been determined in separate measurements.

Single site fraction

Without a constraint on the fraction at which an event from a certain background or signal component is single site, the partition into single site and multi site data sets would add no discrimination power between signal and background. Since there is no practical way of measuring the single site fractions of each component separately, the values predicted by simulation are used as central values in the constraint. The error on these numbers is then given by how accurately the simulation can reproduce the correct single site fractions. This has been discussed in section 8.2.2, and the averaged maximum deviation in the single site fractions between simulation and calibration source data (9.6 %) is used as the square root of the parameters' variance. Despite the different energy spectra, the three calibration sources show the same energy behaviour in the single site fractions, as is apparent in figure 8.3(a). We therefore impose a 90 % correlation on the constraint's covariance matrix.

Neutron capture components

The fractional contribution of each neutron capture PDF to the total number of neutron capture events is constrained to the prediction of the FLUKA simulations with a tolerance of 20 %.

Amount of ^{222}Rn in the liquid xenon

The amount of ^{222}Rn in the active detector region is continually monitored by an independent analysis, which estimates its activity to $3.65(37) \mu\text{Bq kg}^{-1}$. Following the discussion in section 8.2.1, a large part of the decay daughters of the radon in the liquid xenon accumulate on the surface of the cathode due to the high electric field. Studies using the ^{214}Bi - ^{214}Po coincidence have shown that 83 % of these events occur on the cathode, with the remaining 17 % decaying in the bulk of the xenon. The expected number of decays corresponding to the radon PDFs in the liquid xenon are constrained by these fractions times the measured total activity, with a tolerance of 10 % and a correlation across the different radon components of 90 %.

Signal efficiencies

The two normalization parameters N and α are constrained by Gaussians whose standard deviations correspond to the overall uncertainty in signal efficiency and the systematic errors on background determination in the ROI, respectively. The effects contributing to the overall signal efficiency were summarized in table 7.2. Three effects contribute to the error on background determination in the ROI. The uncertainties in the expected number of background events due to a possible incompleteness of the model were discussed in section 8.2.1. Adding the different contributions in quadrature yields a 5.7 % error. The effect of uncertainties in the PDF shapes produced by the simulation was studied in section 8.2.2 and contributes a 9.2 % error to the background determination in the ROI. Another contribution to the ROI background error stems from the error of the energy resolution parameters. While they are known to vary slightly in time, only average parameters can be used to smear the PDFs. This contributes a 1.5 % error to the background determination in the ROI. Adding the three parts in quadrature yields a total error of 10.9 %.

8.3.2. Fit results

Since the energy calibration is done with gamma sources, we examine the possibility that the energy scale for beta decay events (signals as well as backgrounds) is slightly different. The hypothesis that is tested is that the energy scale for beta decays is directly proportional to the energy scale for gamma decays, where the constant of proportionality is referred to as the β/γ scale. Since $2\nu\beta\beta$ is the largest background, it is the main tool to constrain the β/γ scale. The likelihood space for different β/γ scales is explored manually, by performing fits with different sets of PDFs, where beta decay PDFs were scaled by the β/γ scale compared to gamma background PDFs. Figure 8.8(b) shows the profile likelihood for the β/γ scale, which yields a value of 0.999 ± 0.002 . This is consistent with the assumption, that gamma decays and beta decays have the same energy scale in the detector.

The minimization of the likelihood in each fit is performed using MINUIT. Figure 8.7 shows the best fit of the PDFs to the data set. Similar backgrounds have been grouped

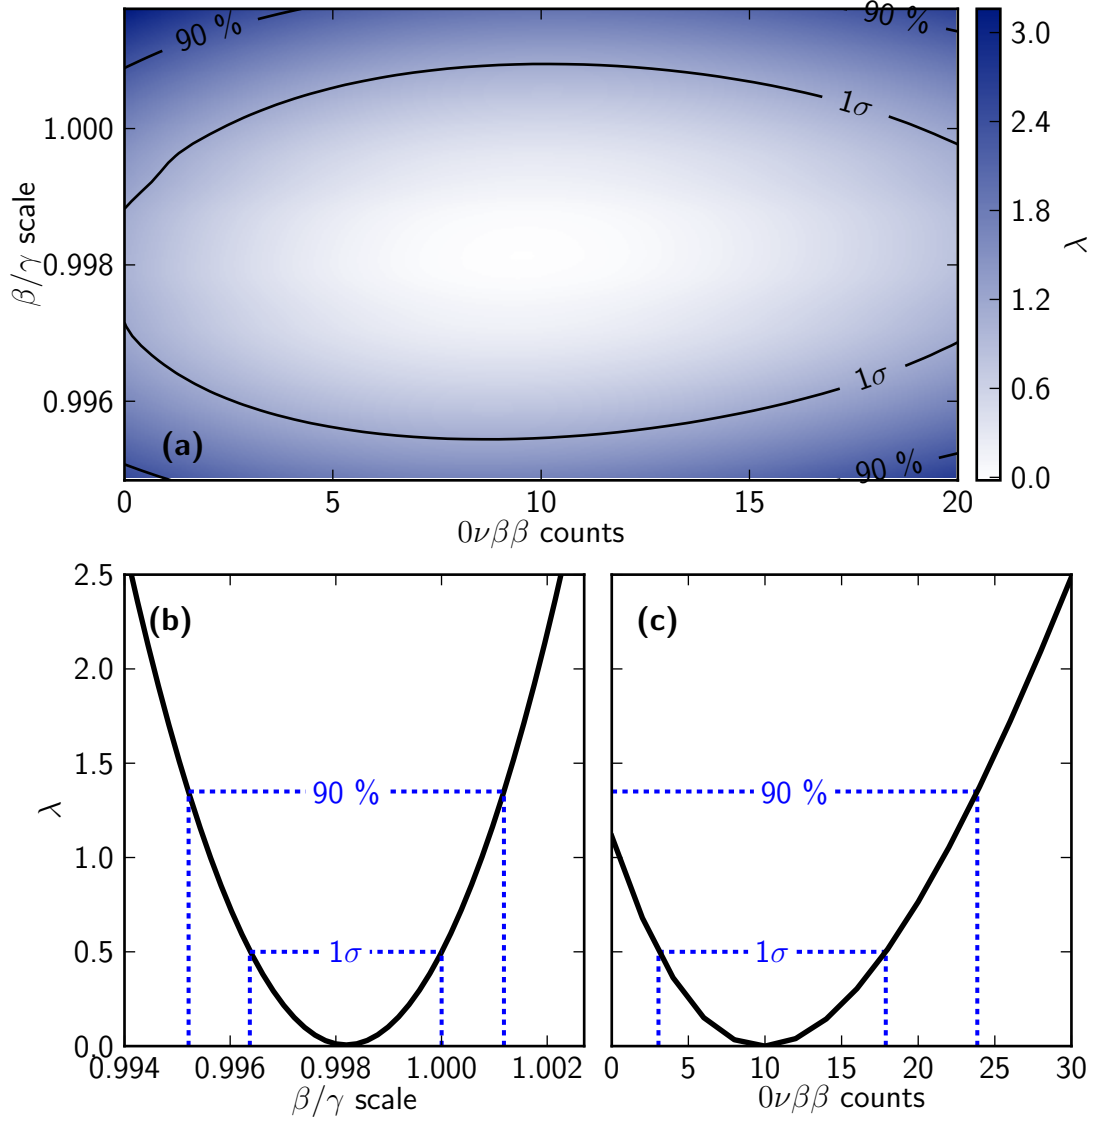


Figure 8.8.: Panel **(a)** shows the two-dimensional profile likelihood of the number of $0\nu\beta\beta$ counts and beta scale. Panels **(b)** and **(c)** show one-dimensional profiles in the β/γ scale and the number of $0\nu\beta\beta$ counts respectively, where the other parameter is treated as floating in the fit. 1σ and 90% confidence regions for the two parameters are shown.

