First Search for Majorana Neutrinos with LEGEND-200

George Geoffrey Marshall

University College London

Physics and Astronomy

Submitted to University College London (UCL) in partial fulfilment of the requirements for the award of the degree of Doctor of Philosophy.

Primary supervisor: Matteo Agostini

Secondary supervisor: Jonathan Butterworth

Examining committee: **Andreas Korn and Cheryl Patrick**

Thesis submission date: **30th September 2024**

Declaration

I, George Geoffrey Marshall, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

George Geoffrey Marshall
London, United Kingdom
30th September 2024

Abstract

The discovery of neutrinoless double beta decay ($0\nu\beta\beta$) would revolutionise our understanding of physics. It would prove that neutrinos are their own antiparticles, lepton number is not conserved and matter can be produced without balancing antimatter. The LEGEND collaboration is pursuing a discovery-oriented $0\nu\beta\beta$ search using hundreds of high-purity germanium detectors made of Ge-76, a $0\nu\beta\beta$ candidate isotope. The running stage, LEGEND-200, is projected to provide the world's best $0\nu\beta\beta$ discovery power for the next 5 years.

This thesis work focuses on the design and implementation of the LEGEND-200 analysis pipeline, and its application both to the characterisation of novel detectors and its first physics data. It shows how the design performance of LEGEND-200 has been largely met, in particular in terms of the resolution and stability of the energy scale, livetime, and background rates. The background across the array was thoroughly characterised and proved uniform within the statistical uncertainties. This enabled a single parameter treatment in the $0\nu\beta\beta$ statistical analysis, significantly reducing its complexity.

Building on this optimised analysis pipeline and the detector performance characterisation and background contributions extracted through it, a blinded $0\nu\beta\beta$ analysis was enabled on the first LEGEND-200 data set corresponding to an exposure of 48.3 kg yr. In addition to performing the entire data reduction and sample preparation, the work presented in this thesis also covers the definition of the main sources of systematic uncertainties and key performance parameters entering the collaboration's statistical analysis. A background index of $5.3 \pm 2.2 \times 10^{-4}$ counts / (keV kg yr) was measured proving that LEGEND-200 will deliver its physics programme. Combining this data set with those of past Ge-76 experiments, a combined constraint on the $0\nu\beta\beta$ half life of $> 1.9 \times 10^{26}$ yr at 90% CL was derived. This is the world leading limit in Ge-76 and one of the most sensitive ones in the field.

Impact Statement

The discovery of neutrinoless double beta decay would be revolutionary for particle physics, nuclear physics and cosmology. It would be evidence of beyond Standard Model physics providing proof of both the Majorana nature of neutrinos, of lepton number violation and of a purely matter creating process explaining why the universe is here rather than empty. Due to the extreme difficulty in detecting this decay the experiments are carried out in large collaborations with the focus of this thesis on my work for the LEGEND experiment which consists of more than 200 scientists across more than 50 institutions.

The focus of this thesis work is on the design, development and implementation of a sophisticated analysis pipeline addressing the challenging reliability needs of a discovery experiment along with its application to perform the first search for $0\nu\beta\beta$ with LEGEND, achieving world-leading sensitivity to constrain this process. The limit obtained is the most stringent limit in Ge-76 and only behind KamLAND-ZEN globally. This limit is just the first step in the LEGEND-200 experiment with future limits passing into previously unexplored parameter space. LEGEND-200 will likely be the leading neutrinoless double beta decay experiment for the next few years and the work outlined in this thesis will form the basis of these future results. Beyond LEGEND-200 is the next generation experiment LEGEND-1000, the work outlined here will both inform the development of this exciting project and likely form the basis of the data processing for it. In particular, the work on the energy optimisation and calibration which places LEGEND-200 at the leading edge of energy resolution for these experiments, the work validating the performance of the large high purity germanium detectors which will be crucial to the success of LEGEND-1000 and the work on pulse shape analysis with novel approaches to enhance the performance of these new large detectors. LEGEND-1000 was the top project in the US Department of Energy's downselection for neutrinoless double beta decay.

Outside academia germanium detectors are used in a variety of applications related to nuclear spectroscopy, including nuclear safeguarding and medical imaging. In particular in the UK, they see use in nuclear decommissioning which is of utmost importance as older reactors are being taken offline and new reactors start to be built again. The work in this thesis, especially on the validation and

troubleshooting of large, novel geometry detectors will inform the work of the manufacturers to improve the detector design, potentially leading to more advanced detectors which require less time to make measurements and therefore lower costs. The novel signal processing developed by LEGEND as well as advanced energy reconstruction will be of wider interest.

Acknowledgements

My first and greatest thanks must go to my supervisor, Dr. Matteo Agostini. I consider myself hugely privileged and lucky to have had you as my supervisor. Though the work has of course been challenging as all PhDs are, you have always made it immensely fun and enjoyable even if we laugh too much during calls. This has also hugely been helped by the wonderful LEGEND group at UCL: David Waters, Ruben Saakyan, Will Quinn, Abi Alexander, Toby Dixon and Difei Xu. I have loved working with you all and have learned so much from you.

Of course this work would not have been possible without the wider LEGEND collaboration. There are of course too many people involved to be able to list everyone here, but I would like to highlight Luigi Pertoldi for his help throughout including many 2 a.m. slack discussions to get the data processed and analysed, and Ian Guinn and Jason Detwiler for all their help with the software design and implementation. Your expertise was invaluable.

The UCL HEP group is a truly fantastic place to work with too many awesome people to name here but to just mention a few: Sim, Mustafa, Em, Lesley (honorary), Fern, Emil and Andrius. I will miss the meandering lunchtime discussions and Friday pub sessions. I knew when joining UCL that it had a great reputation for physics, but it is also a hugely sociable and fun place to work.

On a personal note, I would like to thank my family for their unwavering support throughout; my parents Roo and Penny for always encouraging and pushing me to follow my passion, proofreading strange physics essays and humouring my weird book requests. Also my sister Imogen and brother Henry for putting up with having to live with me and my strange work hours. Finally, to my friends for all their encouragement and support.

Lastly, I would like to thank two teachers who had a huge impact on my life and putting me in the position to write this thesis. Firstly Paul Clark who lit the spark of my interest in 2010 when he visited to give a talk at my prep school on particle physics from that moment on I have been hooked. In the following years, I was lucky enough to also be taught by him at Oundle, and he was just as inspirational there. Secondly, my housemaster Adam Langsdale who went out of his way to encourage this passion. It was because of him that I did an extended project, the topic of which I chose to be on neutrinos a field I am now lucky enough to work in.

I have undoubtedly missed many people here and for that I can only apologise and express my gratitude to you all.

Contents

Table of Acronyms

Notation	Description
$0\nu\beta\beta$	Neutrinoless double beta decay.
$2\nu\beta\beta$	2 neutrino double beta decay.
ADC	Analog to digital converter.
BEGe	Broad Energy Germanium detector: a detector with a small p+ electrode developed for GERDA.
BI	Background Index.
BSM	Beyond Standard Model.
CC4	The amplifier board of the charge sensitive amplifier.
CL	Confidence Level.
COAX	Semi COAXial detector with a large p+ electrode..
DAQ	Data Acquisition.
DEP	Double Escape Peak.
DSP	Digital Signal Processing.
FCCD	Full Charge Collection Depth.
FEP	Full Energy Peak.
FWHM	Full Width Half Maximum.
HPGe	High Purity Germanium.
ICPC	Inverted Coaxial Point Contact: the newest type of Ge detector for legend.
LAr	Liquid Argon.
LMFE	Low Mass Front End.
LNGS	Laboratori Nationale di Gran Sasso. Underground laboratory in central Italy..
LQ	Late Charge cut, a pulse shape discrimination cut used for some detectors to remove surface events.
$m_{\beta\beta}$	Effective Majorana neutrino mass.
MSE	Multi site event.

Notation	Description
PEN	Poly Ethylene Naphthalate: a scintillating plastic used in LEGEND.
PMT	Photomultiplier Tube.
PPC	P+ Point Contact detector: a point contact style detector developed for MAJORANA.
PSD	Pulse Shape Discrimination.
$Q_{\beta\beta}$	Q-value of the double beta decay.
ROI	Region of Interest.
SEP	Single Escape Peak.
SiPM	Silicon Photomultiplier.
SSE	Single site event.

Chapter 1

Introduction

The discovery of neutrinoless double beta decay has the potential to revolutionise our understanding of particle physics. By finding this decay, we would determine not only the existence of a Majorana component to the neutrino mass (a long hypothesised but not yet realised new theory of fermion masses) but also lepton number violation and the first purely matter creating process leading to an explanation for the baryon asymmetry of the Universe.

In the first few chapters of this thesis I will lay out the path both in theory and in experiment to the current situation in the field with an overview of both the current theoretical understanding and experimental status. Before introducing the LEGEND experimental program with the current generation LEGEND-200 the focus of this thesis. In the following chapters I will then discuss my work on the analysis routines and data processing alongside their application to both characterisation data and LEGEND-200 data. Finally, the results of the last year of data taking will be presented in terms of the performance of the experiment, the background measured and the final constraint on the $0\nu\beta\beta$ half life.

I will start with a high level overview of the pivotal moments in the field extracted from Refs. [1, 2].

History of Neutrino Mass

In 1914, James Chadwick found that the beta decay spectrum was continuous, in contrast to the discrete lines of alpha and gamma decay [3]. This was a problem as it would mean that the conservation laws of energy and momentum were violated. A number of solutions were proposed such as Bohr's idea that the conservation laws did not apply in particle interactions, but it was Wolfgang Pauli's idea that with the passage of time proved to be correct. In a now famous letter, he put forward the idea that a new particle was being emitted carrying away some of the energy and momentum which he originally named the neutron but came to be called the neutrino (little neutral one) [4]. This came with the small issue that the neutrino

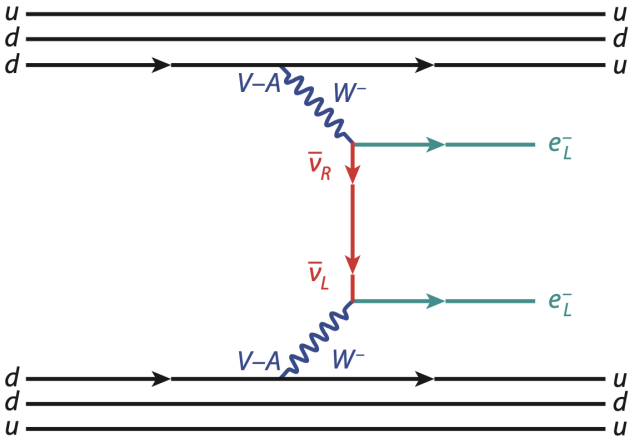

Figure 1.1: Feynman diagram for $2\nu\beta\beta$ from Ref. [9]. Here 2 neutrons decay to 2 protons emitting 2 electrons and 2 neutrinos.

was likely to be so hard to detect that it would never be observed, in fact he was so confident that he bet a case of champagne on it [5].

The idea of a neutrino really gained traction in 1933 due to Enrico Fermi who developed a theory of 4 point interactions for beta decay that included the neutrino [6, 7]. The success of this theory in explaining the beta decay spectrum led to its application to electron capture and beta plus decay but also to Maria Goeppert Mayer in 1935 predicting the existence of another more exotic decay where two neutrons decayed to two protons emitting two electrons and two neutrinos which was christened double beta decay [8] (shown in Figure 1.1).

Subsequently, in 1937, Ettore Majorana showed that this theory of beta decay was unchanged if the neutrino was its own antiparticle [10] which led to Racah proposing that beta decay could be immediately followed by electron capture if the neutrino was a Majorana particle [11]. Although he showed that this decay could occur, he caveated that an experimental observation was not possible due to the extremely low neutrino fluxes involved.

Bringing these two strands together, in 1939, Wendell Furry built on Goeppert-Mayer's double beta decay and Majorana's neutrino to propose the idea of a neutrinoless double beta decay ($0\nu\beta\beta$) where two neutrons decay to two protons emitting two electrons but no neutrinos (shown in Figure 1.2). This would be a single monochromatic peak, as opposed to the spectrum of $2\nu\beta\beta$, and could be used to ascertain the nature of the neutrino, whether it was Dirac like the other particles or something new a Majorana particle.

Up until this point, Pauli's bet was still going with the neutrino only a theoretical construct with no direct detection, but the experimental side was catching up. Perrin and Fermi suggested that its mass could be measured using the endpoint of the Tritium spectrum [6, 7, 12, 13]. This measurement was first done in the 1940s placing the first limit on the lightest neutrino mass [14], indeed this method is still

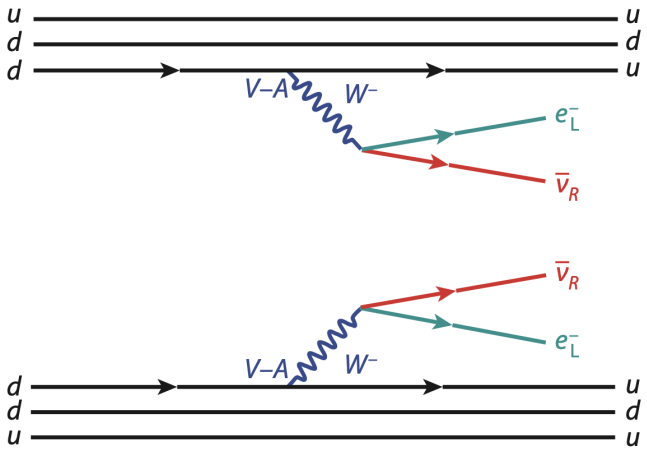

Figure 1.2: Feynman diagram for $0\nu\beta\beta$ from Ref. [9]. Here 2 neutrons decay to 2 protons emitting 2 electrons but no neutrinos.

being used by the state-of-the-art experiment in the field KATRIN. It should be noted that these are all upper limits and therefore at the time it was still not known if the neutrino even had a mass. However, even then the limits being orders of magnitude lower than other particles implied that the coupling of the neutrino had to be much lower than any of the other particles found. This in turn implied that the mass of the neutrino was generated through an alternate mechanism and revived interest in Majorana's work.

Alongside this, the first measurements of the half life of double beta decay were being made on Tin using Geiger counters. At the time it was thought that the half life of $2\nu\beta\beta$ would be longer than that of $0\nu\beta\beta$, so the experiments placed limits on the first. It did not take long for the first claim of a detection to be made in 1949 with a half-life of $4 - 9 \times 10^{15}$ yr [15], but this was ruled out soon after in 1951-53 [16]. Meanwhile, in 1950, double beta decay was experimentally verified in Te using radiochemical methods [17]. These methods measured the number of Xe atoms produced from the decay of Te to Xe to ascertain the half life and therefore are not sensitive to which double beta decay is occurring. It was not until the 1980s, when the first direct detection of $2\nu\beta\beta$ was made in Se-82 using a TPC at T = $1.1 \pm 0.8 \times 10^{20}$ yr [18] with measurements in many other isotopes following soon after.

The years of 1956 and 1957 were a turning point for the field of neutrino physics, in 1956 with the advent of nuclear fission technology it was possible for the neutrino to become experimentally verified which was achieved by Cowan and Reines in Project Poltergeist using the neutrino flux from a nuclear reactor [19]. In the next year (1957), two beautifully simple experiments radically changed the field of particle physics and in particular of weak interactions. Firstly, it was discovered that parity was violated in β decay by Wu et al. [20]. Secondly, Goldhaber et al. using electron capture on Eu-152 showed that only left-handed neutrinos

were produced in weak interactions [21]. These results directly impacted $0\nu\beta\beta$ as it meant that the half life of $0\nu\beta\beta$ could be much longer than that of $2\nu\beta\beta$, meaning more advanced experiments with larger masses would be needed to discover $0\nu\beta\beta$. It also meant that radiochemical methods would not be able to measure the half life of $0\nu\beta\beta$ as they were only sensitive to the total of both decays.

Neutrinos had now been seen from an artificial source but not from a natural one, in the 1960s, Ray Davis set out to change this by measuring the solar neutrino flux using a large tank of cleaning fluid (detecting neutrinos via inverse beta decay on the Cl atoms, producing Ar) in the Homestake mine in South Dakota. He found that the flux was only a third of what was expected [22]. This came to be known as the solar neutrino problem. A solution had in fact already been put forward by Pontecorvo in 1957. At the time two leptons were known the electron and the muon [23], with the tau not to follow until 1975 [24] (completing the 3 generations of particles), each hypothesised to have a corresponding neutrino (the confirmation of the muon neutrino would come in 1962 [25] and the tau in 2001 [26]). Pontecorvo therefore speculated that these neutrinos could oscillate between flavours which meant they must have a mass [27, 28] (as otherwise there would be no evolution of the mass states with time). Incidentally, Pontecorvo had also been the one to propose the idea of using Cl to detect neutrinos in the first place [29]. However, it would take a few decades for this idea to be experimentally verified and for the community to believe the result of Ray Davis (he eventually won the Nobel prize in 2002).

Moving back to the theory side, in the 1970s the Standard Model of particle physics was developed by Salam, Glashow and Weinberg [30, 31, 32]. In the Standard Model all masses are Dirac masses arising from the coupling of left-handed particles to their right-handed counterparts (this is in contrast to Majorana masses which couple particles of the same chirality). Therefore, as Goldhaber had shown that the only neutrinos produced in weak interactions were left-handed, and no right-handed neutrino had been detected, the Standard Model did not include neutrino masses, meaning the neutrino is massless. It took until the discovery of neutrino oscillations in 1998, independently by Super-Kamiokande and SNO [33, 34], definitively showing that neutrinos must have 3 distinct mass states (with at least 2 non-zero) that the first evidence that the Standard Model was incomplete arose. Neutrino oscillation experiments have continued to measure both the mass splittings and the mixing angles to higher and higher precision, with the latest results at the percent level. The confirmation of neutrino masses reinvigorated the question of the neutrino mass generation mechanism and in particular whether it is Majorana or Dirac in nature. In turn this revived interest in neutrinoless double beta decay to answer this question.

This brings us to the current situation. We know there are three neutrinos, at

least two of which must have a mass. The nature of this mass is however still unknown although the limits on both the mass of the neutrino, the half life of $0\nu\beta\beta$ and the sum of neutrino masses are improving all the time. In the next chapters I will lay out the current understanding of the parameter space with the advances made in the theoretical understanding and the experimental techniques now used to constrain the parameter space further.

Chapter 2

Background: Theory

In this chapter I will elucidate the theoretical background of the neutrino mass and the reasoning for the expectation of the existence of neutrinoless double beta decay. I will also discuss the theory of this decay and how it can occur. A full description of the Standard Model and anomaly cancellation can be found in Ref. [35], while more details on neutrino physics can be found in Ref. [36].

2.1 Symmetries of the Standard Model

The Standard Model is currently the most complete description we have of the fundamental particles and their interactions. It consists of 3 generations of fermions (6 quarks and 3 leptons with their associated neutrinos) along with 4 gauge bosons and the Higgs boson, depicted in Figure 2.1. At its core is the idea of gauge symmetries which are local symmetries that the Lagrangian is invariant under. The 3 gauge symmetries are: $SU(3)_C \times SU(2)_L \times U(1)_Y$, corresponding to the colour/strong and electroweak symmetries. These gauge symmetries are responsible for the interactions between particles through the bosons they generate. Alongside these fundamental symmetries there are a number of accidental symmetries of the Standard Model. These are symmetries that are not explicitly in the model's construction but in writing down the Lagrangian are found to be also conserved. The question then arises whether these are true gauge symmetries with an associated symmetry group and boson or only partial with certain interactions where they are violated. The three accidental symmetries are baryon number, lepton number and lepton flavour number. However, due to constraints from anomaly cancellations in the standard model (Figure 2.2) these symmetries do not come by themselves but in linear combinations. The four possible combinations are therefore:

$$B - L, \quad L_e - L_\mu, \quad L_e - L_\tau, \quad L_\mu - L_\tau \tag{2.1.1}$$

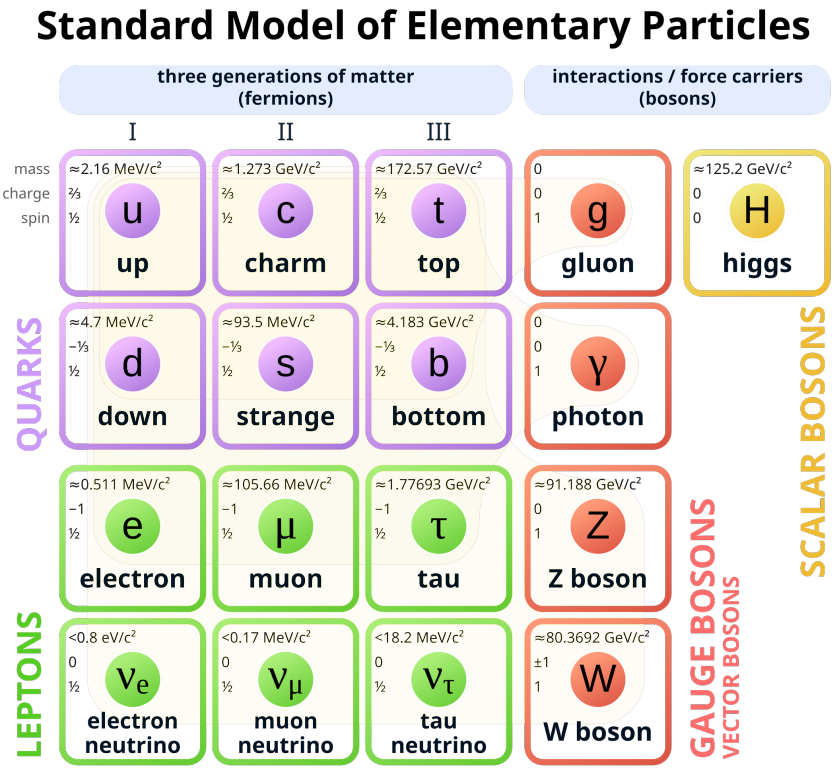

Figure 2.1: The particles of the Standard Model consisting of 6 quarks, 3 leptons, 4 gauge bosons and the Higgs boson arranged into 3 generations of particles.

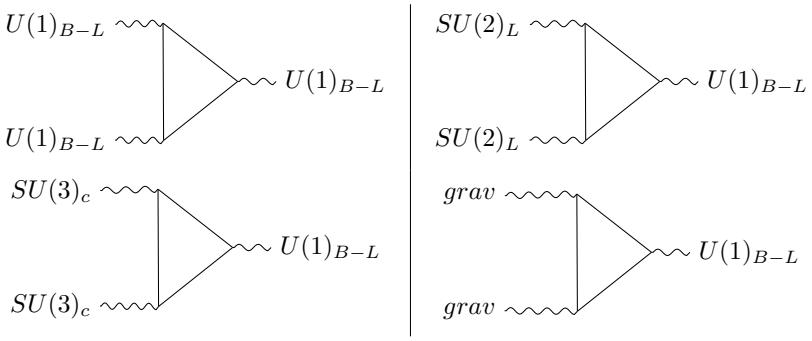

Figure 2.2: The 4 relevant triangle diagrams for showing that B-L is anomaly free. On the left side of each diagram are the known forces of the Standard Model and gravity while on the right is the new U(1) force.

The discovery of neutrino oscillations [34, 33] means that lepton flavour number is not conserved, and therefore the only accidental symmetries left are baryon number and lepton number. The only accidental symmetry group left is then $B-L$. Numerous experiments have, and are being, carried out to find if this symmetry is also violated. In fact one of the discoverers of neutrino oscillations, Super-KamiokaNDE, was an experiment originally designed to search for proton decay which would violate baryon number [37, 38] (the proton is the lightest baryon with baryon number 1, any decay it has would by necessity violate baryon number conservation as quarks can only be arranged in particles with baryon number 1 or 0). Alongside proton decay $B-L$ violation is searched for through exotic muon decays (such as to 2 positrons and an electron [39]) and of course neutrinoless double beta decay.

The reason so many experiments are searching for this violation is to try and explain the baryon asymmetry of the universe. We know there is more matter than antimatter in the local universe and the Standard Model does not provide a mechanism for this because matter and antimatter are always produced in equal amounts if this conservation holds. Many theories have been put forward to solve this problem such as sphalerons [40] and leptogenesis models [41], but all require that $B-L$ is not a fundamental symmetry.

2.2 Majorana and Dirac Mass

Neutrino oscillation show that neutrinos have a mass (the simplest explanation for this is that a massless particle travels at the speed of light and does not have a sense of time so its wavefunction does not evolve), but how can this be added into the Standard Model? The simplest way would be to simply add in a right-handed neutrino for the left-handed to couple to, via the Higgs, giving a Dirac mass term for the neutrino:

$$\begin{aligned}\mathcal{L}_D &= y_\nu \bar{L}\tilde{H}\nu_R \\ &= m_D(\bar{\nu}_L\nu_R + \bar{\nu}_R\nu_L)\end{aligned} \tag{2.2.1}$$

where y_ν is the Yukawa coupling, L is the lepton SU(2) doublet, $\tilde{H}$ is the conjugate Higgs doublet, ν_R/ν_L is the right-handed/left-handed neutrino and m_D is the Dirac mass acquired after spontaneous symmetry breaking of the Higgs field so $m_D = y_\nu v$. This is the simplest way to add a mass for the neutrino, but it has the issue that the Yukawa coupling for this mass has to be much, much smaller than for any of the other fermions in the Standard Model as illustrated in Figure 2.3. An alternative way to add a mass term for the neutrino is to instead add a Majorana mass term.

All the particles in the Standard Model are Dirac spinors this means they have

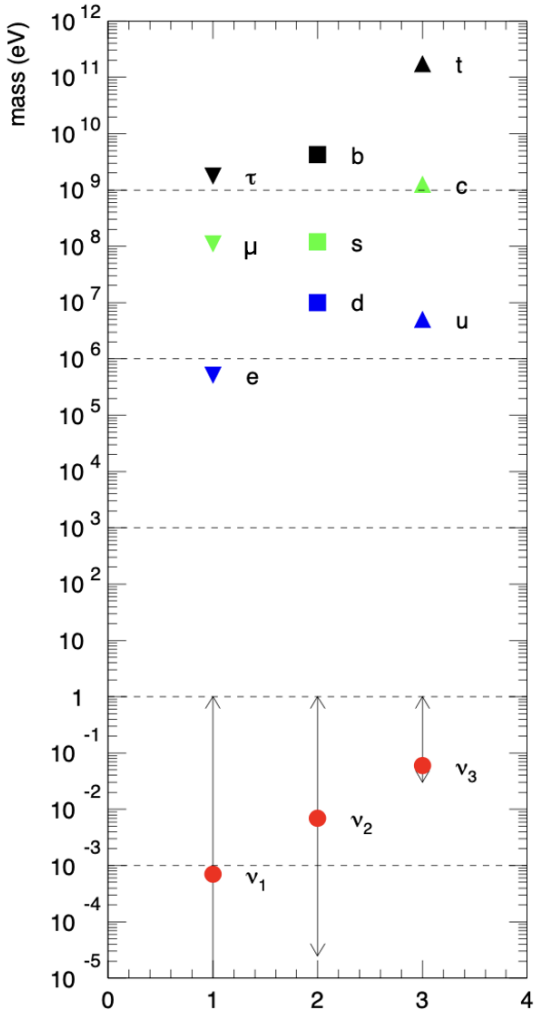

Figure 2.3: The masses of the fermions. It can be seen that the neutrino has a mass 10 orders of magnitude smaller than the other fermions. From Ref. [42]

a left and a right chiral component (Weyl Spinor). Instead, a Majorana spinor is a spinor that only has a left-handed or right-handed component [10]. This means a Majorana particle has to be the same as its antiparticle, and it cannot be charged under any of the Standard Model symmetries as otherwise the charges would be inverted for its antiparticle making them non-identical. This can easily be seen for a U(1) symmetry where when transformed $\Psi \rightarrow \exp(i\theta)\Psi$ and $\Psi_c \rightarrow \exp(-i\theta)\Psi_c$ so Φ cannot equal Φ_c.

As stated above, Majorana spinors only involves spinors of the same chirality and therefore the mass term is:

$$\mathcal{L} \supset m\bar{\Psi}_L\Psi_L. \tag{2.2.2}$$

Or equivalent for a right-handed particle. In the Standard Model all left-handed particles are doublets under $SU(2)_L$ and all right-handed particles are singlets. This means all left-handed particles have a charge under the weak force whereas all right-handed particles do not (the opposite is true for antimatter particles). Therefore, this term cannot be in the Lagrangian for the Standard Model as the neutrino has charges: it is charged under both $SU(2)_L$ and $U(1)_Y$. It could however be written down for a theoretical right-handed neutrino and the implications of this will be discussed next.

2.2.1 Seesaw Model

One way of adding masses to the neutrinos is the so-called seesaw model [43]. In this mechanism a heavy right-handed neutrino 'gives' mass to the SM neutrino. There are many ways that this can be achieved, but first the simplest case will be discussed; the type 1 seesaw. This mechanism relies on rotating between the mass and flavour basis of neutrinos. While in the flavour basis the neutrino might not have an explicit mass term, when rotated into the mass basis it picks up a mass through mixing. For example if considering the 2×2 matrix:

$$\begin{pmatrix} \bar{\nu_L} & \bar{N_R} \end{pmatrix} \begin{pmatrix} 0 & M \\ M & B \end{pmatrix} \begin{pmatrix} \nu_L \\ N_R \end{pmatrix}, \tag{2.2.3}$$

where $\nu_L/\bar{\nu_L}$ are the left-handed neutrino and its antiparticle, $N_R/\bar{\nu_R}$ is the right-handed neutrino. The top left term of the matrix tells us there is no Majorana mass term for the standard model neutrinos, the bottom right (B) would be the Majorana mass term for the new neutrino (e.g. $m/2\bar{N}N$) and the other terms (M) could be related to a Yukawa coupling between the two (e.g. $(L\dot{H})N$) such as the Dirac mass term discussed above. Calculating the eigenvalues of this matrix

gives (2.2.4).

$$\lambda = \frac{B \pm \sqrt{B^2 + 4M^2}}{2}. \tag{2.2.4}$$

The eigenvalues must multiply to the determinant, which is how the seesaw name arises, if one eigenvalue increases the other must decrease. For neutrinos the case where B is much larger than M is considered. In this case the larger eigenvalue is approximately B and the smaller $\frac{-M^2}{B}$. The mass is then the magnitude of these eigenvalues. This means that if the new neutrinos are sufficiently massive or if the coupling is sufficiently small, the SM neutrinos only gain a small mass as observed.

This is the simplest case as only a new right-handed neutrino has been added. There are 2 other similar types of seesaw, imaginatively named type 2 and 3 which add extra scalars and extra fermions respectively but operate on similar principles [44, 45]. The seesaw mechanism is a very attractive way of explaining the smallness of the neutrino masses.

2.3 Neutrinoless Double Beta Decay

Neutrino oscillations mean that there is a mixing between the various neutrino flavours. It is therefore useful to think of the neutrino mass eigenstates (ν_1, ν_2, ν_3) as a linear combination of the flavour eigenstates (ν_e, ν_μ, ν_τ) where the mixing is defined by the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [46]:

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} e^{i\alpha_1/2} & 0 & 0 \\ 0 & e^{i\alpha_2/2} & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{2.3.1}$$

Here the cos and sin terms of each mixing angle have been abbreviated such that $c_{13} = cos(\theta_{13})$, $s_{13} = sin(\theta_{13})$ and so on. This PMNS matrix, (2.3.1), consisting of three mixing angles ($\theta_{12}, \theta_{13}, \theta_{23}$) and a CP violating phase (δ_{cp}), occurs alongside 3 mass splittings ($\Delta m_{12}, \Delta m_{13}, \Delta m_{23}$). In addition, we can add 2 Majorana phases to the matrix to parameterise this new Majorana mass without affecting the oscillation observables [47] as they only depend on the splitting between the neutrino masses and are not sensitive to its mechanism or absolute scale. This gives a total of 9 parameters. Currently, from oscillation experiments, the 3 mixing angles have been measured along with the 3 mass splittings to fairly high precision. The CP violating phase is unknown as is the ordering of the mass states (shown in Figure 2.4, in normal ordering ν_3 the heaviest mass eigenstate while in inverted it is the lightest) and the Majorana phases. The first two are the main goals of the next generation of neutrino oscillation experiments such as DUNE [48], JUNO [49]

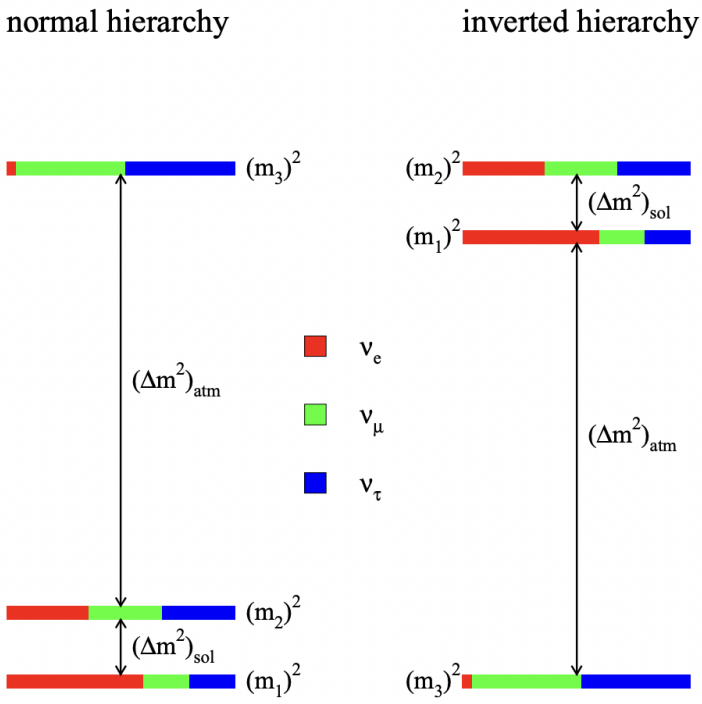

Figure 2.4: The two possible hierarchies of the neutrino masses: normal and inverted. In normal ordering m_3 is the heaviest state whereas in inverted ordering it is the lightest. Shown here are the masses squared as the mass splittings squared are measured from oscillations. From [51]

and Hyper-Kamiokande [50] while the last can not be measured by oscillation experiments.

The Majorana phases are not observable in neutrino oscillation experiments, but they can be probed in neutrinoless double beta decay experiments. In double beta decay two neutrons in the nucleus decay into two protons releasing two electrons and two neutrinos. However, if a neutrino is a Majorana particle then an alternative form is possible with no neutrinos emitted. I have laid out one method of generating a Majorana mass for the neutrino but in fact many exist. Thankfully, Schechter Valle or the black box theorem [52] states that if $0\nu\beta\beta$ exists then whatever the mechanism is that generates it the neutrino must have a Majorana component to its mass.

The simplest mechanism that can generate $0\nu\beta\beta$ is light Majorana neutrino exchange. This would contribute a term to the Lagrangian:

$$\mathcal{L} = \frac{1}{2}\sum \nu_l^t C^{-1} M_{ll'} \nu_{l'} \tag{2.3.2}$$

where C is the charge conjugation matrix and $M_{ll'}$ is the Lepton mass matrix. The term of interest involves the two electron neutrino part which is proportional to M_{ee} and often written as $m_{\beta\beta}$. This matrix can be written in terms of the known mass eigenstates and the mixing matrix allowing us to define an effective Majo-

rana neutrino mass, $m_{\beta\beta}$, as [53]:

$$m_{\beta\beta} = |\sum_{i=1,2,3} U_{ei}^2 m_i| \tag{2.3.3}$$

where the U_{ei} are the PMNS matrix elements and m_i are the neutrino masses. Experimentally though the parameter measured is the half life of the decay which is related by:

$$(t_{1/2})^{-1} = \frac{1}{m_e^2} G_{0\nu} |\mathcal{M}|^2 (m_{\beta\beta})^2 \tag{2.3.4}$$

Where $G_{0\nu}$ is the phase space factor which describes the kinematics of the decay, m_e the mass of the electron and $\mathcal{M}$ is the nuclear matrix element which will be described next.

2.3.1 Nuclear Matrix Elements

The nuclear matrix element encodes the nuclear structure physics of the decay. It is highly dependent on the isotope used and is the main source of the uncertainty on $m_{\beta\beta}$. There are a number of different ways of calculating it which give slightly different values making the conversion from half life to effective mass challenging. These nuclear matrix elements are still being updated with uncertainty in the quenching of the axial coupling and in additional nucleon current contributions [54]. In recent years there has been the development of ab-initio methods for calculating these matrix elements which show promise in answering some of the remaining questions and remove the need for quenching.

These matrix elements make comparing experiments more difficult as the conversion from half life to effective mass carries a large uncertainty due to the matrix elements used. Furthermore, not all isotopes have been calculated for all the different methods with larger isotopes being more challenging than smaller ones due to the larger number of nucleons with their corresponding interactions. For example the matrix element for Mo-100 has not been calculated for the nuclear shell model (NSM). Figure 2.5 is a summary of the situation for one part of the matrix element for the major isotopes showing the variation in the result between the different calculation methods and uncertainties for each method.

2.3.2 Possible Isotopes

In most isotopes that can undergo $0\nu\beta\beta$ decay the signal is swamped by the single beta decay background. For this reason it is beneficial to use isotopes where the single beta decay is forbidden. This means you want an even-even nucleus so that the beta decay would change it to odd-odd. In some cases this odd-odd

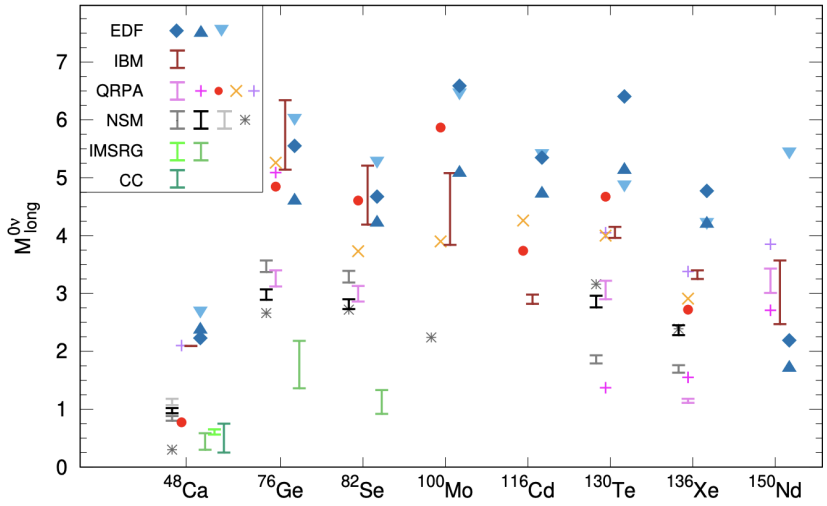

Figure 2.5: Matrix element calculations for some of the major isotopes being used in $0\nu\beta\beta$ searches. From: [2].

nucleus is energetically unfavourable, so the beta decay is forbidden and only the double beta decay to the next even-even nucleus is possible. An example of this is shown in 2.6 for ^{76}Ge. Alternatively, an isotope where the single beta decay is suppressed due to angular momentum conservation considerations can also be used.

2.4 Neutrino Mass Parameters

As well as the effective Majorana mass defined above there are two other parameters generally considered in the neutrino mass space which is the absolute mass scale of the electron neutrino and the sum of the masses.

The parameter space for $0\nu\beta\beta$ is summarised in Figure 2.7 for both the inverted and normal orderings. One feature of this plot is the presence of cases where $m_{\beta\beta}$ goes to extremely small values making $0\nu\beta\beta$ detection extremely challenging. It should be noted, however, that these values, while possible, require more fine-tuning of the parameters and are therefore disfavoured if one assumes naturalness as can be seen in Figure 2.8. Indeed, this shows that by covering the inverted ordering the parameter space for the normal ordering is also covered for around 50% of cases. Neutrino oscillation experiments such as JUNO will hopefully be able to resolve this mass ordering in the next decade which will narrow down the parameter space further [49]. Additionally, many experiments are seeking to measure the other 2 parameters which will limit the $0\nu\beta\beta$ parameter space even more and these will be described now.