8. Search for neutrinoless double beta decay

for better visibility. Several gamma full absorption peaks can be seen, most prominently those of ^{40}K (1460.8 keV [88]) and ^{208}Tl (2469 keV). The two neutrino decay mode of ^{136}Xe is clearly the most dominant background, accounting for over 75 % of the events in the data set, followed by ^{40}K , which contributes less than 10 % of events.

For each value of the beta scale, the profile likelihood around the best fit value of the number of neutrinoless double beta decay events is calculated. This gives a 2D profile likelihood which is shown in figure 8.8(a). To construct a confidence interval for the number of $0\nu\beta\beta$ events, the β/γ scale is treated as a nuisance parameter like the other background parameters. This is accomplished by minimizing the 2D profile likelihood over the β/γ scale for each number of $0\nu\beta\beta$ events. The resulting profile is shown in figure 8.8(b). The 90 % confidence upper limit on the observed number of $0\nu\beta\beta$ events obtained from this profile is 23.9 counts. In total, 39 events were observed within the 2σ ROI. The number of background events within the 2σ ROI expected from the best fit model is $31.1 \pm 1.8(\text{stat}) \pm 3.3(\text{sys})$ events, which corresponds to a background rate of $1.7(2) \times 10^{-3} \text{ keV}^{-1} \text{ kg}^{-1} \text{ yr}^{-1}$. Table 8.2 gives a summary of the number of background

Component	Estimated counts in 2σ ROI
^{232}Th	16.0
^{238}U	8.1
^{137}Xe	7.0
Total background	31.1
$0\nu\beta\beta$	9.9
Total observed events	39

Table 8.2.: Summary of backgrounds in the 2σ ROI estimated by the fit.

events in the 2σ ROI as expected by the fit. The best fit value of 9.9 $0\nu\beta\beta$ events is consistent with the hypothesis of no neutrinoless double beta decay within 1.2σ . The upper limit on the observed number of $0\nu\beta\beta$ events is translated to a 90 % confidence lower limit on the half-life in the following way:

$$T_{1/2} = \ln 2 \cdot \frac{Mt\epsilon N_A}{mN}, \quad (8.11)$$

with the total exposure, $Mt = 100 \text{ kg yr}$, the signal efficiency $\epsilon = 0.846$, the Avogadro constant N_A , the molar mass of ^{136}Xe m and the upper limit on $0\nu\beta\beta$ events N , the lower limit on the $0\nu\beta\beta$ half-life is $1.1 \times 10^{25} \text{ yr}$.

Cross-checks

To verify the above result, several cross-checks have been performed. The event accumulation in the ROI was found to be consistent with a constant rate, and performing the analysis on two separate subsets of the total data set gives estimates of background rates that are consistent with each other and with the values obtained from the total data set. Both checks are in accordance with the assumption that no short lived background polluted the region of interest. In addition the waveforms of the 39 events in




Figure 8.9.: Cross check of the result using Monte Carlo studies. **(a)** shows the distribution of $\lambda_{\text{no-}0\nu\beta\beta}$ (see text for definition) for a set of toy data sets generated from the best fit background model. In less than 90 % of the toy cases, $\lambda_{\text{no-}0\nu\beta\beta}$ attains a smaller value than the one obtained from data. **(b)** estimates the experimental sensitivity. The 90 % confidence upper limit on the number of $0\nu\beta\beta$ counts is calculated for each toy data set. The median of these values is taken as the experimental sensitivity: $\sigma_{\#0\nu\beta\beta} = 14$. Study done by Ryan Killick. Figure adapted from [89].

the 2σ ROI have been investigated by eye, and no reconstruction irregularities were found. Although the background model is believed to be complete, the presence of several additional backgrounds was checked, motivated mainly by the possible presence of $^{110\text{m}}\text{Ag}$ and ^{88}Y in the KamLAND-Zen experiment [38]. Both backgrounds are found to produce high multiplicity events in EXO-200, with a single site fraction of less than 10 %. Subsequently their contribution to the ROI is expected to be very small. Separate profile likelihood scans with these components added separately yields 1σ upper limits of 0.06 counts in the 2σ ROI for $^{110\text{m}}\text{Ag}$ and 0.03 counts in the 2σ ROI for ^{88}Y .

Assuming that the background model is correct and complete, one can validate the conclusions drawn from the likelihood profile in figure 8.8 using a toy Monte Carlo study. The value of the likelihood profile at its best fit value $\lambda_{\text{best-fit}} = 0$ differs from the value of the no- $0\nu\beta\beta$ hypothesis $\lambda_{\text{no-}0\nu\beta\beta} = 1.134$. The distribution of $\lambda_{\text{no-}0\nu\beta\beta}$ when no neutrinoless double beta decay is present can be estimated by calculating $\lambda_{\text{no-}0\nu\beta\beta}$ for a set of toy data sets generated from the best fit distribution of backgrounds. The resulting distribution of $\lambda_{\text{no-}0\nu\beta\beta}$ is shown in figure 8.9(a). In less than 90 % of the cases, $\lambda_{\text{no-}0\nu\beta\beta}$ attains a value which is smaller than the value obtained from data. This means that the observed value is not strong enough to reject the hypothesis of no neutrinoless double beta decay at the 90 % confidence level.

The limit reported by a low background experiment depends on the number of events observed in the region of interest. A statistical downward fluctuation in this number leads to a stronger limit, and the opposite case will lead to a weaker limit. This is a fundamental property of low statistics experiments. A magnitude that aims to characterize the strength of an experiment independent of such fluctuations is called the sensitivity of the experiment. It is appropriate to define the sensitivity of the experiment as the median upper limit, the experiment would report if it were repeated a large number of times. Assuming the background model is correct and complete, this can again be estimated using a toy Monte Carlo study. From each toy data set, randomly generated from the best fit composition of backgrounds, a 90 % confidence upper limit on the number of $0\nu\beta\beta$ events is calculated. The resulting distribution of limits is shown in figure 8.9(b). The median upper limit estimated in this way is 14 counts, compared to the limit of 24 counts obtained from the data. The observed 24 events are consistent with an unlucky upward fluctuation of the backgrounds in the region of interest with a p-value of 0.14. The experimental sensitivity of 14 counts corresponds to a half-life of 1.9×10^{25} yr.

8.4. MST metric as observable in the fit

The MST metric presented in chapter 5 gives a finer graded description of the event multiplicity than the default single site / multi site categorization, and therefore possibly more information on backgrounds. In this section, an analysis that uses this new metric as an additional observable in the fit is presented. The most important requirement, reasonable agreement in the metrics distribution between what is observed in data and what is predicted by simulation, has been demonstrated in chapter 5 (figure 5.3).

The introduction of this observable supersedes the splitting of the data set and PDFs




Figure 8.10.: Best fit of the background model to the low background data in a two dimensional fit with energy and the MST metric as observables. The main panel shows the full range of the PDFs and data projected onto the MST metric dimension. The inset shows a zoom into the region of low values of the MST metric. The Lower panel shows residual between data and the best fit PDF, normalized to the Poisson error of the data bins.