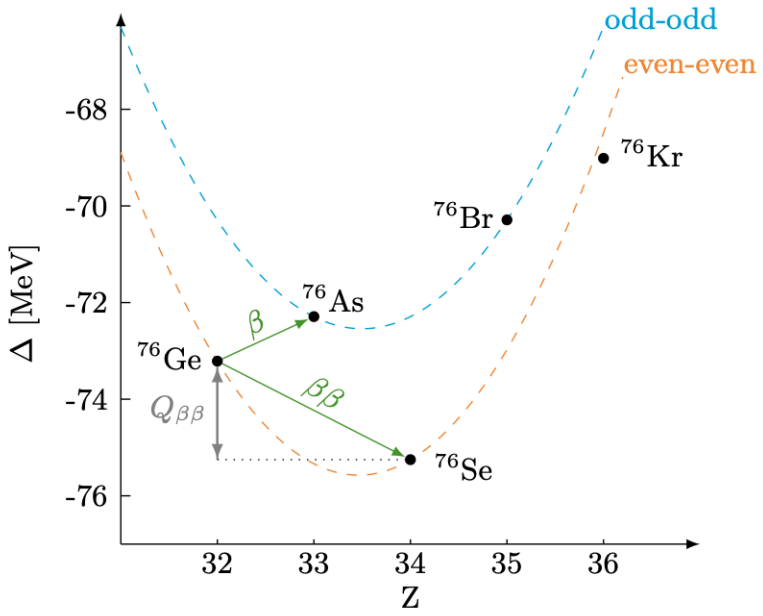

Figure 2.6: Energy diagram for the possible decays of ^{76}Ge [2]. The decay to As-76 is energetically forbidden making the double beta decay to Se-76 observable.

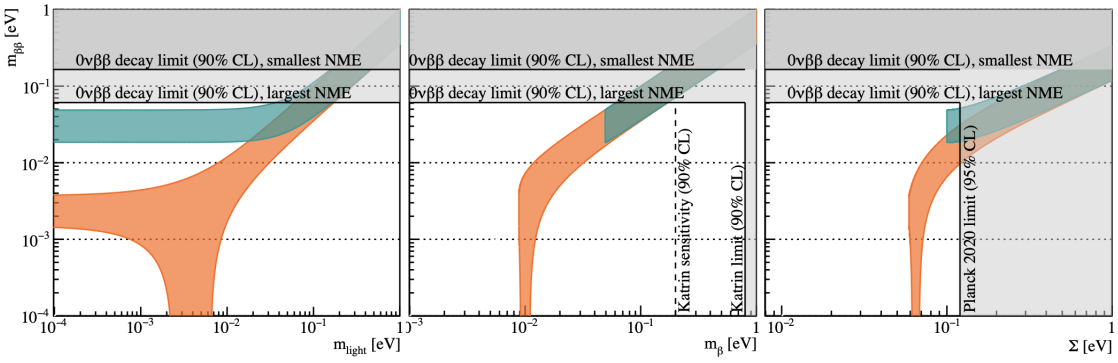

Figure 2.7: Parameter space for the neutrino mass parameters. For each plot the inverted ordering is shown in blue while the normal ordering is shown in red/orange. On the left is $m_{\beta\beta}$ in terms of the lightest neutrino mass m_{light}, which is the so-called Lobster plot. In the middle is m_β which is measured in tritium beta decay experiments, the best limit from KATRIN is shown. On the right is the sum of neutrino masses measured in cosmological experiments, the limit from Planck is shown which is close to ruling out the inverted ordering space. In all cases the limits from current $0\nu\beta\beta$ experiments are also shown. Lower $m_{\beta\beta}$ values are allowed in the normal ordering than the inverted ordering due to there being 2 low mass states and 1 higher mass state in the inverted case while the opposite is true in the normal case as seen in Figure 2.4. From ref. [2].

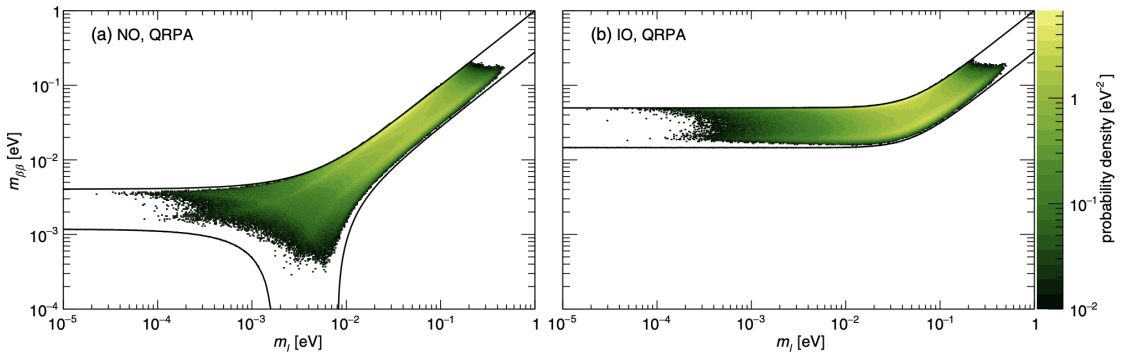

Figure 2.8: Discovery sensitivity parameter space for $m_{\beta\beta}$ for both the inverted and normal orderings for one matrix element (QRPA: the quasiparticle random phase approximation). From: Ref. [55].

2.4.1 Absolute Neutrino Mass

A number of experiments are currently trying to measure the absolute neutrino mass with the current leader being KATRIN which is a kinematic measurement of the endpoint of the Tritium beta decay spectrum. Its newest limit is $m_\nu < 0.45$ eV at 90 % CL [56]. Interestingly the best fit value for the mass squared is currently negative. The eventual goal is to achieve a sensitivity of 0.3 eV. Beyond that a new experimental approach could be needed as there is a limit to how large a mass spectrometer can be built. Ideas include measuring the neutrino mass through cyclotron radiation emission spectroscopy (CRES) being pioneered by the PROJECT8 collaboration [57] and using the endpoint of the spectrum from the electron capture on Ho-163, such as HOLMES [58].

2.4.2 Sum of Neutrino Masses

With the advent of precision cosmology, it became possible to measure the sum of the neutrino masses. It can be calculated through the equation:

$$\Omega_\nu h^2 = \frac{n_\nu \Sigma m_\nu}{\rho_c} \tag{2.4.1}$$

where Ω_ν which is the fraction of the critical density of the universe in neutrinos, h is the Hubble constant, n_ν is the number density of neutrinos, Σm_ν is the sum of the neutrino masses and ρ_c is the critical energy density which depends on the number of neutrinos. Experiments therefore seek to measure Ω_ν then convert this to a limit on the sum of the neutrino masses.

Ω_ν affects cosmic structure formation and the effect of it can be seen in the Cosmic Microwave Background (CMB) and in large scale structures such as Baryon Acoustic Oscillations (BAO). These two measurements are combined to give the best limits. On the CMB side this generally means using data from the Planck satellite and on the BAO side using data from the Dark Energy Spectroscopic

Instrument (DESI). Other measurements such as from gravitational lensing, Supernova data and Lyman Alpha Forest data can also be added to this to improve the constraints.

The current best limits are $\Sigma m_\nu < 0.072(0.113)$ eV at 95% CL [59] in the ΛCDM paradigm, excluding the inverted ordering and placing it very close to the minimum sum in the normal ordering. Interestingly some analyses extending the analysis to $\Sigma m_\nu < 0$ even place the best fit value at 160 ± 90 meV at 68% CL [60] implying a negative neutrino mass.

However, it is important to bear in mind a few things: firstly these measurements assume standard cosmological models, are currently very sensitive to the datasets used [61] and the parameter being measured is not exactly the same as the sum of neutrino masses but an effective parameter. With the launch of the Euclid satellite in 2022 the measurements of large scale structure will improve significantly, so this will be a very active area of interest in the coming years.

The field of neutrino mass is, as can be seen, an extremely active one on both the theoretical side, with the new ab initio methods for calculating matrix elements, and on the experimental side for all the different parameters of interest. The constraints placed on the neutrino mass parameter space will provide insights on the mechanism behind fermion mass and also on cosmological questions regarding the evolution of the universe.

Chapter 3

Experimental Overview

In this section I will present a brief overview of how searches for $0\nu\beta\beta$ are performed along with a brief description of the leading experiments excluding the LEGEND project which will be discussed in the next chapter. I will conclude with a look to the future of the field and the next generation of experiments seeking to explore fully the inverted ordering.

3.1 Searching for $0\nu\beta\beta$

As I have laid out in previous chapters, there is a great deal of interest in the search for $0\nu\beta\beta$ and, therefore, there are many experiments searching for it. In standard double beta decay 2 neutrons decay to 2 protons releasing 2 electrons and 2 neutrinos while in $0\nu\beta\beta$ the neutrinos are not emitted. In both types of double beta decay, it is also possible to decay into excited states which may then decay to the ground state emitting additional gamma rays. Unfortunately, these gamma rays would be common to both $0\nu\beta\beta$ and $2\nu\beta\beta$ searches and therefore again cannot be used to separate the two (although it could be used to separate

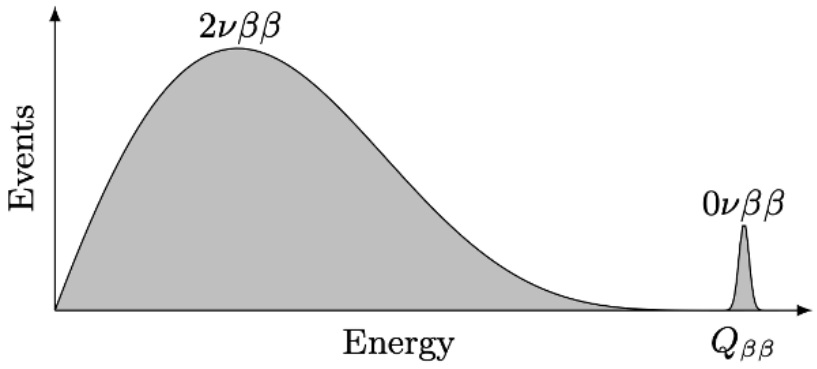

Figure 3.1: An idealised spectrum with the $2\nu\beta\beta$ continuum and the $0\nu\beta\beta$ peak at the end. From Ref. [2].

background). As described in Sec. 1 the first experiments to search for $0\nu\beta\beta$ were radiochemical methods identifying the daughter nucleus. However, these methods cannot separate the $2\nu\beta\beta$ decay from the $0\nu\beta\beta$ decay. Instead, information from the leptons must be used.

As the mass of the daughter nucleus is much larger than that of the other particles emitted, the recoil energy can be neglected, while the neutrinos emitted in $2\nu\beta\beta$ will not be detected. This means that in $2\nu\beta\beta$ the Q-value of the decay is split between the 4 leptons, and therefore the energy spectrum of the summed electrons is continuous up to the Q-value. While in $0\nu\beta\beta$ the summed energy of the electrons is monochromatic as shown in Figure 3.1 at the Q-value, $Q_{\beta\beta}$. The summed energy of the electrons, therefore, is the most used discriminator between the two decays. Here, one is searching for an accumulation of events at the Q-value of the decay (there will still be an irreducible background here of $2\nu\beta\beta$ events which will depend on the energy resolution of the detector).

Broadly there are 2 categories of $0\nu\beta\beta$ experiment designs: the ”source as detector” designs and the ”source and detector” designs. In the first the detector itself is formed from the isotope of interest. This can take the form of a liquid TPC filled with the isotope of interest, a liquid scintillator detector where the scintillator is doped with the isotope or a solid state detector (such as a Ge detector or cryogenic calorimeter) where the isotope is the detector itself and the energy is detected as a charge or phonon. Instead, in the ”source and detector” designs the isotope and detector are separate which can take the form of a source foil surrounded by a tracker and calorimeters or a balloon of isotope in a liquid scintillator detector.

The great benefit of the ”source as detector” designs is that the electrons’ energy is more easily contained within the detector and energy losses are minimised. Although this does mean the electrons’ energy is deposited within a few mm of the decay site and therefore no topological information such as the energy of the individual electrons or the angle between can be gained from the decay. Otherwise, one can use foils, but these must be very thin to ensure the electrons can escape the source with few scatters and therefore minimal energy loss or a low density medium such as gaseous time-projection chambers which are becoming an increasingly promising technology for topological study of the electrons’ kinematics, as their absorption requires centimetres of gas.

$0\nu\beta\beta$ is of course an extremely rare decay with half life limits now greater than 10^{26} yr. This means one needs to observe a very large mass of the isotope of interest (of order 100 moles and greater) for a significant period of time (multiple years with some experiments running for a decade). Additionally, a low enough background is needed to be able to see the signal over it. To take an example, for a half life of 10^{28} yr and an eventual exposure of 10 ton yr, LEGEND-1000 would expect to see around 3 events within 5 keV of $Q_{\beta\beta}$ over the lifetime of the

experiment, so the background in this region needs to be at the single event level to make a discovery.

Experimentally the half life sensitivity scales as:

$$T_{1/2}^{0\nu\beta\beta} \approx \begin{cases} \sqrt{\frac{Mt}{b\ \Delta E}} & \text{with background} \\ Mt & \text{background-free} \end{cases} \tag{3.1.1}$$

where M is the active mass, t the exposure, b the background index at $Q_{\beta\beta}$ and ΔE the energy resolution. If the experiment was completely background free then the sensitivity would grow linearly just as the active mass times the exposure, whereas otherwise it goes as the square root. Of course every experiment has some background and therefore cannot scale exactly linearly. Instead, at a low enough background one enters the quasi-background free region. This is shown in Figure 3.2 for various backgrounds. Here the scaling approaches linear as the background decreases. A low background is therefore crucial to achieving the best discovery sensitivity for a given mass.

Experiments are generally talked about in terms of their discovery sensitivity and exclusion sensitivity. The discovery sensitivity corresponds to the expected number of signal events for which an experiment has 50% chance to observe an excess of events over the background at 99.73% CL. The exclusion sensitivity corresponds to the expected number of signal events that an experiment has 50% chance of excluding at 90% CL [2]. The best exclusion sensitivity on $0\nu\beta\beta$ would be achieved by maximising sensitive exposure while staying under the signal threshold above. Instead, it is more important for the best discovery sensitivity to minimise the background. LEGEND's focus is on the discovery sensitivity and therefore the minimisation of background.

Figure 3.3 shows the parameters of interest for $0\nu\beta\beta$ experiments. Here experiments are characterised by two main quantities related to these sensitivities: sensitive exposure and sensitive background. The sensitive exposure is the product of the number of moles of isotope, the live time and the signal efficiency, while the sensitive background is the number of events in the signal region divided by the sensitive exposure [2]. The mass and total efficiency are self-evident, these define how much isotope is needed and how likely its decays are to be detected. The energy resolution is also important as it allows the $2\nu\beta\beta$ background to be separated from the $0\nu\beta\beta$ signal. Experiments with a low energy resolution will have to deal with this irreducible background. Approximately, the ratio of the number of $2\nu\beta\beta$ events to $0\nu\beta\beta$ events is given by:

$$\frac{S}{B} \approx \left(\frac{Q_{\beta\beta}}{\Delta E}\right)^6 \frac{T_{1/2}^{2\nu}}{T_{1/2}^{0\nu}} \tag{3.1.2}$$

where ΔE the energy resolution and therefore $\Delta E/Q_{\beta\beta}$ is the fractional energy

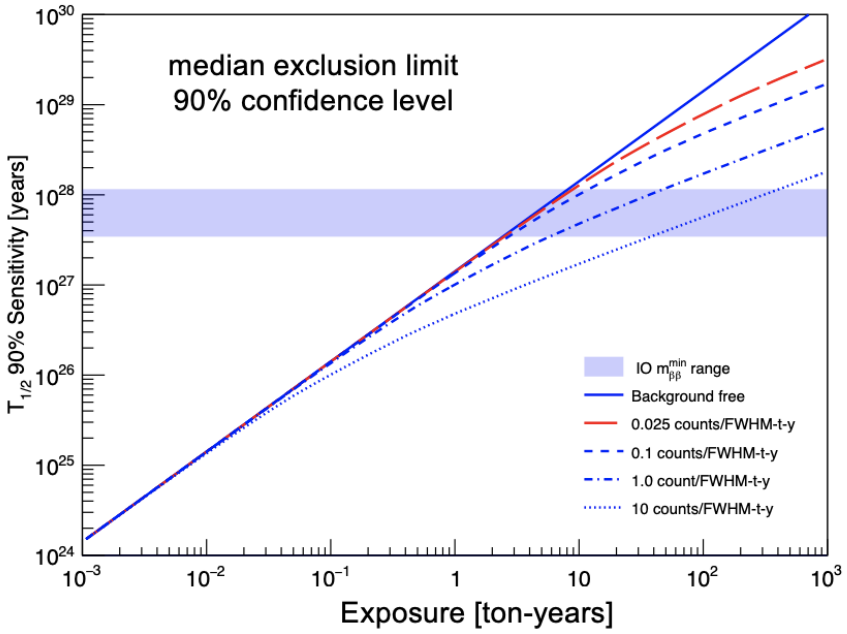

Figure 3.2: Exclusion sensitivity curves for various background indexes. LEGEND-200 will have a total exposure of 1 ton yr and LEGEND-1000 a total exposure of 10 ton yr. From Ref. [62].

resolution. Full details can be found in [63]. However, this also shows that it does depend on the isotope as to how good an energy resolution is needed.

In this way, the isotope used defines the optimal strategy for the experiment. In particular whether to increase the exposure or decrease the background. For example Te experiments can increase the mass more easily due to the high natural abundance of Te (as shown in Figure 3.4). By contrast due to the cost of Ge/Xe isotopic enrichment it makes more sense to reduce the background than to increase the mass. This difference can be seen in Figure 3.3 where Ge experiments have a much lower mass but better energy resolution and sensitive background, while Te experiments have a much higher mass and sensitive background. Both, however, end up with similar sensitive exposures.

The isotope chosen also has a real bearing on the backgrounds the experiment has to deal with. Isotopes with a high $Q_{\beta\beta}$ value (Figure 3.4) will be above much of the background from radioactive isotopes present in the environment. In particular, the uranium and thorium backgrounds which produce a range of gamma lines with the highest being Bi-214 at 2.2 MeV and Tl-208 at 2.6 MeV respectively. Above the Tl-208 gamma line, radioactivity is mainly in the form of alpha decays which necessitates either the use of shielding or a pulse shape discrimination (PSD) method to remove. Building the experiment from clean materials is also essential in mitigating this. Finally, all experiments deal with backgrounds from cosmic rays such as muons passing through the experiment or cosmogenic activation, this is why all experiments are located deep underground and many have some cooldown time for their materials.