8. Search for neutrinoless double beta decay

into SS and MS parts. Consecutively, the total likelihood which is minimized takes a somewhat simpler form than in equation 8.8:

$$-\ln \mathcal{L} = \sum_j \mu_j - k_j \ln \mu_j - \ln \mathcal{C} \quad (8.12)$$

$$\text{with } \mu_j = N \int_{\text{bin } j} \left[\sum_i n_i f_i(\mathbf{x}) + n_{\text{neutron}} \sum_l p_l f_l(\mathbf{x}) \right] d\mathbf{x},$$

with the same conventions as used in equations 8.8 and above. N is the overall normalization factor, n_i and $f_i(\mathbf{x})$ are the total counts and PDF of background or signal component i respectively, n_{neutron} is the number of counts addressed to neutron capture processes and p_l is the fraction at which the neutron capture component l contributes to the total neutron capture PDF. The PDFs $f_i(\mathbf{x})$ are split around the ROI in an analogous way as the single site PDFs in equation 8.9, to allow for a ROI-specific normalization uncertainty. The number of fit parameters is less than in table 8.1, because the single site fractions are obsolete.

To use the MST metric as a third observable in addition to the total event energy and the standoff distance, the number of simulated events for each background component must be increased accordingly to guarantee minimal uncertainties in each bin of the three-dimensional PDFs. The binning in the MST dimension has been optimized in this respect with fine binning for the details at low values and much coarser binning for larger values, totalling 30 bins. To keep the uncertainty per bin at the same level, the number of simulated events would thus have to be increased by a factor of 30, assuming uniform MST distributions. This alone already requires a substantial increase of computational resources. An additional penalty is imposed by the fact that to obtain realistic MST distributions, ionization electron diffusion must be simulated as demonstrated in chapter 5. Since the simulation of the electron drift and waveform generation is the computationally most intensive part in the simulation, this introduces an additional increase in computing time.

For this reason, only a smaller amount of simulated events was produced at the time of this analysis, with approximately 1/4 the number of simulated events that were produced for the analysis presented in the previous sections. The limited statistics do not allow to produce three-dimensional PDFs with tolerable uncertainties in their shape. Instead, only energy and the MST metric are used as observables in the fit and the likelihood profile on $0\nu\beta\beta$ counts is compared to what is obtained if energy and standoff distance or only energy is used as observables in the fit. Due to the smaller number of simulated events and increased number of bins, the uncertainties in the shape of the PDFs are larger than in the analysis presented in the last section. Figure 8.10 shows the best fit of the energy and MST metric fit, projected onto the MST metric dimension and table 8.3 compares the estimates of the largest backgrounds in the 2σ ROI obtained from the energy + MST metric fit to the values obtained by the default fit. In this case the 2σ energy ROI is defined on the whole data set, not just the single site data set as in the previous sections. The best fit values of the main backgrounds in the ROI are comparable between the energy + MST metric and the energy + standoff distance fit.




Figure 8.11.: Profile likelihood in the number of $0\nu\beta\beta$ counts for the energy + MST metric fit and the energy + standoff distance and energy only fits. Systematic errors are assumed to be those presented in section 8.3.1 for all three cases. For the energy + standoff distance profile, the systematic error due to the uncertainties in the single site fractions is not taken into account.

Generally, the systematic effects related to the MST metric observable would need to be evaluated. The detailed investigation of these is deferred for later analyses, since the introduction of diffusion into the simulation potentially affects several of the efficiencies described in section 8.3.1 in a non-trivial way. Instead, the systematic errors in the energy + MST metric fit are assumed to be the same as those determined for the energy + standoff distance fit. Since the spectral differences between simulation and calibration source data are comparable for the standoff distance and MST metric observables, this assumption seems reasonable. Because this means that the error in the MST metric spectrum is not explicitly taken into account, the systematic error associated with the single site fraction uncertainties is also not taken into account, when comparing the energy + standoff and energy only fits to the energy + MST metric fit. This is incorporated in the profile likelihood curves in figure 8.11, which shows the profile of the number of $0\nu\beta\beta$ events obtained using the energy + MST metric fit and compares it to those obtained in the energy + standoff and energy only fits. The curvature of the energy + MST metric profile is comparable to the one for the energy + standoff distance fit, and clearly larger than for the energy only fit. In this configuration, the energy + MST metric fit gives the best 90 % CL upper limit on the number of $0\nu\beta\beta$ events of 18.5, which translates into a lower limit on the half-life of neutrinoless double beta decay in ^{136}Xe of

$$T_{1/2} > 1.4 \times 10^{25} \text{ yr.} \quad (8.13)$$

8. Search for neutrinoless double beta decay

Component	Estimated counts in 2σ ROI		Difference
	Energy + MST	Energy + Standoff distance	
^{232}Th	57.6	75.6	−26.9 %
^{238}U	57.5	51.3	11.4 %
^{60}Co	45.9	48.3	−5.1 %
^{137}Xe	9.4	10.5	−11.0 %
Total background	233.7	242.0	−3.5 %
$0\nu\beta\beta$	5.0	12.0	

Table 8.3.: Backgrounds in the 2σ energy ROI for the total (SS + MS) data set as estimated by the energy + MST metric (second column) and energy + standoff distance (third column) fits.

9. Conclusions and outlook

During the course of this work, the event reconstruction for EXO-200, which was discussed in chapter 4, has been largely restructured and put into a more modular design. This allowed for a more organized way of testing different signal finding and signal discrimination techniques. The work on this helped in better understanding the various systematic effects associated with each of the event reconstruction steps. Another focus during this dissertation was to improve the way the position of energy depositions is reconstructed. To achieve this, an algorithm was developed that pairs the different signal types according to a pairing likelihood, which is based on distribution functions that were derived from data. These improvements in the three dimensional event position reconstruction, and studies of its efficiencies and accuracies narrowed down the systematic effects related to event position reconstruction and resulted in a smaller uncertainty in the size of the fiducial volume, which was previously the main systematic error.

Much work has also gone into methods to further process the reconstructed data and into the analysis of the final data set. The introduction of a new processing scheme and format for the corrected and calibrated data, enabled a more flexible way of placing event selection cuts for the final data set and simplified the determination of the event selection efficiencies presented in chapter 7. Contributions have also been made to studies of systematic effects that were presented in chapters 7 and 8, such as the error in the size of the fiducial volume, or the effect of background PDF shape uncertainties.

The analysis that was presented in the previous chapters was a collaborative effort that poses a 90 % lower limit on the $0\nu\beta\beta$ half-life of ^{136}Xe of 1.1×10^{25} yr, first published in [15]. The limit on the half-life corresponds to an upper limit on the effective Majorana neutrino mass of 190 meV to 450 meV, depending on the choice of nuclear matrix element. This is summarized in figure 9.1.

It was found that effects of drift charge diffusion on the event multiplicity can not be neglected on the order of the detector's resolution, and the transverse diffusion coefficient for EXO-200 has been approximately determined to be around $80 \text{ cm}^2/\text{s}$. A dedicated study will be necessary to measure the exact transverse diffusion coefficient for liquid xenon at the conditions of the EXO-200 TPC.

Finally, a finer graded description of the event multiplicity has been developed during the course of this work. The additional information guides the fit to obtain better background discrimination, and in a first re-analysis of the data set this yielded a slightly better lower limit on the $0\nu\beta\beta$ half-life of ^{136}Xe of 1.4×10^{25} yr. This is a work in progress, and additional studies need to be performed to quantify the systematic effects introduced by using this metric as an observable in the fit and the addition of drift charge carrier diffusion to the simulation. Additionally, the effect on the signal efficiency in the region of interest of the residual shape disagreement in the distribution of this metric




Figure 9.1.: Confidence bands for the effective Majorana neutrino mass vs the mass of the lightest neutrino. The red band marks the limit on the effective Majorana neutrino mass obtained in this analysis. The limit is a band instead of a line due to the uncertainty in the calculations of the nuclear matrix elements. The confidence bands have been calculated using the values in table 2.1.

between simulation and data needs to be studied, like it was quantified for the energy and standoff distance distributions. Usage of the multiplicity metric in addition to the standoff distance metric will require a substantial increase in computing time to produce the background PDFs in three dimensions. Nevertheless this seems like a viable way to reach better signal sensitivity for future analyses.

A. Clustering parameters




Figure A.1.: Time difference vs channel distance of V-wire signals from the signal with the largest amplitude. Due to electrostatic shielding by wires close to the charge drift path, signals on wires further away from the drift path get reconstructed to earlier times.

The parameters used in the clustering procedure are all empirical and based on studies done in low background as well as source calibration data. Figure A.3 shows the distribution of V-wire signal times versus their distance (in units of channels) from the largest signal. It is apparent that signals further away from the central signal get reconstructed at earlier times. This is addressed to electrostatic shielding by wires closer to the charge drift path and to the way the signals are reconstructed. The trend is approximately linear in the distance from the central signal. A fit of this trend gives:

$$t_i = t_c - 2.97 \mu\text{s} |\Delta_{ic}|, \quad (\text{A.1})$$

where t_i is the time of the individual V-wire signal, t_c is the time of the signal with largest amplitude and Δ_{ic} is the distance (in units of channels) of the individual signal from the signal with the largest amplitude.

The top plot in figure A.2 shows the summed V-wire signal amplitude over the summed U-wire signal amplitude per event. The profile of this histogram is used to extract a

A. Clustering parameters

linear relationship between U- and V-signal amplitudes:

$$\hat{E}_V(E_U) = \begin{cases} 0 \text{ ADC}_V & E_U < -b/a \\ aE_U + b & E_U \geq -b/a \end{cases} \quad (\text{A.2})$$

where \hat{E}_V is the estimated V-signal amplitude given a U-signal amplitude E_U , and $a = 0.2378 \text{ ADC}_V/\text{ADC}_U$ and $b = -30.79 \text{ ADC}_V$ have been determined from the profile by a fit. The width of the distribution is modeled as

$$\sigma_E(E_U) = \begin{cases} c & E_U < 350 \text{ ADC}_U \\ dE_U + e\sqrt{E_U} & E_U \geq 350 \text{ ADC}_U \end{cases}, \quad (\text{A.3})$$

where the parameters $c = 20.22 \text{ ADC}_V$, $d = 0.0101 \text{ ADC}_V/\text{ADC}_U$ and $e = 0.892 \text{ ADC}_V/\text{ADC}_U^{1/2}$ are again extracted by a fit. \hat{E}_V and σ_E are used to construct the PDF which relates U-signal amplitudes to V-signal amplitudes:

$$f_{\text{amplitude}}(E_U, E_V, Z) = \begin{cases} A & |Z| > 160 \text{ mm} \\ B \exp \left[-1/2 \left(\frac{\hat{E}_V(E_U) - E_V}{\sigma_E(E_U)} \right)^2 \right] & |Z| \leq 160 \text{ mm} \end{cases}, \quad (\text{A.4})$$

where B is a normalization constant. The PDF is just a constant A for $|Z| > 160 \text{ mm}$, because the amplitude relationship becomes Z -dependent in this region, as shown in the bottom plot of figure A.2. This essentially has the effect that the relationship between V- and U-signal amplitudes is not applied in the clustering procedure for values of $|Z| > 160 \text{ mm}$.

It was found that the time of only the largest V-wire signal in a bundle is the best estimator of the time at which the charge cloud passes the induction plane. Figure A.3 shows the time difference of the largest V-wire signal from the U-wire signal of events that contain only one U-wire signal. The time difference shows a Z -dependence, which is modeled as

$$\delta_t(Z) = \begin{cases} 0 \mu\text{s} & |Z| \leq 185.2 \text{ mm} \\ a_t + b_t Z' + c_t Z'^2 + d_t Z'^3 & 185.2 \text{ mm} < |Z| \leq 194.1 \text{ mm} \\ 3 \mu\text{s} & |Z| > 194.1 \text{ mm} \end{cases}, \quad (\text{A.5})$$

where $Z' = |Z| - 190 \text{ mm}$ and $a_t = 2.73 \mu\text{s}$, $b_t = 0.55 \mu\text{s mm}^{-1}$, $c_t = -0.065 \mu\text{s/mm}^2$ and $d_t = -0.013 \mu\text{s/mm}^3$ are obtained from a fit to the histogram. This is used to model the time difference PDF in the following way:

$$f_{\text{time}}(t_U, t_V, Z) = C \exp \left[-1/2 \left(\frac{t_U - t_V - \delta_t(Z)}{\sigma_t} \right)^2 \right], \quad (\text{A.6})$$

where C again is a normalization constant and σ_t is a constant width extracted from the histogram.




Figure A.2.: Relationship between U-signal amplitude and V-signal amplitude. The top plot shows the summed amplitudes of all V-wire signals vs the summed amplitude of all U-wire signals. The data is from a calibration source run. The linear trend is fitted and used to relate V-signal amplitudes to U-signal amplitudes. The bottom plot shows the summed V-signal amplitudes minus the scaled sum over U-amplitudes. The amplitude relationship is independent of Z for values of $|Z| < 160$ mm.




Figure A.3.: Time difference between the V-wire signal with the largest amplitude and the U-wire signal of events that contain only one U-wire signal.

B. MST metric for different diffusion parameters

Figure B.1 compares the distribution of the MST metric at low values as observed in data from the ^{228}Th calibration source to simulations of the calibration source with different values of the transverse electron diffusion during drift. The calibration source is placed at the cathode plane to make full use of the maximum possible drift length. The agreement especially in the one-wire (0 mm) and two-neighbouring-wires (9 mm) bin gradually improves with increasing diffusion coefficient. The overall best agreement seems to be reached at a diffusion coefficient of around $80\text{ cm}^2/\text{s}$.