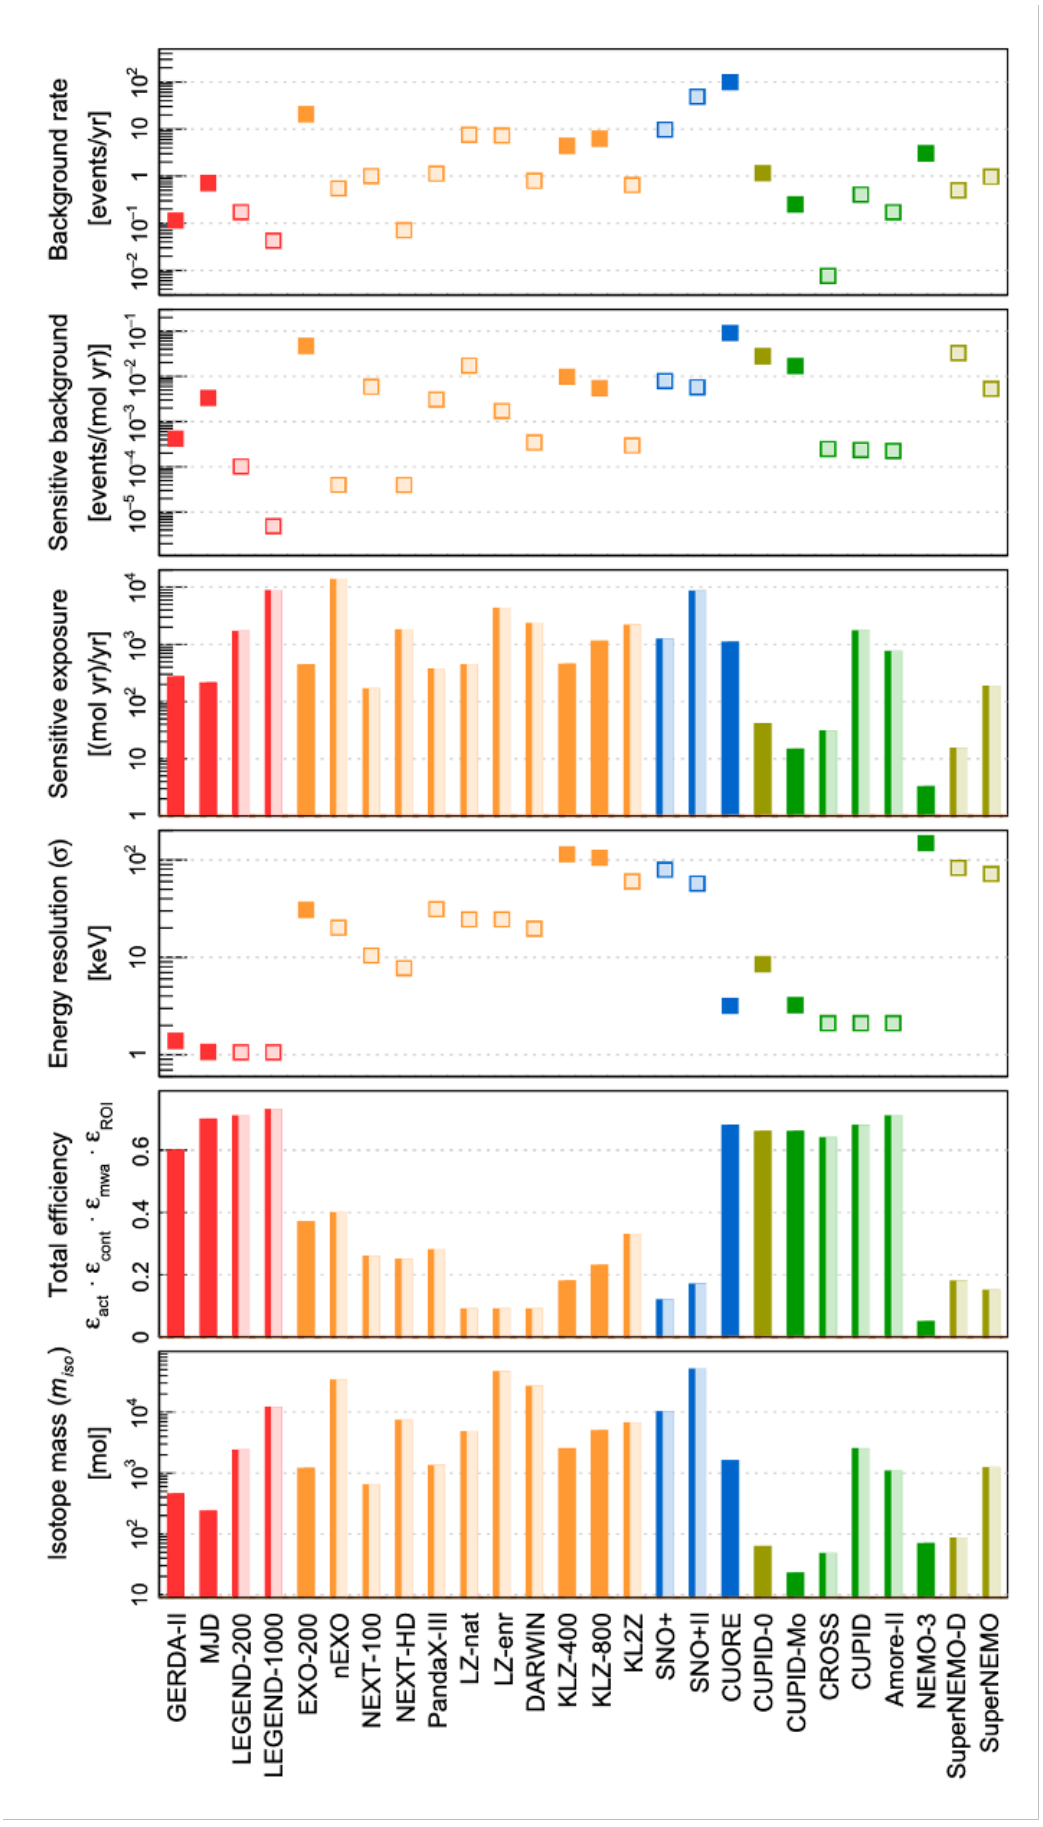

Figure 3.3: Experimental parameters of current and future $0\nu\beta\beta$ experiments (current filled bars and points, future split/unfilled bars and points). The colours refer to the different isotopes used: Red for Ge-76, Orange for Xe-136, Blue for Te-130, Green for Mo-100 and olive for Se-82. From: [2].

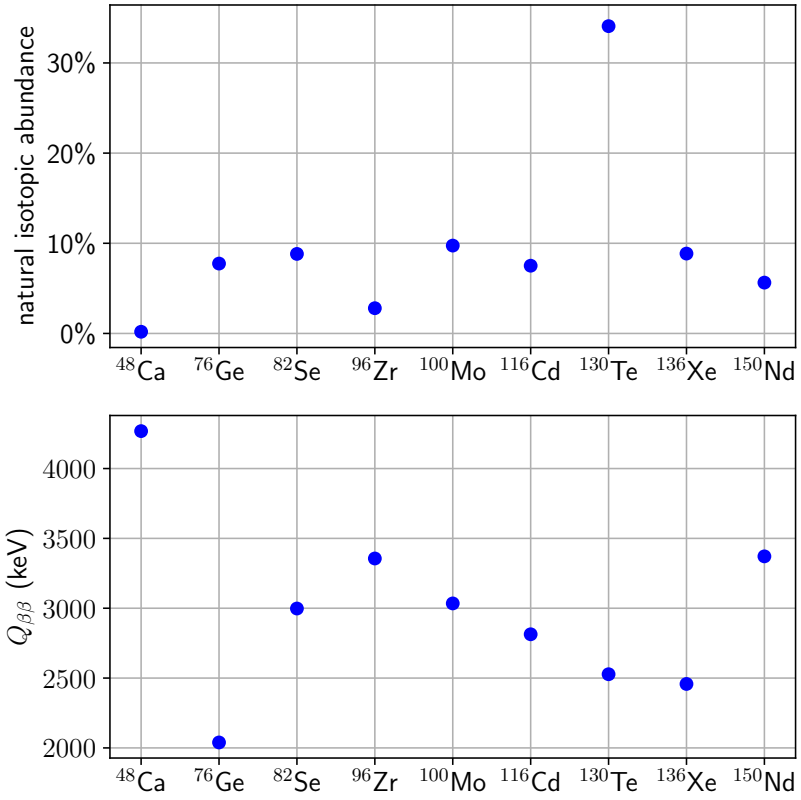

Figure 3.4: The natural isotopic abundance and the $Q_{\beta\beta}$ value for the leading isotopes used in $0\nu\beta\beta$ searches. Values from: [2].

Finally, the isotope defines the conversion from the half life to the ultimate parameter of interest the effective Majorana mass, $m_{\beta\beta}$, through the phase space factor and nuclear matrix elements. For example, isotopes with a higher nuclear matrix element value will have a lower $m_{\beta\beta}$ for the same half life.

Below is a table of the leading isotopes with their $Q_{\beta\beta}$ value and the current lower limits on the decay half life, the experiment that gave this limit along with the main experimental techniques used for each isotope. This highlights the main techniques in the field, High Purity Germanium (HPGe) will be discussed in the next chapter, LS is liquid scintillator which can be doped with isotope and the light detected with photomultiplier tubes, TPC is a time projection chamber these can be single or dual phase in both cases the ionisation is drifted to a readout plane with dual phase also having light readout, finally bolometers are cryogenic calorimeters which will be discussed more below.

3.2 Current Experiments

I will now describe the current leading experiments in the field of $0\nu\beta\beta$ searches.

Isotope	$Q_{\beta\beta}$ (keV)	$T_{1/2}^{0\nu\beta\beta}$ Lower Limit 90 % CL (yr)	$m_{\beta\beta}$ Upper Limit 90 % CL (meV)	Experiment	Techniques
Ge-76	2039.0	$> 1.9 \times 10^{26}$	79–180	GERDA [64]	HPGe
Se-82	2995.5	$> 4.4 \times 10^{24}$	263-545	CUPID-0 [65]	Tracking Calorimeter, Bolometer
Mo-100	3034.4	$> 3.0 \times 10^{24}$	210-610	AMoRE [66]	Bolometer
Te-130	2527.5	$> 3.8 \times 10^{25}$	70-240	CUORE [67]	LS, Bolometer
Xe-136	2457.8	$> 3.8 \times 10^{26}$	28-122	KamLAND-ZEN [68]	LS, TPC

Table 3.1: Table of the experimental parameters of the leading isotopes, featuring the Q-value for the decay in keV, the current best constraint on the half life with the experiment that achieved that constraint and the techniques generally used for the isotope.

KamLAND-ZEN

The current leading limit is set by the KamLAND-ZEN experiment. Based in the Kamioka mine in Japan, it reuses the KamLAND infrastructure (a neutrino oscillation experiment looking at antineutrinos from nearby nuclear reactors [69]), adding a nylon balloon containing Xe dissolved in liquid scintillator. This rests in an even larger volume of liquid scintillator. The outer sphere is instrumented with photomultiplier tubes to detect scintillation light. Currently, 745 kg of Xe is being used which is isotopically enriched to 90% in Xe-136. The main backgrounds come from: the $2\nu\beta\beta$ due to the poor energy resolution (6.7%), radioactive impurities on the balloon and spallation products from cosmic muons. The final result used 2 ton yr of exposure and is the leading constraint in the field with a half life lower limit of 3.8×10^{26} yr at 90% CL corresponding to an upper limit of 28–122 meV on $m_{\beta\beta}$ [68].

CUORE

Based at the Laboratori Nationale di Gran Sasso in Italy, CUORE uses bolometers (cryogenic calorimeters) to search for $0\nu\beta\beta$. These bolometers are made from Te O_2 (Te-130 being the isotope of interest) which are coupled to thermal sensors (neutron transmutation doped Ge thermistor) to detect heat via phonons. 742 kg of Te which is 34% Te-130 corresponding to 200 kg of the isotope of interest. This is split into 988 bolometers which are arranged in 19 towers and placed in a large dilution refrigerator to keep them at very low temperatures (10 mK). Finally, this is surrounded by multiple layers of lead to reduce backgrounds with the innermost of these consisting of lead from a Roman shipwreck which is depleted in Pb-210 (this isotope is at the bottom of the uranium chain which features an alpha from the decay of Po-210 which would be a background for the experiment). Due to the large number of phonons produced in particle interactions, bolometers have a very good energy resolution in CUORE this is (7.32 ± 0.02) keV (FWHM), cor-

responding to 0.3% relative resolution at $Q_{\beta\beta}$. They also have the ability to use PSD and the high granularity of the detector to further suppress backgrounds.

CUORE has run for 5 years, achieving 2039 kg yr of total exposure (567 kg yr exposure of the isotope of interest). It has a background index of background index of (1.42 ± 0.02) counts / (keV kg yr) in the region of interest which is dominated by degraded alphas from close by components [70]. This allowed them to set a lower limit of 3.8×10^{25} yr (90% CL) which corresponds to a upper limit of $m_{\beta\beta}$ < 70 – 240 meV [67].

Other Experiments

There are a number of other experiments currently running of which I will only mention a couple. SNO+ is another liquid scintillator experiment. Based in Canada it uses Te which is dissolved in the liquid scintillator. It is currently loading to 0.5% Te by mass, beginning next year, corresponding to 1.3 tonnes isotope, with the eventual aim to reach 3%. At 0.5% this would give a projected sensitivity of 9.20×10^{25} yr (90% CL) after 1 year of data taking [71]. Finally, SuperNEMO uses a very different technique. Based at LSM in France it uses a source foil sandwiched between 2 tracking detectors and calorimeters. This, uniquely, would give the ability to study the kinematics of the decay if found. It is currently using Se-82 with 6 kg of isotope. The projected sensitivity is 4.6×10^{24} yr (90% CL) [72].

3.3 Future Experiments

The current range of experiments aim to touch the top and start delving into the parameter space of the inverted ordering. The next generation of experiments will push further to reach the bottom of the inverted ordering and beyond. The major contenders are discussed here. Figure 3.5 shows the projected sensitivities for some of these experiments along with a band for lowest $m_{\beta\beta}$ value for the inverted ordering. The width of this band depends on the values of the neutrino oscillation parameters while the distribution of the sensitivities for each experiment are due to the nuclear matrix elements uncertainties.

CUPID

The next version of CUORE, CUPID will use the same bolometer technology but with a switch from Te to Mo as the isotope of interest. It will be increasing the mass used to 253 kg of Mo-100. It will improve its background suppression by adding particle identification capabilities through the addition of light readout which will allow it to separate alphas (the main background in CUORE). The shift from Te to

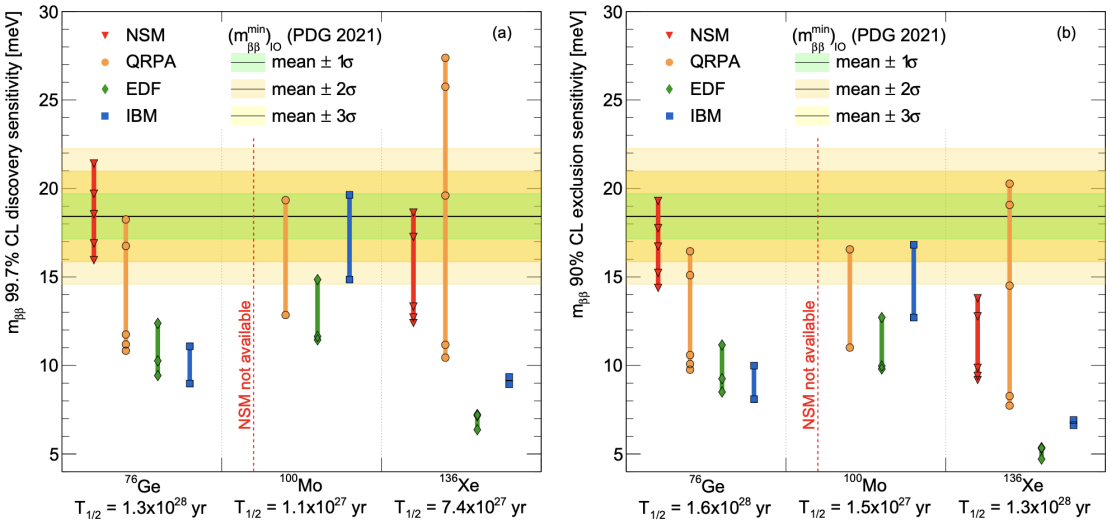

Figure 3.5: Discovery and exclusion sensitivities for the next generation of $0\nu\beta\beta$ experiments (LEGEND-1000, CUPID and nEXO). The bands show the bottom of the inverted ordering for variations of the neutrino oscillation parameters while the points are the values of the sensitivity for each of the various matrix elements calculations. From: [73].

Mo will also mean the Q-value is higher (3034.4 keV as compared to 2527.5 keV) placing it above the gamma line from Tl-208 the highest energy gamma line in the Thorium and Uranium decay chains, further reducing backgrounds. The background target is 10^{-4} counts / (keV kg yr) and the sensitivity is improved by a factor of 5 to 1.5×10^{27} yr (90% CL) [74].

KamLAND2-ZEN

KamLAND-ZEN will also have a newer version, KamLAND2-ZEN. This will use a higher mass of 1 tonne of Xe-136 and again improvements to the energy resolution by using new light concentrators, new PMTs and improved scintillator purity. This will lead to improvements in background as the $2\nu\beta\beta$ will be better resolved from $0\nu\beta\beta$. Alongside this there will be further improvement from improved balloon purity. The background target is a reduction by a factor of 20 over KamLAND-ZEN and the sensitivity target is 1.1×10^{27} yr [75].

NEXO

NEXO is the latest version of the liquid Xe TPC technology developed in EXO-200. It uses a 2 phase read out of both ionisation and scintillation light giving it particle identification capabilities for background discrimination. The total mass is 5 tonnes of Xe isotopically enriched to 90% in Xe-136 (4.5 tonnes). The background is expected to be improved by a factor of 1000 over EXO-200, giving a discovery

sensitivity of $T_{1/2}$ = 7.4×10^{27} yr after 10 years [76].

Other Experiments

There are a number of other experiments in the pipeline which are simply listed here, parameters can be found in Figure 3.3:

- PandaX is a high pressure gas time projection chamber using Xe to search for dark matter as well as $0\nu\beta\beta$ [77]
- NEXT is a high pressure gas time projection chamber using Xe, developing daughter isotope tagging to reduce backgrounds, the current version is NEXT-100 with 87 kg of enriched Xe-136 under construction [78]
- The dark matter time projection chambers such as LZ [79], DARWIN [80] and XLZD [81]. Due to their very large masses these will be competitive in $0\nu\beta\beta$ searches, especially with the possibility of runs with enriched isotopes.

3.4 Conclusion

As can be seen, the field of $0\nu\beta\beta$ searches is very active and a high priority in particle physics. In the near term (next 5 years) the main experiments will be LEGEND-200, KamLAND-ZEN and SNO+ which will start to probe the inverted ordering region. Then, hopefully towards the end of this decade, the next generation of experiments will start to come online, featuring masses of over a tonne these will run into the 2030s probing in some models the entire inverted ordering region and a considerable part of the normal ordering.

Chapter 4

The LEGEND Experiment

In this chapter I will introduce the LEGEND experimental program, giving a description of the current phase: LEGEND-200 and a look to the future with LEGEND-1000.

4.1 History of Ge-76 based experiments

There has been a long history of using Ge detectors for $0\nu\beta\beta$ searches as can be seen in Figure 4.1. The first constraint on the half life of $0\nu\beta\beta$ decay in Ge was set in 1968 by Fiorini et al. [83] using a single natural Ge(Li) detector (this is Lithium-Drifted Germanium which is Ge counter doped with Li). Since then, there have been many experiments using Ge detectors to search for $0\nu\beta\beta$ decay. The development of High Purity Germanium crystals (HPGe) and of improved crystal geometries (semi-coaxial to point contact to inverted point contact) has helped improve the sensitivity of these experiments. It would be remiss to talk about Ge experiments without mentioning the Heidelberg-Moscow experiment which is the most recent experiment to claim a discovery of $0\nu\beta\beta$ in 2002 at 1.5×10^{25} yr (this claim was made by a subgroup of the experiment [84] and not accepted by the community). Which brings us to the most recent experiments in GERDA and MAJORANA.

GERDA, based at the Laboratori Nationale di Gran Sasso in Italy, achieved both the lowest background of any $0\nu\beta\beta$ experiment and set the best limit on $0\nu\beta\beta$ decay in Ge at $T_{1/2} > 1.8 \times 10^{26}$ yr at 90% CL [64] (ruling out the Heidelberg-Moscow claim). It innovated the use of bare Ge crystals immersed in liquid argon which acts both as a coolant, a shield against external radiation and, with instrumentation, as an active veto against background events due to radioactive contaminations on the materials surrounding the detectors. 15.6 kg of detectors were used in the first phase of the experiment and 45.2 kg in the second phase giving a total exposure of 160 kg yr.

Meanwhile, MAJORANA has the best energy resolution at 2.5 keV (FWHM) at

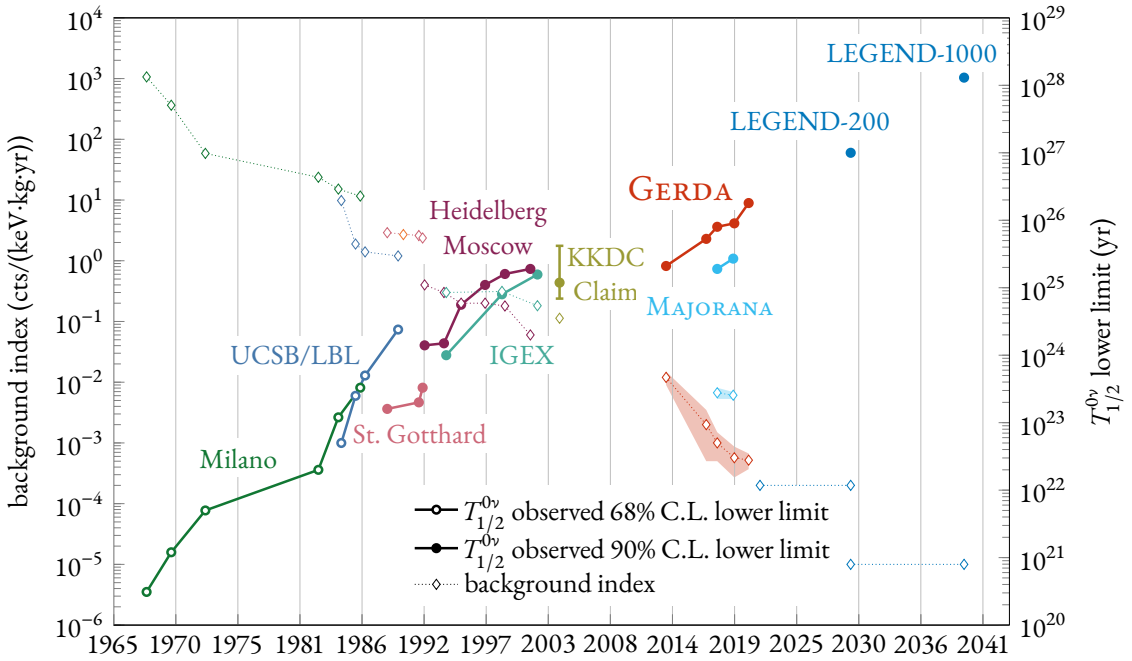

Figure 4.1: History of Ge-76 experiments in terms of the lower limit set and the background index achieved. The solid lines and circles show the progression of the lower limit on the half life while the dashed lines and diamonds show the background index. Adapted from Ref. [82].

$Q_{\beta\beta}$ and pioneered the use of various ultra-clean components such as low mass front end electronics and underground electroformed copper for detector holder-s/shielding. It operated up to 40.4 kg of detectors achieving a total exposure of 64.5 kg yr. Its final limit was 8.3×10^{25} yr (90% CL) [85]. These 2 experiments joined forces to form the LEGEND collaboration building the next generation of Ge $0\nu\beta\beta$ experiments.

4.2 LEGEND-200

LEGEND-200 is the first phase of the LEGEND program. Reusing the GERDA infrastructure in hall A of LNGS (Figure 4.2), it will eventually use 200 kg of HPGe detectors. It is designed as a series of layers to reduce background with the Ge detectors at the centre as shown in Figure 4.3.

The first step in the background reduction is the location, with the mountain overburden of 1800 m water equivalent reducing the muon flux and therefore background from cosmogenically activated isotopes in the detector materials. The first layer of the experiment consists of a large water tank instrumented on the bottom and sides with PMTs. This water tank serves both as a neutron moderator and as an active veto where muons passing through the experiment will emit Cerenkov light which can be detected by the photomultiplier tubes (PMTs), a time based veto can then be applied to the whole experiment.

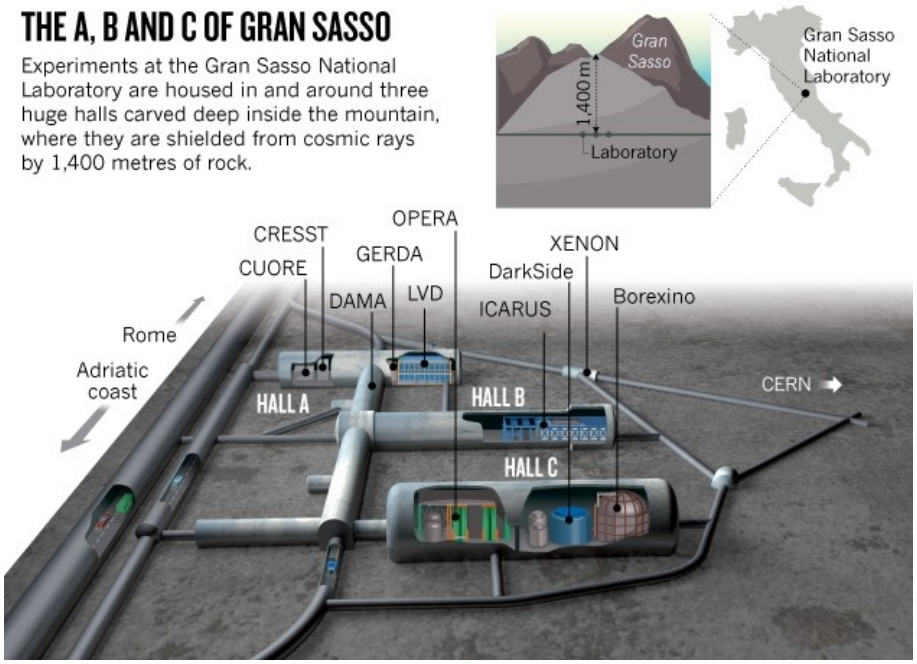

Figure 4.2: Render of the LNGS underground laboratory. LEGEND-200 is located in hall A reusing the GERDA infrastructure. While LEGEND-1000 will be located in hall C taking over the BOREXINO infrastructure, from Ref. [86].

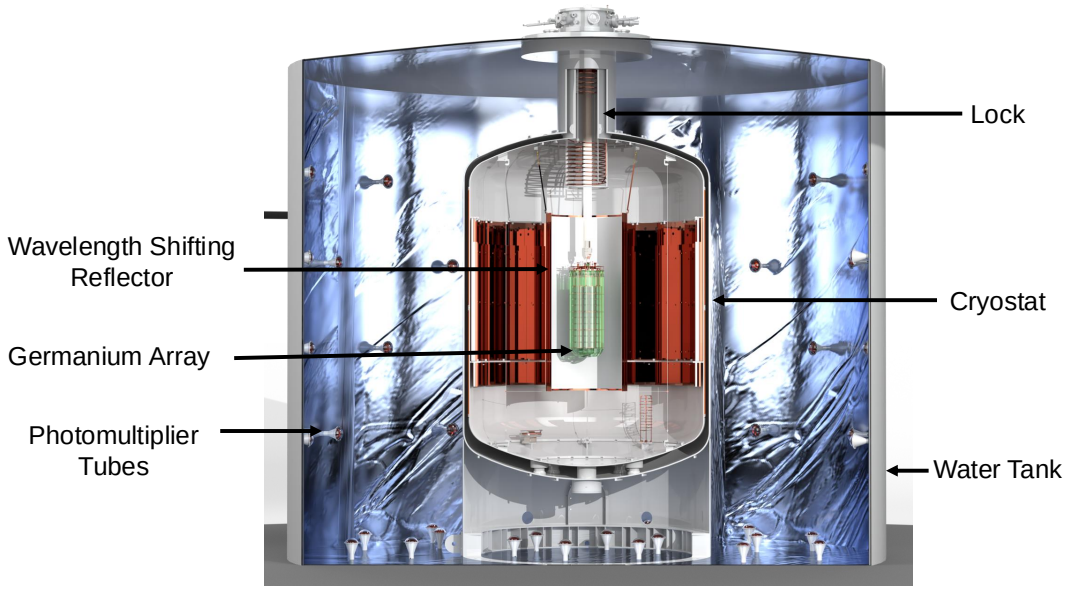

Figure 4.3: Render of the LEGEND-200 experiment with the major components labelled.

The next layer is a large 64 m^3 cryostat containing liquid argon (LAr). LAr is a scintillator so any background event which deposits part of its energy in the LAr will produce scintillation light. This light will be read-out by optical fibres immersed in the LAr which sit both on the inside (the inner barrel) and outside (the outer barrel) of the Ge detector array and are coupled at the top and bottom to silicon photomultipliers (SiPM). Compared to GERDA the LAr subsystem underwent a number of upgrades to improve its light collection including: the addition of this inner barrel, a higher number of fibres and SiPMs, and greater spacing between the germanium detector strings. All these changes should enhance the efficiency of the LAr subsystem.

Finally, at the centre of the experiment sits the Ge detector array. An ideal Ge detector would be large mass to reduce the number of detectors needed with their associated support material and cabling which contribute to background, have excellent energy resolution and PSD capabilities to separate signal from background events, and finally be fully depleted at a reasonable voltage (we require <5000 V). Over time developments in Ge detector technologies have sought to improve these properties. LEGEND-200 utilises detectors from throughout this development and for this reason uses 4 different categories of Ge detectors (with each detector subtly different) which are shown in (shown in Figure 4.4). Two different manufacturers were used for the detectors Mirion and Ortec. These are:

- Semi-coaxial detectors (COAX) originally from IGEX and Heidelberg-Moscow are characterised by a well with a large p+ electrode and an n+ electrode wrapping the outside with the well enabling a large mass detectors at around 2 kg,
- Broad Energy Germanium (BEGe) detectors from GERDA and made by Mirion have instead a small p+ electrode which creates common drift paths for charges enhancing the PSD capabilities when compared to COAX detectors where charges in the detector all follow different path (this will be discussed more in the next chapter), they however can only be made at relatively low masses of around 1 kg,
- Point Contact Germanium (PPC) detectors from MAJORANA and made by ORTEC are similar to the BEGes with a small p+ electrode for the PSD properties but instead of the n+ electrode wrapping totally around have a passivated surface on the bottom of the detector,
- Inverted Coaxial Point Contact (ICPC) detectors which are new for LEGEND and combine the small p+ electrode design with a borehole on the opposite side to enabling larger mass detectors with the improved PSD properties. They are made by both Mirion and Ortec with the Ortec detectors having a smaller p+ contact than Mirion detectors.

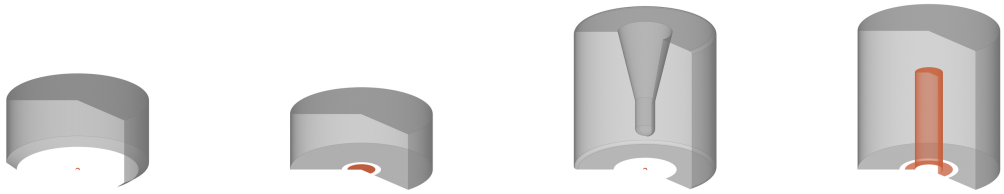

Figure 4.4: Renders of the different types of germanium detectors operated in LEGEND-200. The n+ electrodes are shown in grey while the p+ electrodes are shown in red. From left to right the detector types are: PPC, BEGe, ICPC and COAX.

An additional benefit of these small p+ electrodes is the reduced capacitance of the detectors which leads to lower noise and therefore better energy resolution. These detectors all have different charge collection characteristics and therefore signal shapes which will be discussed later.

The Ge detectors are organised in strings, varying from 5 to 12 detectors per string, which are hung from a top copper plate. They are supported by a baseplate made of Poly Ethelyne Naphthalate (PEN) which is a low background plastic with scintillating properties (shown in Figure 4.5). Each detector is connected to a charge sensitive amplifier which is split into two pieces with the low mass front end (LMFE) consisting of the most radiopure components bonded to the detector and the amplifier (CC4) located further away in the birds nest. The baseplate is held by copper rods, these are screwed together to build the string and along which the cables are run for the high voltage and readout of each detector. Finally, each string is then surrounded by a nylon mini shroud to reduce drift of cosmogenically activated isotopes in the LAr onto the detector surfaces which was a major background seen in GERDA Phase I. 4 additional nylon tubes are also used to guide the calibration sources when they are lowered into the array. At the top of the array, above the copper plate, is the birds-nest where the electronics already mentioned are situated, this ensures they are further away from the detectors to reduce backgrounds. These electronics are connected to the data acquisition system which is located in the cleanroom above the cryostat and consists of FlashCam [87] digitizers recording the waveforms at 16 ns sampling.

The components of the experiment (and therefore the possible sources of background in the array) can be thought of in terms of their distance from the germanium detectors. At close range we have the detector holders, copper rods, LMFE, cables and mini shrouds. At mid-range, the fibres and copper of the inner and outer barrels that comprise the LAr subsystem and finally, at long range, the wavelength shifting reflector (WSLR) which is a copper sheet coated in tetraphenyl butadiene (TPB) outside the outer barrel to reflect light back to it, as well as shifting its wavelength to minimise attenuation.

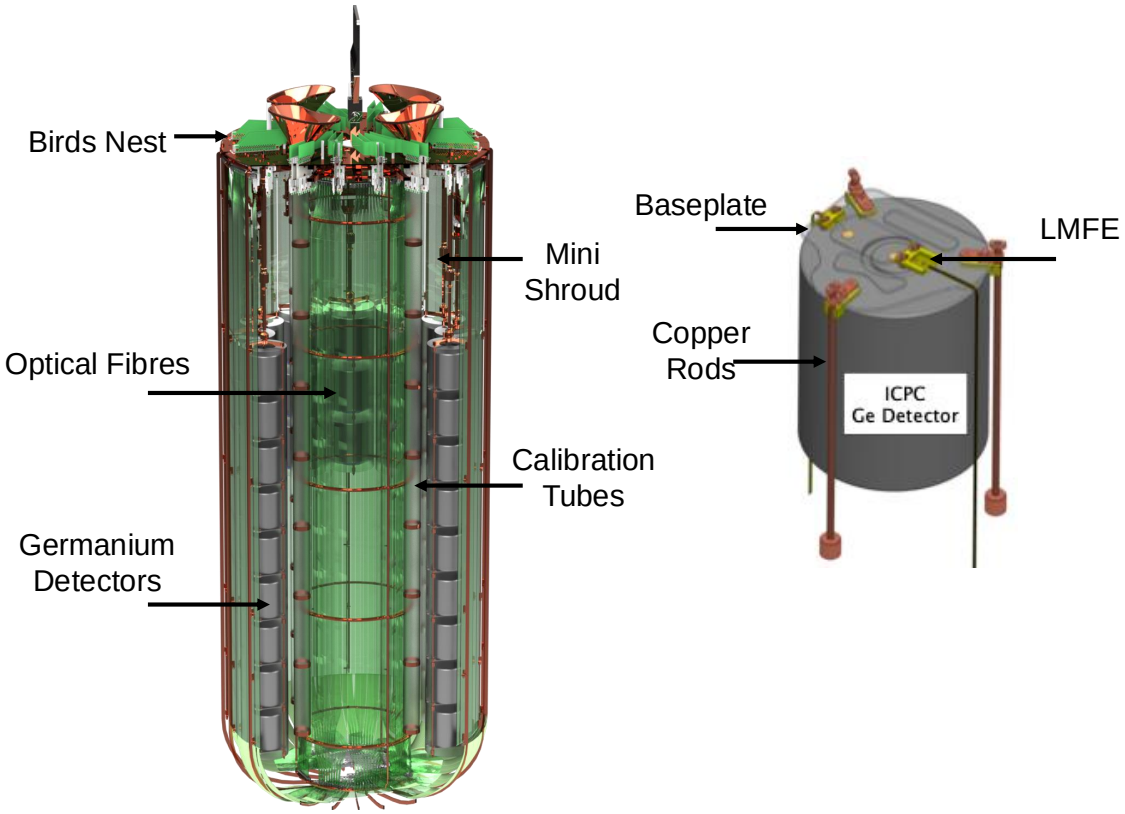

Figure 4.5: On the left is a render of the germanium detector array surrounded by the optical fibres with the major components labelled. On the right is a render of a germanium detector and its holder. This render is inverted compared to the deployment in LEGEND-200 where the plate and LMFE sit on the bottom.

All the materials used in LEGEND-200 underwent a radioassay campaign to ensure low radioactivity components. There are broadly two main types of radioactive contamination that can be present in the materials: a bulk contamination where the material itself has high radioactivity levels, and surface contamination where just the top few microns of the material are contaminated. Surface contamination comes from the handling of the material where radioactive isotopes are introduced while machining or handling the material. To combat this many materials in LEGEND are surface etched with acid to remove the top few microns of material and cleaned before being used in the experiment.

Alongside these passive background reduction techniques, LEGEND also uses active background mitigation, as depicted in Figure 4.6. The design philosophy of LEGEND is to minimise the amount of inactive material in the experiment so that we can use active vetoes to remove events that do not follow the $0\nu\beta\beta$ topology. This is that a candidate $0\nu\beta\beta$ event should deposit 2039 keV of energy in a single detector. Any coincident energy deposition in another detector or LAr subsystem means the event has more than 2039 keV of energy and is therefore not a $0\nu\beta\beta$ event and can be vetoed. Additionally, for a $0\nu\beta\beta$ event the two electrons will deposit their energy within a few mm of each other in the detector, so if our pulse shape analysis sees multiple scatters in a single detector the event is removed. Finally, if the event is near the detector surface we may not see the full energy due to dead layer effects (the outer surface of the detector has a Lithiated layer to form the n+ contact, but this is only partially active) so these events are also removed by PSD. Therefore, we remove at the analysis level any events which do not pass these conditions.

LEGEND-200 was first commissioned using 60 kg of ICPC detectors in 2022, before increasing to 140 kg of detectors at the end of 2022. It entered physics data taking in March 2023 and took data for just less than 1 year. A series of background runs were then taken with various components removed (first the outer barrel and then the mini shrouds) to study the background. Later this year more detectors will be added with the eventual plan to reach 200 kg of detectors sometime next year. The goal of LEGEND-200 is a background index of 2×10^{-4} counts / (keV kg yr) and a discovery sensitivity of 10^{27} yr.

4.3 LEGEND-1000

The next phase of the program is LEGEND-1000. This will be a larger experiment with 1000 kg of Ge detectors. It will still be located at LNGS with the benefits this gives as described above but instead will be taking over the BOREXINO infrastructure in hall C. A render of the experiment is shown in Figure 4.7. The main changes compared to LEGEND-200 will be: the move to just ICPC detectors,

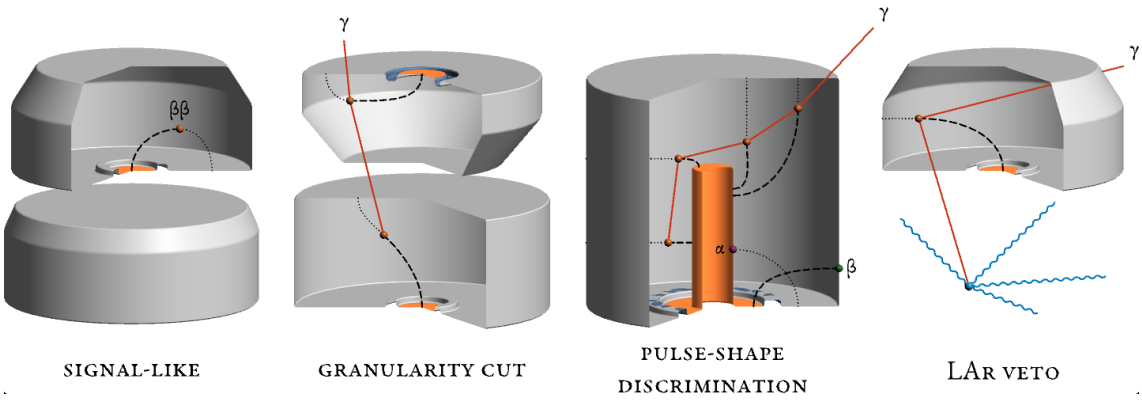

Figure 4.6: Comparison of a signal like event to the background events removed through active vetoing. A $0\nu\beta\beta$ event should consist of a single energy deposit in a single detector as depicted on the far left. Background events can be removed from left to right via: the granularity cut when a hit is seen in multiple detectors, LAr anticoincidence when scintillation light is seen alongside the energy deposit in the Ge detector, and finally PSD cuts which separate single from multiple energy deposits in a single detector.

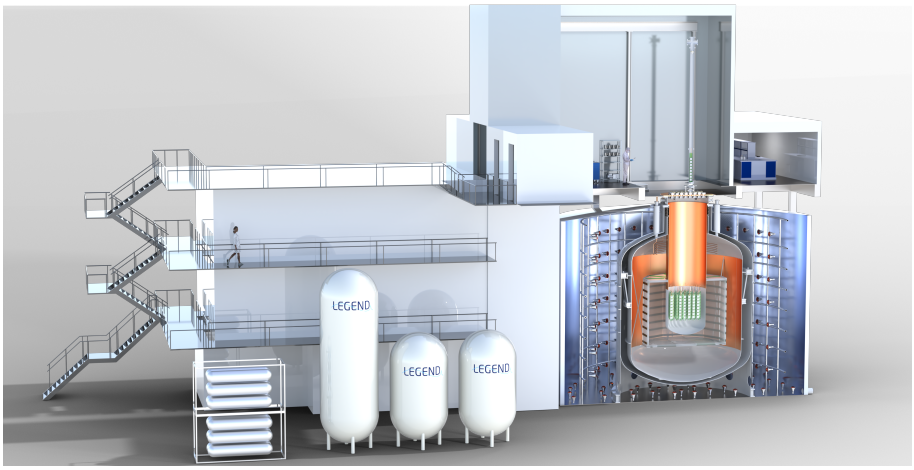

Figure 4.7: Render of the future LEGEND-1000 experiment. The detector is located on the right with the cleanroom for assembly above.