B. MST metric for different diffusion parameters

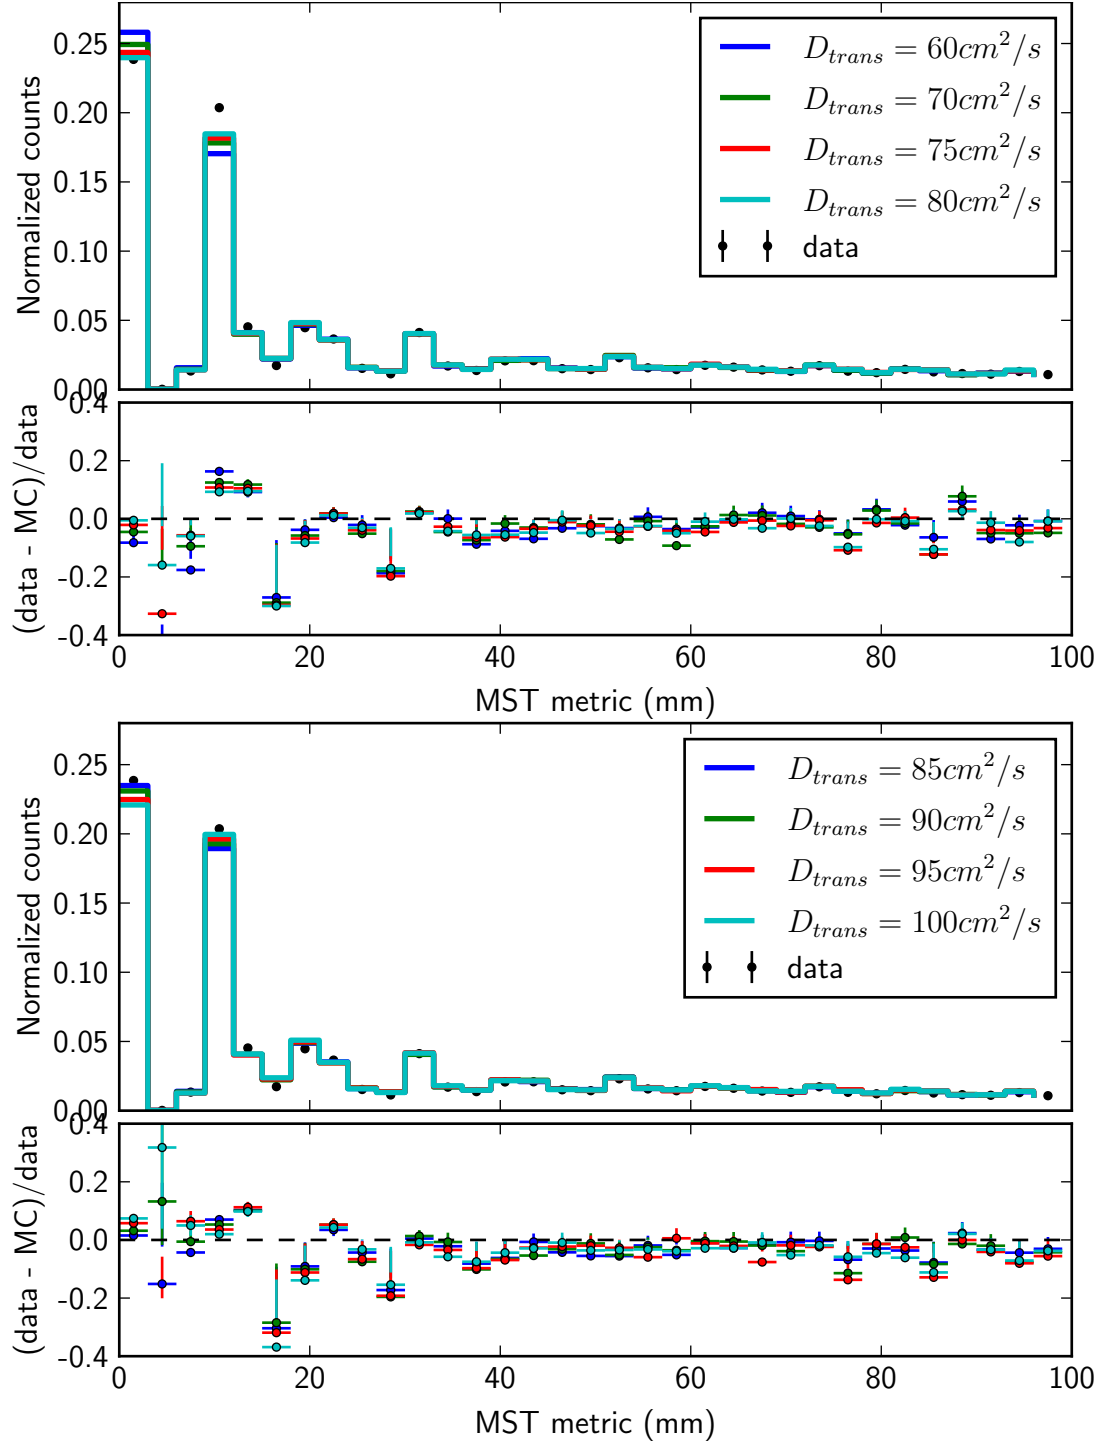


Figure B.1.: Comparison of MST metric distributions between data and simulation for different simulated transverse diffusion coefficients. The data are split between two panels for better visibility. Sub panels show the normalized residuals defined as (simulation - data) / data.

Bibliography

- [1] G. Bellini et al. “Neutrinos from the primary proton–proton fusion process in the Sun”. In: *Nature* 512.7515 (Aug. 27, 2014), pp. 383–386. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature13702.
- [2] John Bahcall, Marc Kamionkowski, and Alberto Sirlin. “Solar neutrinos: Radiative corrections in neutrino-electron scattering experiments”. In: *Physical Review D* 51.11 (June 1995), pp. 6146–6158. ISSN: 0556-2821. DOI: 10.1103/PhysRevD.51.6146.
- [3] C. L. Cowan et al. “Detection of the Free Neutrino: a Confirmation”. In: *Science* 124.3212 (July 20, 1956), pp. 103–104. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.124.3212.103.
- [4] G. Danby et al. “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos”. In: *Physical Review Letters* 9.1 (July 1962), pp. 36–44. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.9.36.
- [5] K. Kodama et al. “Observation of tau neutrino interactions”. In: *Physics Letters B* 504.3 (Apr. 2001), pp. 218–224. ISSN: 03702693. DOI: 10.1016/S0370-2693(01)00307-0.
- [6] Raymond Davis, Don S. Harmer, and Kenneth C. Hoffman. “Search for Neutrinos from the Sun”. In: *Physical Review Letters* 20.21 (May 1968), pp. 1205–1209. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.20.1205.
- [7] Y. Ashie et al. “Evidence for an Oscillatory Signature in Atmospheric Neutrino Oscillations”. In: *Physical Review Letters* 93.10 (Sept. 2004). ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.93.101801.
- [8] Q. Ahmad et al. “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory”. In: *Physical Review Letters* 89.1 (June 2002). ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.89.011301.
- [9] Ch. Weinheimer. “KATRIN, a next generation tritium beta decay experiment in search for the absolute neutrino mass scale”. In: *Progress in Particle and Nuclear Physics* 48.1 (Jan. 2002), pp. 141–150. ISSN: 01466410. DOI: 10.1016/S0146-6410(02)00120-5.
- [10] Planck Collaboration et al. “Planck 2013 results. XVI. Cosmological parameters”. In: *arXiv:1303.5076 [astro-ph]* (Mar. 20, 2013). (Visited on 05/23/2014).