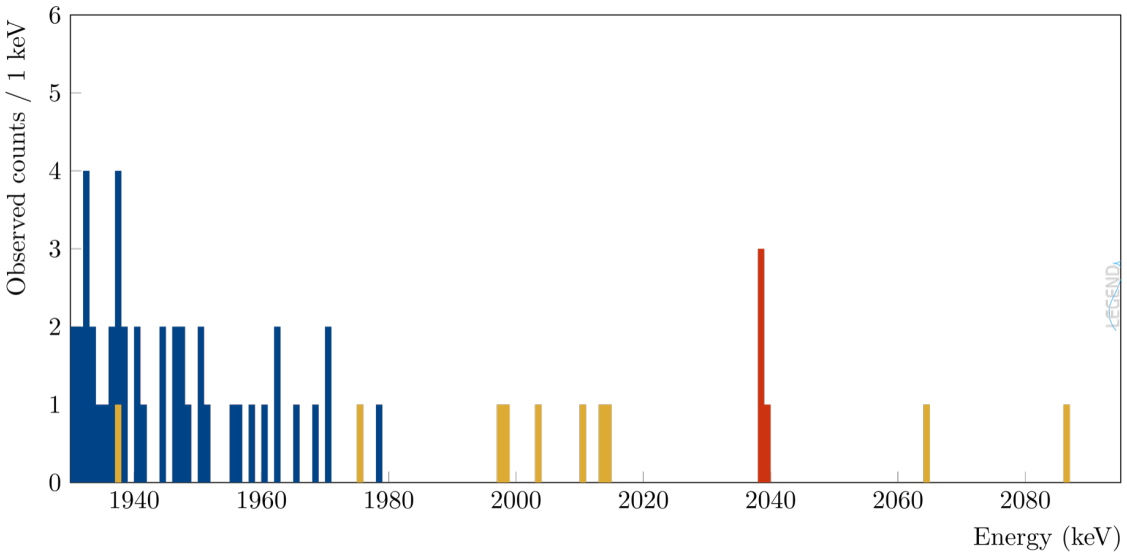

Figure 4.8: Example spectrum of the LEGEND-1000 experiment after 10 years of data taking with a signal at 10^{28} yr. This shows one of the advantages of the experiment where a signal can easily be seen by eye above the background.

which will reduce the background as fewer detectors will be needed so less support material and less cabling and also simplify the analysis, the use of a single string lock design which will ease maintenance and allow issues in single strings to be fixed without lifting the entire array enabling a greater uptime and number of working detectors, and most importantly the innermost volume inside the copper reentrant tube will be using underground liquid argon to reduce the background from cosmogenically activated isotopes in the LAr. On top of this, all the subsystems will see many upgrades from: upgraded electronics with the move to ASICs from LMFEs, a fibre shroud around each detector string and an improved neutron tagging system. An example of what the spectrum could look like after 10 years of data taking, when 10 ton yr of exposure is reached, is shown in Fig.4.8. The ultimate discovery sensitivity of the experiment will be 10^{28} yr with a background index of 10^{-5} counts / (keV kg yr) .

LEGEND, therefore, is a very exciting program aimed at discovery of $0\nu\beta\beta$, with scalable and well understood technology while the low background and great energy resolution would make for a clean, unambiguous signal.

Chapter 5

HPGe Detector Modeling and Analysis Routines

In this chapter I will describe the analysis routines for the germanium detectors used in the LEGEND-200 experiment. I led and personally carried out most of the design, implementation and testing of these routines which drew on the experience of MAJORANA and GERDA. Some parts such as the pulse shape analysis are very close to the GERDA approach [88] while others such as the data cleaning and energy optimisation are almost entirely new.

5.1 Ge Detector physics

5.1.1 Signal Formation

The detectors used in LEGEND are High Purity Germanium (HPGe) crystals enriched up to 92% in the isotope of interest, Ge-76. These detectors are grown using the Czochralski method where a seed crystal is slowly pulled out of the molten material while rotating [89]. This produces an exceptionally pure crystal with a low concentration of impurities at the level of 1 part in 10 million (although some impurities are required to give the necessary charge drift characteristics).

Two electrodes: a small p+ electrode, made through boron implantation and a larger n+ electrode, made through lithium diffusion, are added along with a groove between the two which is passivated to reduce leakage current. The dead layer will grow with time due to continued diffusion of the Li and this depends on the temperature of the detector. In LEGEND, we assume no dead layer growth while in the cryostat but during storage the detectors are kept around room temperature, so the dead layer will grow. A large reverse bias is applied to the n+ electrode to fully deplete the crystal, the size of which will depend on the geometry of the crystal and its impurity profile. A diagram of a detector is shown in Figure 5.1.

When a particle interacts with the crystal, it creates, through ionisation, electron-

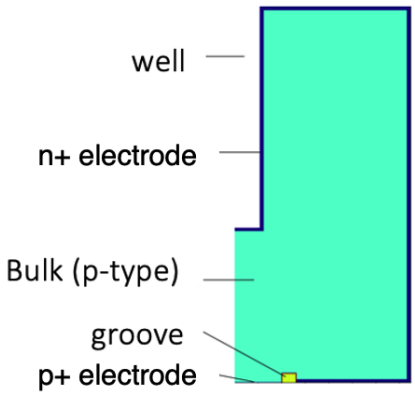

Figure 5.1: Diagram of the main parts of a germanium ICPC type detector.

hole pairs (the number of which is proportional to the energy deposited) which are then drifted to the electrodes by the electric field. Germanium's low band gap at only 0.67 eV at 300 K [90] means that the number of charge carriers created is very large which is why germanium has such great energy resolution. Holes drift to the p+ electrode while electrons drift to the n+ electrode. Their drift induces on the read-out electrode a change of voltage which translates into a current signal described by the Shockley Ramo theorem [91, 92].

$$I = q\ v\ E_v \tag{5.1.1}$$

Where I is the current induced, q is the charge, v the velocity and E_v the electric weighting field. As the electrons move twice as fast (due to the effective mass of the holes being twice as large) and generally have a shorter path to the n+ electrode, they are collected quickly. They therefore contribute negligibly to signals in the bulk of the detector [89]. The holes take longer to drift to the p+ electrode and as they pass through the strongest electric field form the majority of the signal. The charge sensitive amplifier then converts this current into a voltage signal giving a characteristic exponentially rising shape as shown in Figure 5.2.

In LEGEND-200, we digitise a 148 μs long waveform with this rising edge section at its centre. This ensures we have a baseline section before the pulse of around 45 μs and a tail section after the pulse of around 68 μs. This is shown in Figure 5.3. The tail section consists of an exponential RC decay. When processing this signal the main quantities to extract are: the amplitude of the signal which is proportional to the total energy release within the detector, the start/end time of the rising edge as it is proportional to the time of the energy deposition and its distance from the electrode thus needed for vetoing and timing based correction and the maximum current and other shape information which is used in the pulse shape discrimination. Other quantities related to the physical nature of the signal are also calculated.

The first steps in the processing of the signals is to shift them down to 0 and

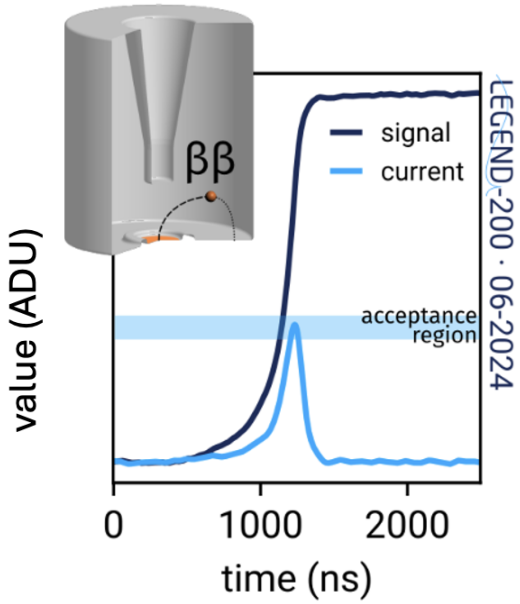

Figure 5.2: Example of an event in a germanium crystal and the associated voltage generated in dark blue with the current in light blue. The interaction point is shown with the holes drifting to the p+ electrode in orange and the electrons to the n+ electrode outside.

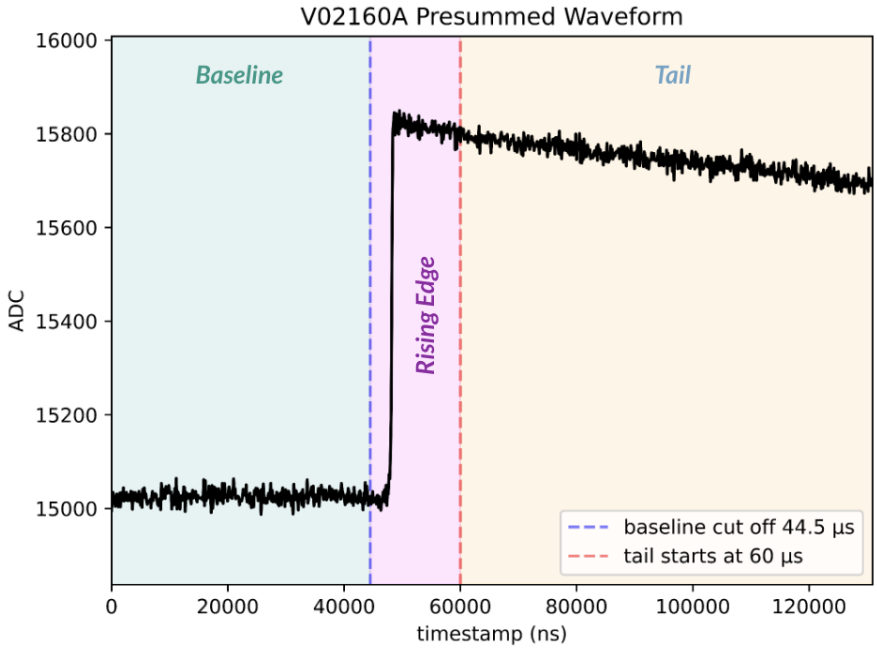

Figure 5.3: Example waveform with the 3 main sections. The waveforms in LEGEND-200 are 148 μs long, the first 45 μs is the baseline section followed by 15 μs of leading edge and 88 μs of tail. The shape of the rising edge is determined by the interaction point in the detector. For most interaction points the electron contribution is negligible and the hole contribution dominates. The rising edge section is set conservatively large as slow pulses from surface betas can be several μs long.

correct the RC decay (pole zero correction). Once this is done, more advanced filters can be run over the signal. Shown in Figure 5.4, as an example, is the filter used to extract the start point (t0) of the signal which is using a modification of an asymmetric trapezoidal filter I developed.

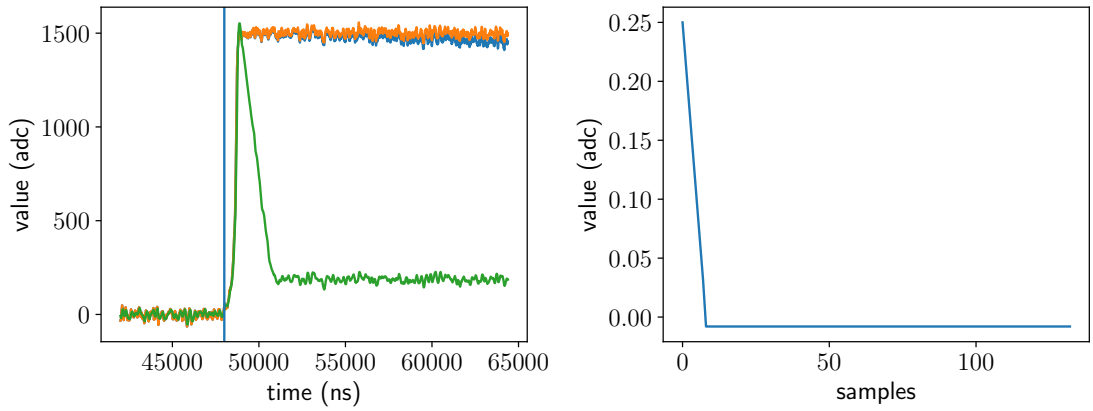

Figure 5.4: Example of the digital signal processing of a waveform. Here in blue is the waveform after baseline subtraction (bringing baseline section to 0), in orange is it after the pole zero correction to remove the RC decay and in green is the output of the t0 filter. Finally, the vertical line shows the t0 value which is obtained by walking down the filter output until a threshold is reached. On the right is the kernel of the t0 filter used which is convolved with the pole zero corrected waveform. It consists of a linear rising section to average the leading edge while preserving the shape followed by a negative flat section to average out the noise.

Every detector operated in LEGEND-200 is slightly different due to the manufacturing method. We have 2 detector manufacturers Mirion and Ortec who produce subtly different detectors. The shape and size of each detector is informed by the concentration gradient of the impurities in the crystal which means that every detector has a different electric field and therefore different charge collection characteristics, meaning the signal shape generated is different for each detector. Consequently, the analysis of the detectors needs to account for this by developing routines to optimise the performance individually. These routines need to be exceptionally reliable and robust with some such as the calibration applied every week on over 100 detectors; this means in the last year they were applied over 4000 times. Additionally, they need to be able to handle well a broad spectrum of problems which may arise when applying to such varied detectors.

5.1.2 Energy Reconstruction

When an interaction occurs in the crystal a number of electron hole pairs are created. Due to the low band gap in germanium (0.67 eV at 300 K [90]) this number is very large which is why germanium has such good energy resolution. Naively, the number of charge carries created would be a Poisson process with

mean $\lambda = E_i/E_g$, where E_i is the energy deposited and E_g is the germanium band gap, and therefore the energy resolution would be the square root. However, in reality the resolution is even better as the creation of these electron hole pairs are not independent. This improvement in energy resolution is encapsulated in terms of a fudge factor known as the Fano factor.

On top of this fundamental limit on the energy resolution, there are 2 other factors that can degrade the energy resolution the electronics noise and incomplete charge collection. Incomplete charge collection effects consist mainly of charge trapping (where defects in the crystal trap an electron or hole and release it after some time delay) and ballistic deficit (when a filter is too short it can integrate only part of the charge). The first is being increasingly mitigated by the producers and is nowadays sub-dominant, any lingering effects can be mitigated through corrections while the second is almost completely removed with good filtering choices.

Finally, the electronics noise affects the accuracy in which we can estimate the number of charge carriers created. Ideally it would be completely white where all frequencies are equally represented and therefore easily minimised. However, in reality, it will, depending on the experimental setup, vary with frequency which will determine the optimal filtering to use in reconstruction. Putting this all together the total energy resolution is:

$$FWHM^2 = FWHM^2_{det} + FWHM^2_{noise} + FWHM^2_{charge} \tag{5.1.2}$$

where $FWHM_{det} = 2.35\sqrt{\epsilon.F.E}$ is the detector resolution (ϵ is the germanium ionisation energy, F is the Fano factor and E is the energy deposited, the 2.35 is the conversion factor from sigma to FWHM), $FWHM_{noise}$ is the electronic noise and $FWHM_{charge}$ is the charge collection resolution [93].

In our pulses the height of the pulse corresponds to the energy deposited. Therefore, the simplest estimate of the energy would be the difference between the maximum and minimum of the pulse. However, this would mean the noise term above would be large as the energy output would be dependent on the statistical fluctuations of the noise. Therefore, to mitigate this term and the charge collection term we make use of optimal filtering and corrections. These terms vary from detector to detector which motivates the development of routines to optimise the energy resolution for each detector. This also motivates the decision to record as long traces as possible using the DAQ. Longer pulses mean more samples to characterise the different frequencies that are present in the noise and remove them.

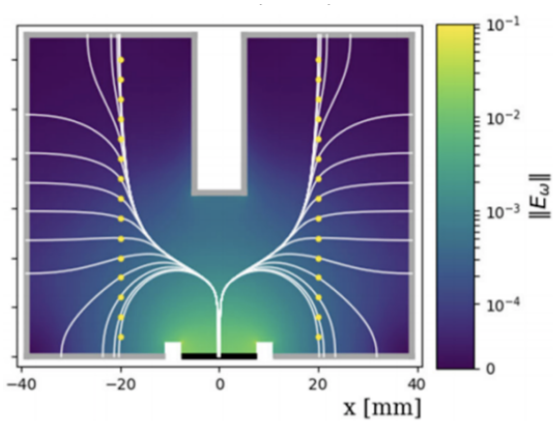

Figure 5.5: Weighting field of an ICPC detector, with lines showing the drift paths of electrons/holes [94].

5.1.3 Pulse Shape Discrimination

GERDA and MAJORANA pioneered the use of small p+ electrode detectors with their excellent pulse shape discrimination (PSD) properties. The benefit of these detectors can be seen from its weighting field (Figure 5.5).

Due to this, the holes are funnelled into the middle of the detector and then towards the p+ electrode. This gives the detector powerful PSD properties as all charges follow a common path. Figure 5.6 shows how the signal generated for a single site event (SSE) such as neutrinoless double beta decay is different to that of a multi-site event (MSE) which would be a background e.g. from Compton scattering. This motivates the use of the A/E parameter which is the maximum current of the pulse divided by the energy. An acceptance band can be defined for single site events and anything outside this band cut. Multi-site events will be lower than this band as for a given energy they will have a lower maximum current. In addition, we can also identify surface events from either alphas on the p+ electrode where the holes are collected very quickly leading to a very fast rise and an elevated A/E or betas on the n+ electrode which give rise to very slow events due to diffusion in the semiconductor region where the field is null.

As the charge cloud drifts collective effects will affect its shape. In particular self repulsion of the charge cloud will cause the charge cloud to spread out lowering the A/E, and stochastic diffusion means there is some amount of randomness in the path expanding the single site band. This leads to an energy dependence in the position and width of the single site band as more energy means more charge carriers and a wider charge cloud. Additionally, there will be some position dependence as events further from the p+ contact will take longer to drift in the detector and therefore are more likely to have charges removed by trapping. As with everything else this will be detector dependent and means that an A/E calibration will be needed for each detector.

The calibration of a cut on the A/E parameter to discriminate between MSE and

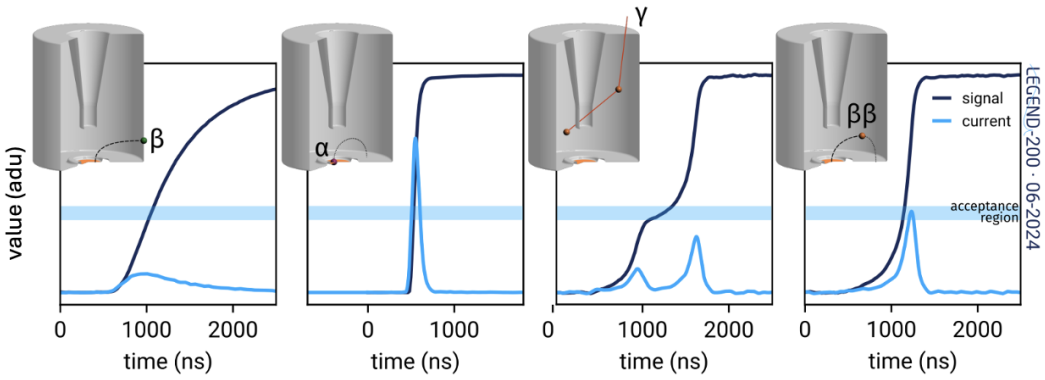

Figure 5.6: Comparison of the signal generated for single-site, multi-site events and surface events. The recorded voltage signal is shown in dark blue with the corresponding current signal in light blue. The energy of an event is proportional to the amplitude of the signal. On the left are surface events composed of betas on the n+ electrode which are characterised by a slower signal due to slower charge diffusion in the Lithiated n+ surface and therefore a lower peak current, and alpha events on the p+ contact on the p+ electrode which have a very fast signal and therefore a high peak current as the holes are collected very quickly. Second from the right are multi-site events which are events with multiple scatters in the same detector. These differ from single site events (on the furthest right) in that they have multiple peaks all with a lower peak current for a given energy.

SSE can be performed using calibration datasets, in particular a good source of single site events is a double escape peak, DEP. These lines are produced when the primary gamma undergoes pair production; the electron's energy is deposited while the positron annihilates in the detector producing two 511 keV gammas back to back. For the single escape peak, SEP, one of these gammas escapes the detector and the other is absorbed, while for the DEP both escape. This gives these peaks unique topologies which can be used for PSD calibration. The DEP is just a single energy deposit in the detector and therefore is a proxy for single site events such as $0\nu\beta\beta$. The SEP is composed almost entirely of multi-site events which are background-like events. However, one thing to bear in mind is that these events are not homogeneous in the detector volume as the probability for gammas to escape is higher at the edges and corners. The final type of gamma line is the full energy peak (FEP) where even if pair production occurs the full energy is deposited in the detector. These are generally more multi-site like due to pair production.

In Figure 5.7 we can see the different shapes in the A/E spectrum for different event populations. For the DEP we expect a gaussian shape in the A/E spectrum as it consists mainly of single site events. Compton events will be a mixture of single and multi-site events depending on how many times the gamma has scattered. We therefore see a single site gaussian with a tail of multi-site events which we want to remove with the cut.

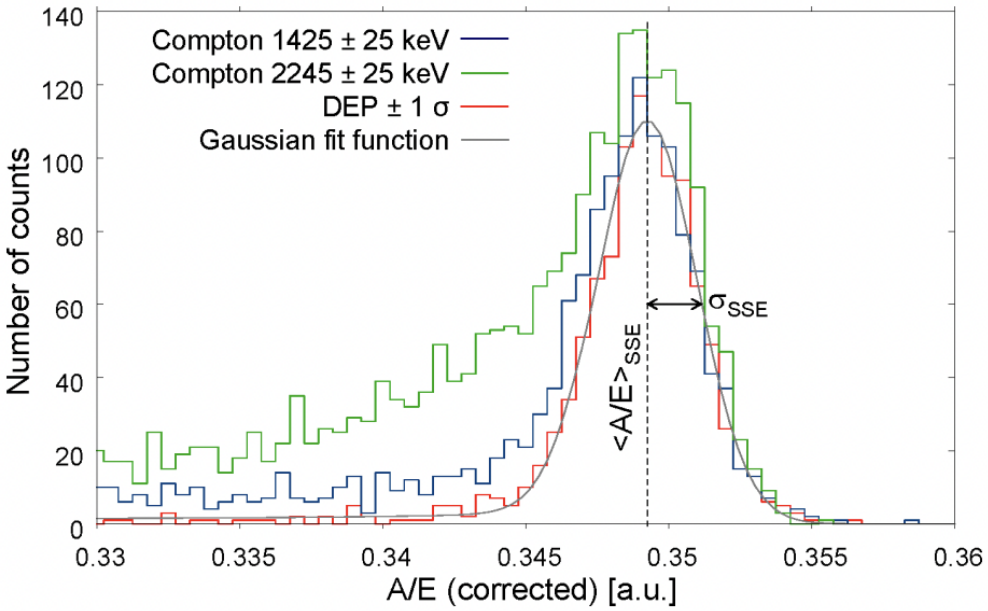

Figure 5.7: Example of the shapes in the A/E spectrum for different event populations in a Th-228 calibration. The DEP is a single site proxy so has a gaussian shape while the Compton bands have both a multi and single site component so have a gaussian with a tail to lower A/E values. SSE here is single site event. Two different Compton regions, below and above the DEP, are plotted to show the difference in peak to tail ratios for different energy regions (below the DEP the peak to tail ratio is higher than above only a single gamma is scattering in the detector while above multiple gammas are scattering so multi-site events are more likely). From Ref. [95].

5.2 Data Cleaning

One of the first steps in the analysis of the data is the cleaning. This work was carried out by a team of people with my focus on the choice of the parameters and digital signal processing routines used, as well as the implementation of the cut logic. The data cleaning is essential to remove non-physical events such as discharges, micro-discharges and noise bursts. In LEGEND, we could potentially claim a discovery of $0\nu\beta\beta$ with a single digit number of events, therefore it is important that the only events we keep are fully physical on all channels (no discharges or noise bursts). Even if one channel has a normal looking pulse if another channel has a non-physical pulse it is not possible to be certain that the non-physical pulse is not affecting the other channels. Therefore, we require, for an event to be valid, that all channels have either physical or empty signals.

A complicating factor is the presence of crosstalk, this is still a valid signal but can be a different shape to physical signals. The main type of crosstalk we see is negative going and looks like an inverted pulse, although some can be positive pulses also (an example is shown in Figure 5.8). We therefore have to make sure that we are not removing these events.

Finally, it is important that all cuts, if possible, are energy independent so that we do not have varying efficiencies. This could cause issues, for example in the background model where a varying efficiency could change the relative intensity of gamma lines causing the background to be misconstructed.

The cuts are applied to both calibration and physics data. The major difference here is the presence of pileup in calibration data. This is where 2 events are close together in time so either a second signal is seen in the trace or the signal sits on the decay tail of a previous event. In all cases the energy of the main signal will be misreconstructed. We could add routines to reconstruct the true energy, but as the event rate is so low in physics data that these events are not expected, we wish to remove these events. This also ensures our calibration sample passed to subsequent routines is similar to physics data.

Shown in Figure 5.3 is an example waveform from a HPGe detector. We split the waveform into 3 main sections: the baseline section which corresponds to the 44.5 μs before the trigger and should be flat and featureless, the leading edge which is the 15.5 μs around the trigger and the tail which is the last approximately 60 μs of the waveform which should be dominated by the pole zero decay. The leading edge window is generously sized to allow for the different event topologies we expect such as surface betas which can have slow pulses of several μs.

The quality cuts are applied in sequence starting from the simplest cut and working up. The cuts in the order applied are:

- Discharge cut

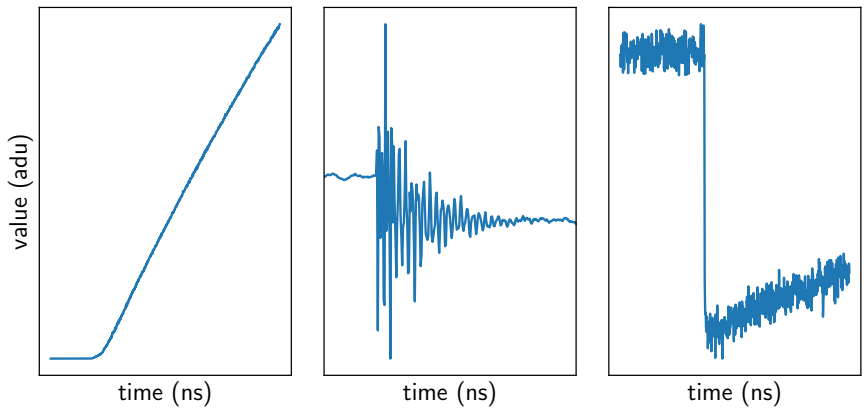

Figure 5.8: On the left are two examples of unphysical waveforms we want to remove: a discharge where the ADC has hit 0 and, in the middle, a noise burst. On the right is a crosstalk event which we don't want to remove but instead to flag. The noise burst here is shown for the windowed waveforms around the leading edge while the others are the full waveform.

- Baseline and tail cuts
- Leading edge cuts

The first two are applied to all waveforms while application of the leading edge cuts varied based on the amplitude of the pulse. I will now describe each of these cuts in more detail.

The simplest cut to apply is to remove discharges, these are identified by a flash ADC channel hitting 0. When this is seen we apply a time veto on all channels after this to remove any recovering signals. It was found using special runs that a 10ms window was sufficient to remove all these types of events. Applying a time veto is also advantageous as this simply adds a dead time into the array the efficiency of which can be easily calculated. Additionally, vetoing the whole array makes sense as often times a discharge on one channel is associated with unphysical events on neighbouring channels. Possibly, in the future, the dead time could be decreased by only vetoing associated channels e.g. same string and/or CC4 (amplifier).

The next cuts are common between empty traces and those with an energy deposition. In both cases we want the baseline and tail sections to follow the expected shape (for the baseline flat in both cases, for the tail either flat or exponentially decaying) with no unexpected features such as small energy depositions. 3 parameters were found that satisfied the energy independent requirement: 2 for the baseline and 1 for the tail (detail in appendix). The baseline parameters were chosen by injecting exponential decays and simulated pulses onto empty traces and investigating parameters which could separate these from the remaining empty traces. In all cases the distributions are gaussian with a tail of non-

physical events. Cuts are determined by performing a fit to the top of the gaussian to extract the centroid and standard deviation and then a double-sided cut is applied at 4 standard deviations from the mean. This method was chosen due to its robustness and simplicity in comparison to more complex fitting methods which were more prone to failure, especially when applied to a large number of detectors.

The cuts effectively remove pileup events in calibrations where 2 energy deposits are close together in time. You can see the effect of this cut clearly on the 2.6 MeV peak where the low energy tail is reduced and even removed in some cases (Figure 5.9). The main type of pileup removed is where the energy deposit sits on the exponential tail of a previous pulse which always results in an underestimate of the energy of the main pulse leading to a tail forming on the low energy side of the peak (Figure 5.10).

The final set of cuts are on the leading edge. For empty traces we simply check that the energy is in line with empty events as determined from forced triggers. Whereas for physical signals, a series of time cuts are applied to ensure the event is of standard shape and will be reconstructed correctly. These cuts were developed using calibration data and are placed conservatively to ensure no physical pulses are removed and that there was no overlap with the pulse shape analysis. Simulated pulses were also used to investigate signal-to-noise dependencies. We check: the pulse is in the position expected in the trace, the point the energy is sampled is in the position expected, the rise and drift times of the pulse are reasonable and that there are no oscillations in the initial section. Additionally, there is an energy cut off for these cuts where they are only applied above 100 standard deviations of the forced trigger energy distribution. Below this we do not expect the start point of the pulse to be reconstructed correctly due to the lower signal-to-noise. Only routines with a high degree of smoothing can be applied at these lower energies so only the requirement on the energy sampling is retained. This does introduce a discontinuity in the efficiency with energy, but it is well below the energy threshold of current analyses. For future analyses looking at this low energy region a different set of cuts should be used.

5.3 Energy Reconstruction

An accurate and high resolution energy reconstruction is essential for the analysis of the data. We tested a number of different filters and decided the CUSP filter [96] was the best for our noise environment as it gave the optimal energy resolution. To improve the resolution further we performed an optimisation of the parameters of this filter using calibration data and also implemented a charge trapping correction to remove the detector position dependencies. This optimised energy

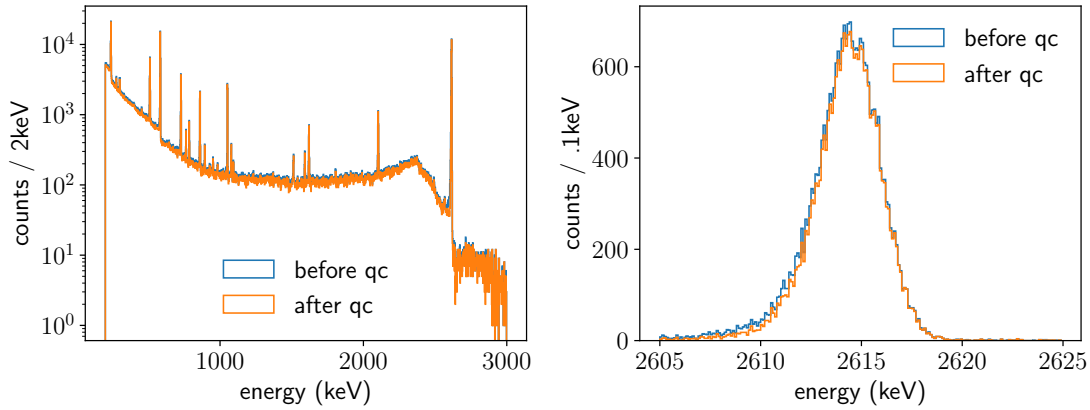

Figure 5.9: Example calibration spectrum for an ICPC detector with and without quality cuts (qc), for the full spectrum on the left and for the 2.6 MeV Tl-208 gamma line on the right. The effect of the quality cuts in removing pileup events on the left tail can clearly be seen.

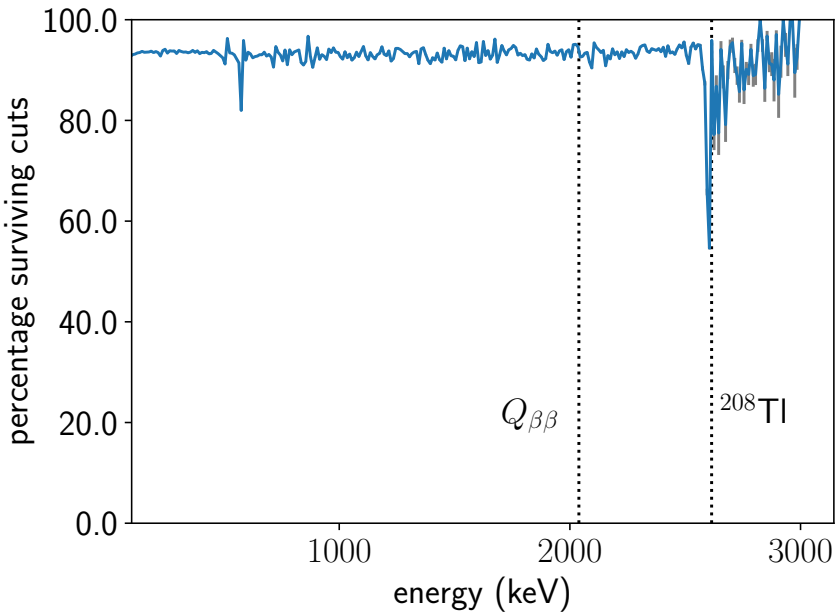

Figure 5.10: Survival fraction vs energy for one detector in calibrations. The errorbars, shown in grey, are increased by a factor of 10 to make them visible above the Tl-208 FEP. The survival fraction is generally flat in energy with dips in the tails of the major gamma lines from pileup events. Above the 2.6 MeV Tl-208 gamma line we are limited by statistics.

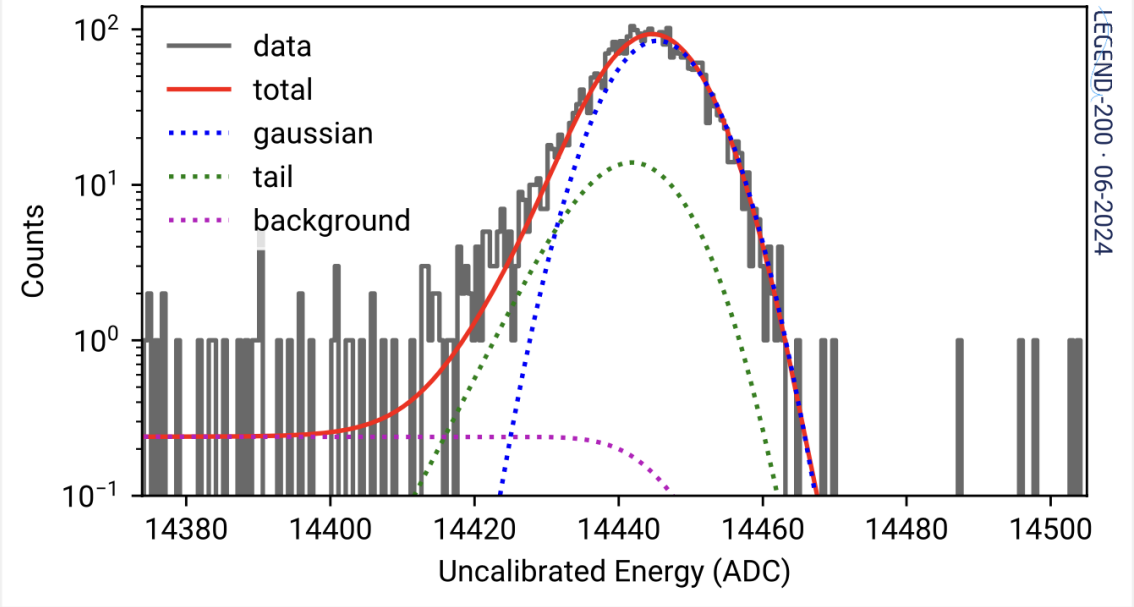

Figure 5.11: An example of a peak fit for the 2.6 MeV Tl-208 gamma-line for 1 detector. The 3 components can be seen: the Gaussian peak with an exgaussian tail and the step background. The exgaussian tail is used to model incomplete charge collection effects in the detector e.g. from charge trapping. While the step models the background from Compton events. In this case a tail was found to be needed for this detector but in some cases this is dropped.

reconstruction was done individually for each detector as they each have unique charge collection and noise properties. The optimal filter was then applied to all waveforms to obtain a raw energy estimate.

The next step is to calibrate this energy estimate into a physical quantity. Naively we would expect that this should be a simple scaling factor as the energy is directly proportional to the number of charge carriers. However, from studies using the pulser and calibration data we found that there were some small non-linearities in the system, and therefore, we determined that a quadratic treatment was necessary. This means the calibrated energy is determined using the following equation:

$$E_{cal} = a + bE_{uncal} + cE_{uncal}^2 \tag{5.3.1}$$

In a stable period of data taking these non-linear terms are expected to be constant in time. A two-step energy calibration procedure was therefore developed. Firstly, for each calibration run, the Tl-208 2.6 MeV peak was fitted to determine the gain of the system as shown in Figure 5.11. This peak was used as it is the most prominent high energy peak in the spectrum and is present in all detectors. It is also reasonably close to the $Q_{\beta\beta}$ value of 2039 keV. From this, a simple scaling could be applied to correct for this gain. Secondly, all calibration data in a stable period were grouped together and all possible peaks fitted. The non-linear terms could then be extracted using a fit to these positions and the peak energy

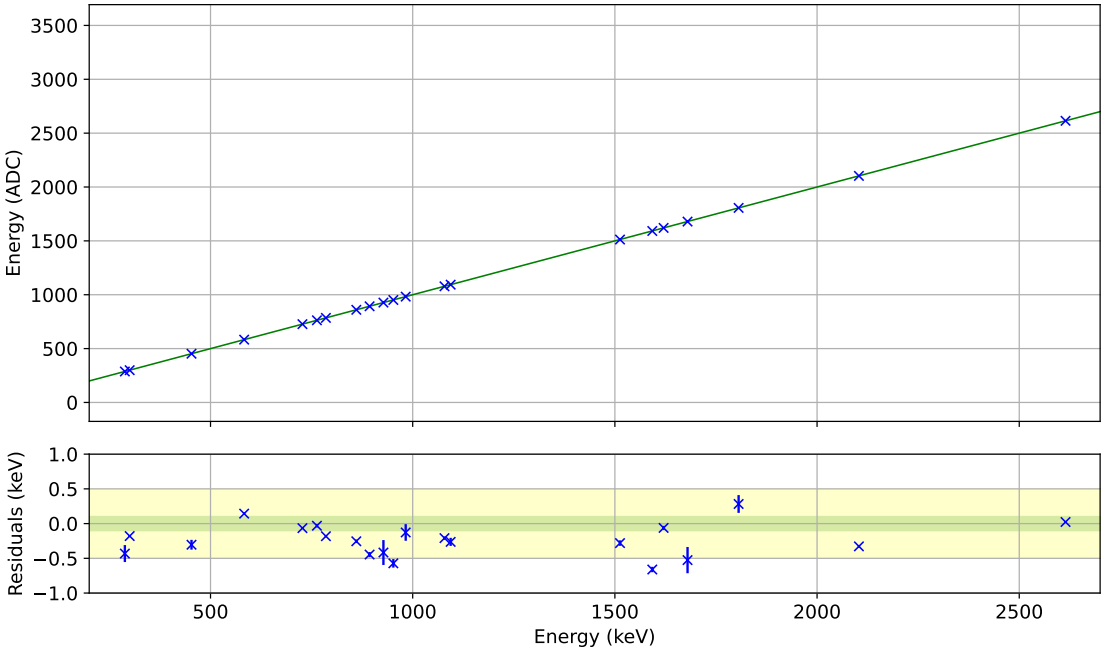

Figure 5.12: Example calibration fit for 1 detector. We fit a quadratic function to the distribution of peak positions vs expected energy. The residuals of this fit are shown in the bottom panel.

as shown in Figure 5.12. By combining the calibration runs in this way many more peaks could be used in this fit, due to the increased statistics, which helped in the determination of the non-linearities.

For detectors put in anti-coincident/veto mode (not deemed suitable for full analysis) only the first step of this energy calibration was applied as the more precise calibration is not required. This was done with a simple peak top fit as often the fit would not converge due to the differing shapes of the peaks.

Additionally, the energy resolution of the detector is also determined by fitting the variation of the energy resolution of the peaks with energy. A simple interpolation can then be performed to extract the energy resolution at $Q_{\beta\beta}$. This is an important input for the statistical analysis and for simulations of the data. Finally, the validity of both the energy resolution and energy calibration were checked using peaks in the physics data which will be discussed in further chapters.