- [11] M. Fukugita and T. Yanagida. “Baryogenesis without grand unification”. In: *Physics Letters B* 174.1 (June 1986), pp. 45–47. ISSN: 03702693. DOI: 10.1016/0370-2693(86)91126-3.
- [12] N. Ackerman et al. “Observation of Two-Neutrino Double-Beta Decay in Xe with the EXO-200 Detector”. In: *Physical Review Letters* 107.21 (Nov. 2011). ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.107.212501.
- [13] J. B. Albert et al. “Improved measurement of the double beta decay half-life of ^{136}Xe with the EXO-200 detector”. In: *Physical Review C* 89.1 (Jan. 28, 2014). ISSN: 0556-2813, 1089-490X. DOI: 10.1103/PhysRevC.89.015502.
- [14] M. Auger et al. “Search for Neutrinoless Double-Beta Decay in Xe with EXO-200”. In: *Physical Review Letters* 109.3 (July 2012). ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.109.032505.
- [15] J. B. Albert et al. “Search for Majorana neutrinos with the first two years of EXO-200 data”. In: *Nature* (June 4, 2014). ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature13432.
- [16] R. N. Mohapatra. *Massive neutrinos in physics and astrophysics*. In collab. with P. B. Pal. 3rd ed. World Scientific lecture notes in physics v. 72. River Edge, N.J: World Scientific, 2004. 451 pp. ISBN: 9812380701.
- [17] Felix Boehm. *Physics of massive neutrinos*. In collab. with Petr Vogel. 2nd ed. Cambridge [England] ; New York, NY, USA: Cambridge University Press, 1992. 249 pp. ISBN: 0521418240.
- [18] M. Fukugita. *Physics of neutrinos and applications to astrophysics*. In collab. with T. Yanagida. Texts and monographs in physics. Berlin ; New York: Springer, 2003. 593 pp. ISBN: 3540438009.
- [19] Palash B. Pal. “Dirac, Majorana, and Weyl fermions”. In: *American Journal of Physics* 79.5 (2011), p. 485. ISSN: 00029505. DOI: 10.1119/1.3549729.
- [20] S. M. Bilenky. “Majorana Neutrino Mixing”. In: *Journal of Physics G: Nuclear and Particle Physics* 32.5 (May 1, 2006), R127–R149. ISSN: 0954-3899, 1361-6471. DOI: 10.1088/0954-3899/32/5/R01. arXiv: hep-ph/0511227. URL: <http://arxiv.org/abs/hep-ph/0511227>.
- [21] S. M. Bilenky. “Neutrinos: A Brief Review”. In: *Modern Physics Letters A* 19.33 (Oct. 30, 2004), pp. 2451–2477. ISSN: 0217-7323, 1793-6632. DOI: 10.1142/S02177323040159.
- [22] Andreas Aste. “A Direct Road to Majorana Fields”. In: *Symmetry* 2.4 (Oct. 27, 2010), pp. 1776–1809. ISSN: 2073-8994. DOI: 10.3390/sym2041776.
- [23] M. Goldhaber, L. Grodzins, and A. Sunyar. “Helicity of Neutrinos”. In: *Physical Review* 109.3 (Feb. 1958), pp. 1015–1017. ISSN: 0031-899X, 1536-6065. DOI: 10.1103/PhysRev.109.1015.
- [24] F. P. An et al. “Observation of Electron-Antineutrino Disappearance at Daya Bay”. In: *Physical Review Letters* 108.17 (Apr. 2012). ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.108.171803.

- [25] V. N. Aseev et al. “Upper limit on the electron antineutrino mass from the Troitsk experiment”. In: *Physical Review D* 84.11 (Dec. 2011). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.84.112003.
- [26] Ch Kraus et al. “Final results from phase II of the Mainz neutrino mass search in tritium beta decay”. In: *The European Physical Journal C* 40.4 (Apr. 2005), pp. 447–468. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s2005-02139-7.
- [27] Peter Minkowski. “ μ to e gamma at a rate of one out of 109 muon decays?” In: *Physics Letters B* 67.4 (Apr. 1977), pp. 421–428. ISSN: 03702693. DOI: 10.1016/0370-2693(77)90435-X.
- [28] M. Magg and Ch. Wetterich. “Neutrino mass problem and gauge hierarchy”. In: *Physics Letters B* 94.1 (July 1980), pp. 61–64. ISSN: 03702693. DOI: 10.1016/0370-2693(80)90825-4.
- [29] R. Foot et al. “See-saw neutrino masses induced by a triplet of leptons”. In: *Zeitschrift für Physik C Particles and Fields* 44.3 (Sept. 1989), pp. 441–444. ISSN: 0170-9739, 1434-6052. DOI: 10.1007/BF01415558.
- [30] D. V. Forero, M. Tórtola, and J. W. F. Valle. “Global status of neutrino oscillation parameters after Neutrino-2012”. In: *Physical Review D* 86.7 (Oct. 2012). ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.86.073012.
- [31] J. Barea, J. Kotila, and F. Iachello. “Nuclear matrix elements for double-beta decay”. In: *Physical Review C* 87.1 (Jan. 14, 2013). ISSN: 0556-2813, 1089-490X. DOI: 10.1103/PhysRevC.87.014315.
- [32] NEMO Collaboration, A. S. Barabash, and V. B. Brudanin. “Investigation of double-beta decay with the NEMO-3 detector”. In: *Physics of Atomic Nuclei* 74.2 (Feb. 2011), pp. 312–317. ISSN: 1063-7788, 1562-692X. DOI: 10.1134/S1063778811020062.
- [33] The GERDA Collaboration et al. “Measurement of the half-life of the two-neutrino double beta decay of ^{76}Ge with the GERDA experiment”. In: *Journal of Physics G: Nuclear and Particle Physics* 40.3 (Mar. 1, 2013), p. 035110. ISSN: 0954-3899, 1361-6471. DOI: 10.1088/0954-3899/40/3/035110.
- [34] M. Agostini et al. “Results on Neutrinoless Double-Beta Decay of ^{76}Ge from Phase I of the GERDA Experiment”. In: *Physical Review Letters* 111.12 (Sept. 2013). ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.111.122503.
- [35] F. Danevich et al. “New results of ^{116}Cd double beta decay study with $^{116}\text{CdWO}_4$ scintillators”. In: *Physical Review C* 62.4 (Sept. 2000). ISSN: 0556-2813, 1089-490X. DOI: 10.1103/PhysRevC.62.045501.
- [36] E. Andreotti et al. “ ^{130}Te neutrinoless double-beta decay with CUORICINO”. In: *Astroparticle Physics* 34.11 (June 2011), pp. 822–831. ISSN: 09276505. DOI: 10.1016/j.astropartphys.2011.02.002.
- [37] A. Gando et al. “Measurement of the double-beta decay half-life of ^{136}Xe with the KamLAND-Zen experiment”. In: *Physical Review C* 85.4 (Apr. 2012). ISSN: 0556-2813, 1089-490X. DOI: 10.1103/PhysRevC.85.045504.

- [38] A. Gando et al. “Limit on Neutrinoless Double-Beta Decay of ^{136}Xe from the First Phase of KamLAND-Zen and Comparison with the Positive Claim in ^{76}Ge ”. In: *Physical Review Letters* 110.6 (Feb. 7, 2013). ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.110.062502.
- [39] J. Schechter and J. Valle. “Neutrinoless double-beta decay in $\text{SU}(2)\times\text{U}(1)$ theories”. In: *Physical Review D* 25.11 (June 1982), pp. 2951–2954. ISSN: 0556-2821. DOI: 10.1103/PhysRevD.25.2951.
- [40] J. Kotila and F. Iachello. “Phase-space factors for double-beta decay”. In: *Physical Review C* 85.3 (Mar. 2012). ISSN: 0556-2813, 1089-490X. DOI: 10.1103/PhysRevC.85.034316.
- [41] Matthew Redshaw et al. “Mass and Double-Beta-Decay Q Value of Xe^{136} ”. In: *Physical Review Letters* 98.5 (Feb. 2007). ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.98.053003.
- [42] R.G Helmer and C van der Leun. “Recommended standards for gamma-ray energy calibration (1999)”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 450.1 (Aug. 2000), pp. 35–70. ISSN: 01689002. DOI: 10.1016/S0168-9002(00)00252-7.
- [43] International Atomic Energy Agency. *Update of X ray and gamma ray decay data standards for detector calibration and other applications*. Ed. by M. Herman and A. L. Nichols. Vienna, Austria: International Atomic Energy Agency, 2007. 2 pp. ISBN: 9201136064.
- [44] W.B. Streett, L.S. Sagan, and L.A.K. Staveley. “An experimental study of the equation of state of liquid xenon”. In: *The Journal of Chemical Thermodynamics* 5.5 (Sept. 1973), pp. 633–650. ISSN: 00219614. DOI: 10.1016/S0021-9614(73)80004-7.
- [45] M.J. Berger et al. *XCOM: Photon Cross Sections Database*. URL: <http://www.nist.gov/pml/data/xcom/index.cfm>.
- [46] M.J. Berger et al. *ESTAR*. URL: <http://physics.nist.gov/PhysRefData/Star/Text/ESTAR.html>.
- [47] Tadayoshi Doke. “Recent developments of liquid xenon detectors”. In: *Nuclear Instruments and Methods in Physics Research* 196.1 (May 1982), pp. 87–96. ISSN: 01675087. DOI: 10.1016/0029-554X(82)90621-8.
- [48] U. Asaf and I. Steinberger. “Photoconductivity and electron transport parameters in liquid and solid xenon”. In: *Physical Review B* 10.10 (Nov. 1974), pp. 4464–4468. ISSN: 0556-2805. DOI: 10.1103/PhysRevB.10.4464.
- [49] T. Takahashi et al. “Average energy expended per ion pair in liquid xenon”. In: *Physical Review A* 12.5 (Nov. 1975), pp. 1771–1775. ISSN: 0556-2791. DOI: 10.1103/PhysRevA.12.1771.

- [50] E. Aprile and T. Doke. “Liquid xenon detectors for particle physics and astrophysics”. In: *Reviews of Modern Physics* 82.3 (July 2010), pp. 2053–2097. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.82.2053.
- [51] T. Doke et al. “Estimation of Fano factors in liquid argon, krypton, xenon and xenon-doped liquid argon”. In: *Nuclear Instruments and Methods* 134.2 (Apr. 1976), pp. 353–357. ISSN: 0029554X. DOI: 10.1016/0029-554X(76)90292-5.
- [52] K. Twelker et al. “An apparatus to manipulate and identify individual Ba ions from bulk liquid Xe”. In: *Review of Scientific Instruments* 85.9 (Sept. 2014), p. 095114. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/1.4895646.
- [53] Thomas Brunner et al. “An RF-only ion-funnel for extraction from high-pressure gases”. In: *arXiv:1412.1144 [nucl-ex, physics:physics]* (Dec. 2, 2014). arXiv: 1412.1144.
- [54] B. Mong et al. “Spectroscopy of Ba and Ba⁺ deposits in solid xenon for barium tagging in nEXO”. In: *arXiv:1410.2624 [hep-ex, physics:nucl-ex, physics:physics]* (Oct. 9, 2014). arXiv: 1410.2624.
- [55] D.S. Leonard et al. “Systematic study of trace radioactive impurities in candidate construction materials for EXO-200”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 591.3 (July 2008), pp. 490–509. ISSN: 01689002. DOI: 10.1016/j.nima.2008.03.001.
- [56] R. Neilson et al. “Characterization of large area APDs for the EXO-200 detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 608.1 (Sept. 2009), pp. 68–75. ISSN: 01689002. DOI: 10.1016/j.nima.2009.06.029.
- [57] M. Yamashita et al. “Scintillation response of liquid Xe surrounded by PTFE reflector for gamma rays”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 535.3 (Dec. 2004), pp. 692–698. ISSN: 01689002. DOI: 10.1016/j.nima.2004.06.168.
- [58] M.J. Terry et al. “The densities of liquid argon, krypton xenon, oxygen, nitrogen, carbon monoxide methane, and carbon tetrafluoride along the orthobaric liquid curve”. In: *The Journal of Chemical Thermodynamics* 1.4 (July 1969), pp. 413–424. ISSN: 00219614. DOI: 10.1016/0021-9614(69)90072-X.
- [59] V. A. Rabinovich and L. S. Veklsler. *Thermophysical Properties of Matter and Substances*. Vol. 2. New Delhi: Amerind Publishing, 1974.
- [60] Gabor Jancso and W. Alexander Van Hook. “Condensed phase isotope effects”. In: *Chemical Reviews* 74.6 (Dec. 1974), pp. 689–750. ISSN: 0009-2665, 1520-6890. DOI: 10.1021/cr60292a004.
- [61] F. LePort et al. “A magnetically driven piston pump for ultra-clean applications”. In: *Review of Scientific Instruments* 82.10 (2011), p. 105114. ISSN: 00346748. DOI: 10.1063/1.3653391.

- [62] A. Dobi et al. “A xenon gas purity monitor for EXO”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 659.1 (Dec. 2011), pp. 215–228. ISSN: 01689002. DOI: 10.1016/j.nima.2011.09.017.
- [63] V. Chisté, M. M. Bé, and C. Dulieu. “Evaluation of decay data of radium-226 and its daughters”. In: EDP Sciences, 2007. DOI: 10.1051/ndata:07122.
- [64] S. Agostinelli et al. “Geant4—a simulation toolkit”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (July 2003), pp. 250–303. ISSN: 01689002. DOI: 10.1016/S0168-9002(03)01368-8. URL: <http://geant4.cern.ch>.
- [65] *Maxwell*. URL: <http://www.ansys.com/Products/Simulation+Technology/Systems+&+Multiphysics/Multiphysics+Enabled+Products/ANSYS+Maxwell>.
- [66] W. Shockley. “Currents to Conductors Induced by a Moving Point Charge”. In: *Journal of Applied Physics* 9.10 (1938), p. 635. ISSN: 00218979. DOI: 10.1063/1.1710367.
- [67] S. Ramo. “Currents Induced by Electron Motion”. In: *Proceedings of the IRE* 27.9 (Sept. 1939), pp. 584–585. ISSN: 0096-8390. DOI: 10.1109/JRPROC.1939.228757.
- [68] S. Herrin. internal EXO-200 presentation.
- [69] Reinhard Diestel. *Graph theory*. 4th ed. Graduate texts in mathematics 173. Heidelberg ; New York: Springer, 2010. 436 pp. ISBN: 9783642142789.
- [70] Rene Brun and Fons Rademakers. “ROOT — An object oriented data analysis framework”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 389.1 (Apr. 1997), pp. 81–86. ISSN: 01689002. DOI: 10.1016/S0168-9002(97)00048-X. URL: <http://root.cern.ch>.
- [71] L. J. Wen. internal EXO-200 presentation.
- [72] A. Göök et al. “Application of the Shockley–Ramo theorem on the grid inefficiency of Frisch grid ionization chambers”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 664.1 (Feb. 2012), pp. 289–293. ISSN: 01689002. DOI: 10.1016/j.nima.2011.10.052.
- [73] Clayton G. Davis. “A Search for the Neutrinoless Double Beta Decay of Xenon-136 with Improved Sensitivity from Denoising”. PhD thesis. College Park: University of Maryland, 2014.
- [74] E. Conti et al. “Correlated fluctuations between luminescence and ionization in liquid xenon”. In: *Physical Review B* 68.5 (Aug. 2003). ISSN: 0163-1829, 1095-3795. DOI: 10.1103/PhysRevB.68.054201.
- [75] E. Aprile et al. “Observation of anticorrelation between scintillation and ionization for MeV gamma rays in liquid xenon”. In: *Physical Review B* 76.1 (July 2007). ISSN: 1098-0121, 1550-235X. DOI: 10.1103/PhysRevB.76.014115.

- [76] C. Licciardi. internal EXO-200 presentation.
- [77] Glenn F. Knoll. *Radiation detection and measurement*. 4th ed. Hoboken, N.J: John Wiley, 2010. 830 pp. ISBN: 9780470131480.
- [78] N. P. Archer and G. L. Keech. “DIRECT MEASUREMENTS OF THE HALF-LIVES OF SOME SHORT-LIVED XENON AND CESIUM ISOTOPES”. In: *Canadian Journal of Physics* 44.8 (Aug. 1966), pp. 1823–1845. ISSN: 0008-4204, 1208-6045. DOI: 10.1139/p66-154.
- [79] F. James and M. Roos. “Minuit - a system for function minimization and analysis of the parameter errors and correlations”. In: *Computer Physics Communications* 10.6 (Dec. 1975), pp. 343–367. ISSN: 00104655. DOI: 10.1016/0010-4655(75)90039-9.
- [80] Wolfgang A. Rolke, Angel M. López, and Jan Conrad. “Limits and confidence intervals in the presence of nuisance parameters”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 551.2 (Oct. 2005), pp. 493–503. ISSN: 01689002. DOI: 10.1016/j.nima.2005.05.068.
- [81] S. S. Wilks. “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses”. In: *The Annals of Mathematical Statistics* 9.1 (Mar. 1938). ISSN: 00034851.
- [82] G. K. Schenter and P. Vogel. “A Simple Approximation of the Fermi Function in Nuclear Beta Decay”. In: *Nuclear Science and Engineering* 83.3 (Mar. 1983), pp. 393–396.
- [83] Steven Herrin. “DOUBLE BETA DECAY IN XENON-136: MEASURING THE NEUTRINO-EMITTING MODE AND SEARCHING FOR MAJORON-EMITTING MODES”. PhD thesis. Stanford University, 2013.
- [84] G. Battistoni et al. “The FLUKA code: Description and benchmarking”. In: *Proceedings of the Hadronic Shower Simulation Workshop 2006*. AIP Conference Proceeding 896, 31-49, (2007). Fermilab, Sept. 2006.
- [85] A. Ferrari, P. R. Sala, and A. Fasso. “FLUKA: a multi-particle transport code”. In: (Oct. 2005).
- [86] W. B. Wilson et al. *SOURCES 4A: A Code for Calculating (α,n), Spontaneous Fission, and Delayed Neutron Sources and Spectra*. Sept. 1999.
- [87] J. Walton. internal EXO-200 presentation.
- [88] R.G. Helmer, R.J. Gehrke, and R.C. Greenwood. “gamma-ray energies for 40K, 108Ag and the 226Ra decay chain”. In: *Nuclear Instruments and Methods* 166.3 (Dec. 1979), pp. 547–549. ISSN: 0029554X. DOI: 10.1016/0029-554X(79)90548-2.
- [89] R. Killick. internal EXO-200 presentation.

List of Figures

2.1.	Comparison of fermion masses.	7
2.2.	Neutrino mass hierarchy.	9
2.3.	Beta decay scheme for isobars with mass $A = 136$	11
2.4.	Feynman diagram of double beta decay.	12
2.5.	Schematic energy spectrum of double beta decay.	14
2.6.	Effective neutrino mass $m_{\beta\beta}$ vs mass of the lightest neutrino.	15
3.1.	Radioactive penetration lengths in liquid xenon.	19
3.2.	Schematic of the EXO-200 TPC.	22
3.3.	Photograph of TPC during insertion into the cryostat (right).	23
3.4.	Photograph of one half of the TPC.	24
3.5.	Photograph of one anode plane of the TPC.	24
3.6.	Photograph of one light collection plane of the TPC.	26
3.7.	Schematic of the read out electronics.	27
3.8.	Simplified schematic of the xenon system.	30
3.9.	Schematic of the detector surroundings.	31
3.10.	Calibration system.	32
3.11.	Simulation model of the TPC.	33
4.1.	Example raw data signal waveforms.	36
4.2.	Signal finding process.	37
4.3.	Extraction of signal parameters.	38
4.4.	Signal reconstruction efficiencies.	40
4.5.	Empirical PDFs used in 3D position reconstruction.	42
4.6.	Position reconstruction accuracy.	45
4.7.	3D reconstruction efficiencies.	47
5.1.	Example of a minimum spanning tree.	50
5.2.	MST metric distribution for ^{60}Co calibration source.	51
5.3.	MST metric distribution in simulation and data.	52
5.4.	MST metric distribution of different background components.	54
6.1.	Schematic of data processing.	57
6.2.	Measurement of electron lifetime.	59
6.3.	Time variation of free electron lifetime.	60
6.4.	Inefficiency of the shielding grid.	61
6.5.	Combination of charge and light signal.	63

6.6.	Comparison of charge-only, scintillation-only, and combined energy spectrum.	64
6.7.	Detector energy resolution.	66
7.1.	Schematic of fiducial volume.	71
7.2.	Backgrounds from anode planes.	72
7.3.	Scintillation vs charge spectrum of the data set.	73
8.1.	Example background PDFs used in the fit.	81
8.2.	Spectral agreement between simulation and data.	84
8.3.	Event multiplicity of calibration sources.	85
8.4.	Residual spectral differences between simulation and calibration source data.	87
8.5.	Example of an unskewed PDF.	88
8.6.	Bias in ROI counts due to spectral disagreements between simulation and data.	89
8.7.	Fit of the PDFs to the low background data set.	90
8.8.	Profile likelihood on number of $0\nu\beta\beta$ counts.	93
8.9.	Coverage and sensitivity studies.	95
8.10.	Fit result with MST metric.	97
8.11.	Profile likelihood in number of $0\nu\beta\beta$ counts with MST metric.	99
9.1.	Limit on effective Majorana neutrino mass and lightest neutrino mass. . .	102
A.1.	Time structure of V-wire signals.	103
A.2.	Relationship between U-signal amplitude and V-signal amplitude. . . .	105
A.3.	Time difference between V-wire and U-wire signals.	106
B.1.	Comparison of MST metric distributions between data and simulation for different diffusion coefficients.	108

List of Tables

2.1. Neutrino mixing parameters.	10
2.2. Half-lives of double beta decay isotopes.	13
2.3. Properties of double beta decay isotopes.	15
3.1. Electronic signal shaping times.	28
3.2. Activities of EXO-200 calibration sources.	32
6.1. Shielding grid inefficiency parameters.	60
6.2. Energy resolution parameters.	67
6.3. Energy region of interest (ROI).	68
7.1. Summary of detector run time and vetoed times.	70
7.2. Signal efficiencies and associated systematic errors.	75
8.1. List of parameters floating in the fit.	91
8.2. Summary of backgrounds in the ROI.	94
8.3. Backgrounds in ROI estimated by fit with MST metric.	100