The last important part of the energy reconstruction is the crosstalk correction. This is essential for the multiplicity determination as a signal on one channel can cause a smaller signal in another channel causing a mislabelling of a multiplicity one event as multiplicity two if not corrected for. This is also important for the background analysis where crosstalk can cause multiplicity two events to move outside the gamma peaks throwing off the background model. Special calibration runs are taken in global mode, so all detectors are triggered to see the effect of a pulse on other channels. From this a crosstalk matrix is built using the method outlined in [97] (example in Figure 5.13). The crosstalk is an event level quantity

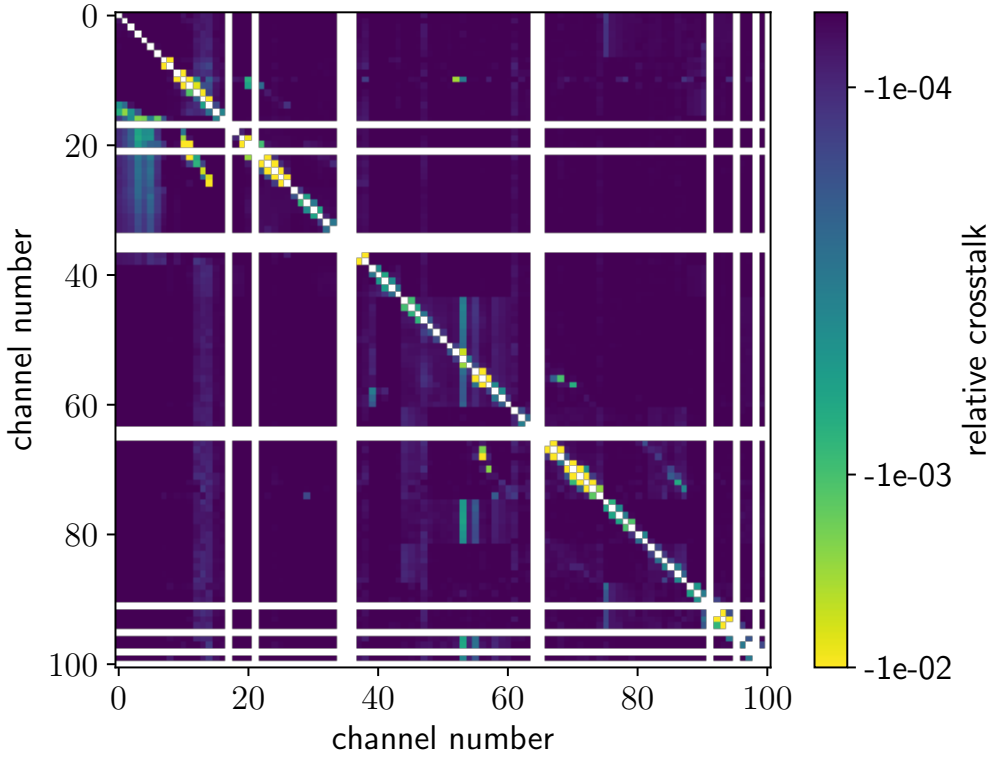

Figure 5.13: Example crosstalk matrix for LEGEND-200. This shows the expected amount of crosstalk between two channel pairs. On the y-axis is the triggered channel while on the x-axis is the response channel.

where all deposits on all germanium channels are needed. A simple correction is then applied to the uncalibrated energy for channels above a certain threshold (in our case 50 keV) using this matrix and then all the channels have the calibration reapplied. The multiplicity of the event can then be determined after this correction.

5.4 Pulse Shape Discrimination

The final step in the calibration of the detectors is the pulse shape discrimination. This is crucial for separating signal events from background. The main parameter used in the PSD is the A/E parameter. This is the maximum current of the pulse divided by the energy (for the PPCs a subtly different parameter was used based on the maximum slope, this is essentially equivalent but with a larger averaging window).

The calibration of the A/E is needed to remove any energy, position or time dependencies and to determine the cut value. This was done using all the calibrations in a stable period combined to give sufficient statistics in the Compton region and in the single site proxy, the DEP. In LEGEND due to the increased

use of ICPC detectors a drift time correction was added for these detectors. This was to ensure we were not biased in the regions of detector being accepted or rejected. The drift time correction was determined from DEP events by fitting in both the drift time and A/E distributions. The energy correction followed the GERDA approach [88] of slicing the Compton region into 20 keV bands and fitting the A/E distribution for each band to extract the centroid and width. These were both then fit vs energy and a classifier built that was independent of energy as shown in Figure 5.14. Finally, the cut was determined as the value that gave a 90% efficiency

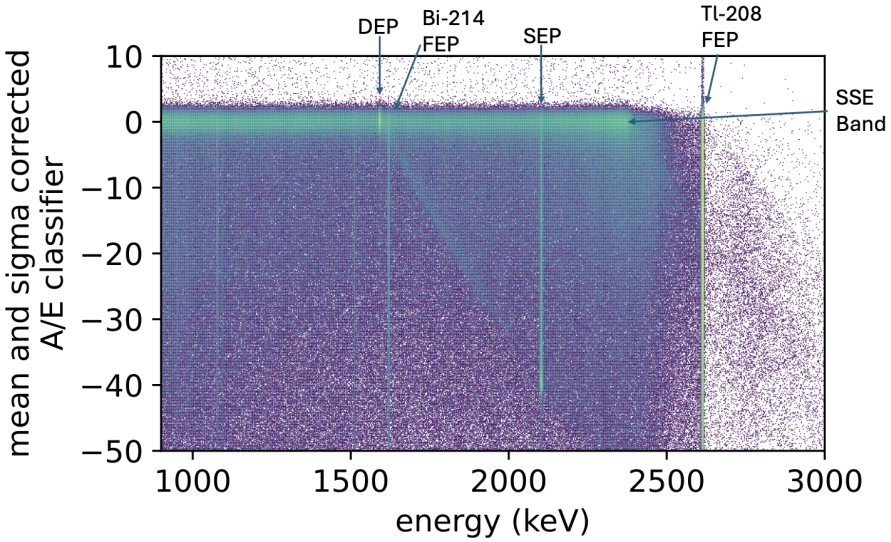

Figure 5.14: A/E Classifier for a single detector. The single site event band is shifted to 0 by our calibration with the gamma lines extending down from it. The DEP can also be seen in this band located at 1592 keV. A diagonal band also extends from the DEP to the SEP composed of pair production events. The gamma lines can also be seen to extend in some cases above the single site band due to events near the p+ contact.

for the DEP. This was determined using fits of the distributions of the surviving and failing events in energy for various cut values (shown in Figure 5.15).

The effect of the cut can be seen in Figure 5.16 where the Tl-208 DEP is preserved while the Bi-214 gamma line at 1620 keV is suppressed. A high cut is also added to remove surface events. This is fixed at a value of 3σ from the single site band. Whereas for the PPCs and Ortec ICPCs it was found that a different parameter, late charge (LQ), was more effective as used in the MAJORANA DEMONSTRATOR [85]. As for the other steps, these calibrations were calculated for each detector individually. The performance was checked using the peaks and Compton region around $Q_{\beta\beta}$ in the calibration spectrum and the $2\nu\beta\beta$ region in physics data.

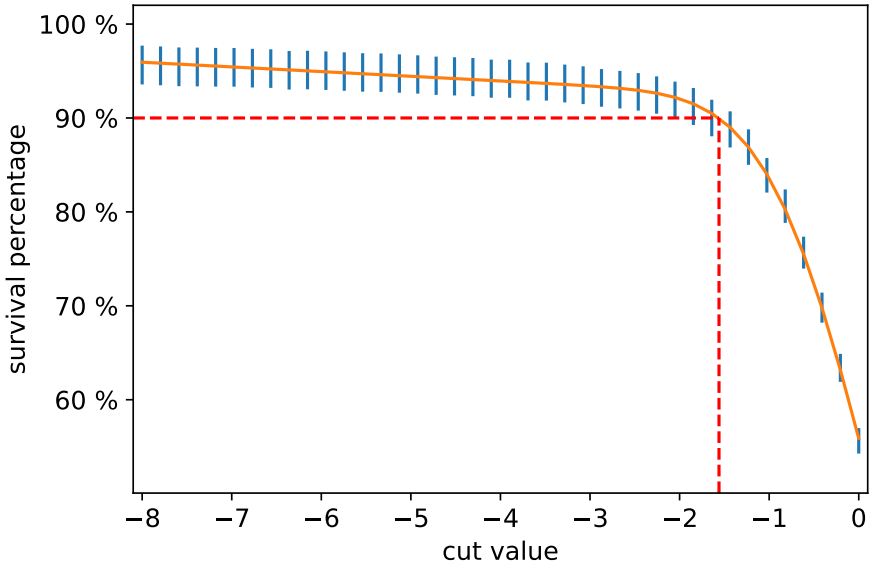

Figure 5.15: An example of the fit to determine the A/E cut value. We sweep through various cut values, calculating the survival fraction for each and then fit with a sigmoid function.

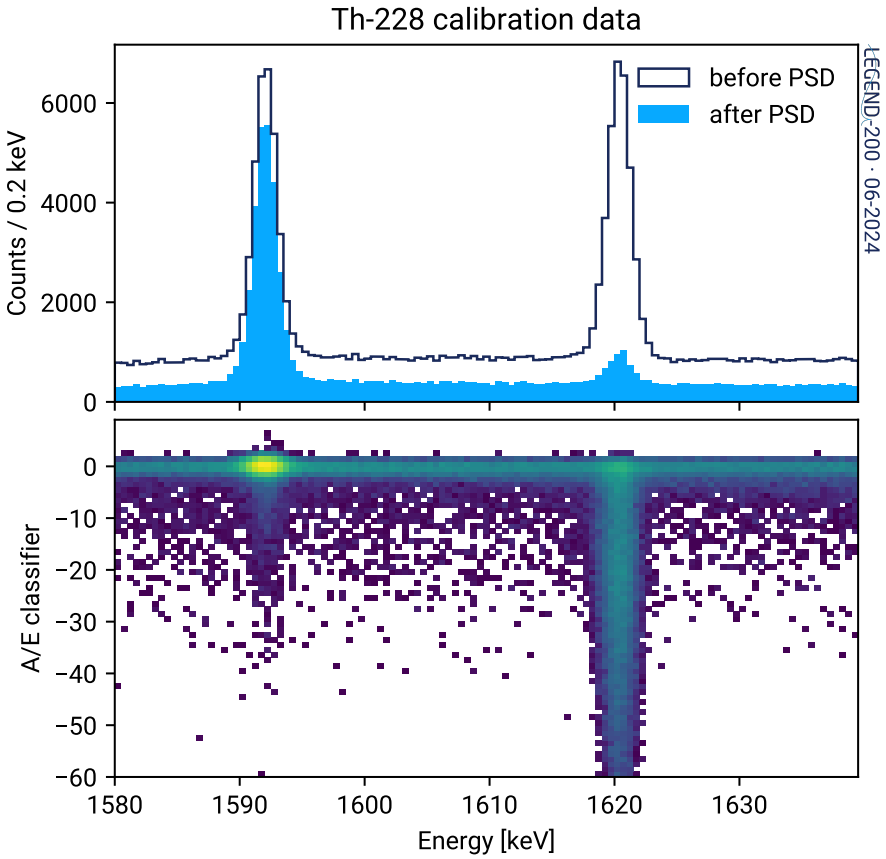

Figure 5.16: The effect of the A/E cut on the DEP and nearby Bi-214 gamma line. The DEP is preserved while the gamma line is suppressed. The top panel shows the energy spectrum while the bottom shows the A/E distribution. A typical cut value is between -1.5 and -2 on the low side.

Chapter 6

Routine Application to Characterisation Data

In this chapter I will discuss my work on the characterisation of the germanium detectors for LEGEND. Many of the routines and methods initially developed here evolved into the LEGEND data production pipeline. My work on the characterisation consisted of the development of the python/pygama based data processing software and the initial analysis of the data with this. The collaboration used this data to extract information such as the active volume of the detector. Previous to my work all data processing was done with an older C++ based software package. By switching to the python based software it brings us more in line with the other LEGEND processing along with improvements to the routines such as in energy resolution. These improvements were essential to assess the performance of the newest detectors which are the largest produced and require more sophisticated routines than used previously to correct for the longer charge drifts. Additionally, it allowed for the validation of these routines before they were used in the LEGEND data processing.

6.1 HADES setup

All new Mirion detectors for LEGEND-200 are first characterised at the HADES underground laboratory in Belgium (Ortec detectors are characterised separately in the US). This is necessary firstly to ensure the detectors meet the required specifications for LEGEND, secondly to understand their performance and thirdly to find issues which the manufacturer may need to fix. The HADES laboratory is located close to the manufacturing site of Mirion at Olen for precisely this reason. As it is underground, it can also serve as a place to store the detectors, so they are shielded from cosmic rays to reduce the production of cosmogenic isotopes such as Ge-68. Minimising the time on the surface for all detectors is crucial to keep their internal backgrounds as low as possible.

The choice of measurements has been developed through the experience of the GERDA and MAJORANA collaborations. In particular certain parameters such as the active volume are very difficult to extract in situ and therefore have to be done in the characterisation campaign. The detectors undergo a series of measurements with a variety of radioactive sources to determine they meet the required specifications for LEGEND in terms of depletion voltage, energy resolution and PSD performance and homogeneity. Additional measurements are taken to determine the dead layer thickness/ active volume of the detector. When the detectors are made an n+ electrode is added via Li diffusion in which no charge collection of deposited events occurs. This is done to reduce the number of surface events, however, it does mean we have some inactive mass. It is therefore essential that this is measured so that in the final statistical analysis we know the amount of exposure of the enriched isotope.

3 types of measurement are made at HADES: flood sources where the uncollimated source is simply placed on the side for lateral measurements and on the top for top measurements, collimated, so the energy depositions are localised and finally scan measurements which use the collimator and a scanning table to perform a series of measurements in steps around the detector. Below is a table which summarises the sources used, the type of measurement and the use of each:

Source	Measurement Type	Uses
Co-56	Flood	Depletion voltage determination
Th-228	Lateral and top flood	Energy and PSD characterisation
Ba-133	Top flood	Dead layer
Am-241	Top flood and surface scan	Dead layer and PSD

As can be seen in Figure 6.1, the data is taken in a vacuum cryostat cooled with liquid nitrogen which is surrounded by a lead castle to reduce the background from the environment. In total 39 new detectors have been characterised at HADES by the characterisation team along with some older detectors to study the dead layer growth as discussed in the previous chapter.

6.2 Processing

HADES data is taken with the same Flashcam DAQ as used in LEGEND-200. The waveforms are considerably shorter than those in LEGEND-200 at 60 μs long with a sampling of 16 ns. As mentioned above, these measurements are performed in a vacuum cryostat at liquid nitrogen temperature (77 K) so the noise conditions are quite different to those in LEGEND-200 where, the immersion of the crystals in LAr (87 K) can produce slightly different behaviour (characterisation is

Figure 6.1: Photo of the HADES setup with the vacuum cryostat in the centre surrounded by the lead castle.

not done in liquid Argon due to the complexity of decryostating the detectors and operating sources in LAr). However, overall the waveforms are very similar and this similarity to the LEGEND-200 data made it an extremely useful test bench for the development of the LEGEND processing routines. It was in particular used for studies on the best baseline subtraction method, t0 estimation and early quality cuts work. Work is now ongoing to incorporate the latest versions of the routines back into the HADES data processing software.

On the dataflow side, HADES data is considerable simpler to process in some ways than LEGEND data as only a single channel is present. The only complication is the large number of different measurements which require different calibrations and different routines. A specialised data processing version was developed for this data which also informed the development of the dataflow software before LEGEND data was available. In total around 1 Tb of data was taken per detector with processing taking a few hours per detector.

6.3 Energy Characterisation

One of the first checks performed on a detector is the energy resolution. In LEGEND we have an overall goal of an energy resolution at or better than 2.5 keV. Additionally, one of the main areas of development has been the size of the detectors with a lot of effort going into producing larger and larger detectors. As of writing the largest detector produced is more than 4 kg which is close to the limit of the manufacturers' methods, the limiting factor being the size of the crystal boule and the voltage needed to deplete the detector. These detectors are very beneficial to LEGEND as it reduces the channel count and therefore the background due to less supporting material and cabling. It is crucial, however, when producing these larger detectors, to avoid a degradation in the performance compared to smaller crystals. As a detectors size increases it is possible that the charge trapping will increase because of the longer drift times so the control of the impurities in the crystal is extremely important. From HADES data we can see that even at the highest masses there is no discernable difference in the energy resolution and, as shown in Figure 6.2, when plotting mass against energy resolution there is definitely no overall trend. This validates the production and gives us confidence in continued production of the larger detectors. The feedback loop from the characterisation to the production as the detectors grew in size was crucial to enabling this development.

The development of the analysis routines as described in the previous chapter could be benchmarked against previous analysis routines with this data. As shown in Figure 6.3 for all the filters included in the new routines an improvement was seen giving us confidence in their performance before application to LEGEND-

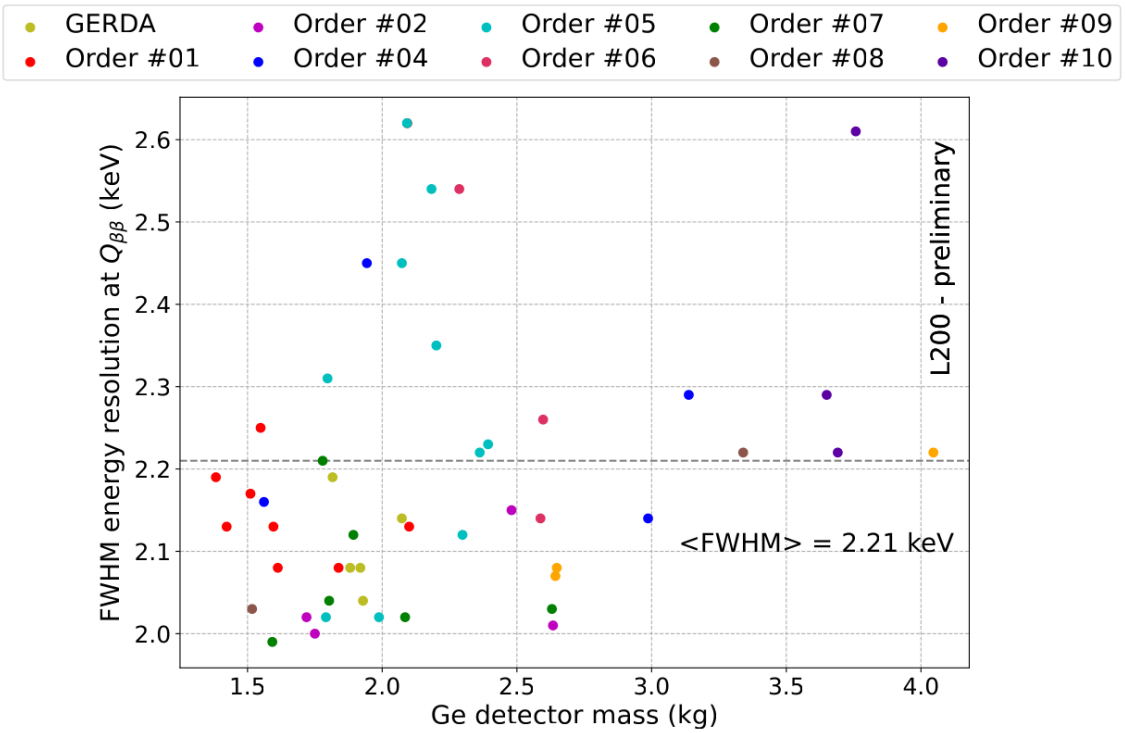

Figure 6.2: Energy resolution vs mass for detectors characterised in HADES. The different colours refer to the different batch numbers of the detectors. This shows we can make larger detectors without losing energy resolution.

200.

6.4 PSD Characterisation

One of the most important parts of the detector characterisation is the PSD. This is crucial as every detector is slightly different in geometry and charge collection characteristics and our routines are extremely sensitive to any inhomogeneity or geometry issues, Again, as above, the new larger detectors are of particular interest. The PSD is characterised using 3 measurements: a top and a lateral flood measurement with a Th-228 source, and a surface scan with an Am-241 source. I will focus here on the lateral flood measurement as the variation with height is the most important characteristic, as this is the axis where the charge collection paths differ and is also along the impurity gradient. To study the response, maps were made of drift time vs A/E for energy ranges of interest, namely the DEP, Bi FEP, Compton continuum at $Q_{\beta\beta}$, the SEP and Tl FEP. For a BEGe detector only a single ”blob” is seen as we only have a single drift time region (most clearly in the DEP) whereas for an ICPC 2 ”blobs” are expected due to drift time saturation effects and DEP events clustering in the corners as the gammas are more likely to escape [98]. The drift time saturation occurs due to the time it takes for a germanium signal to rise out of the noise especially for events at the top of the detector

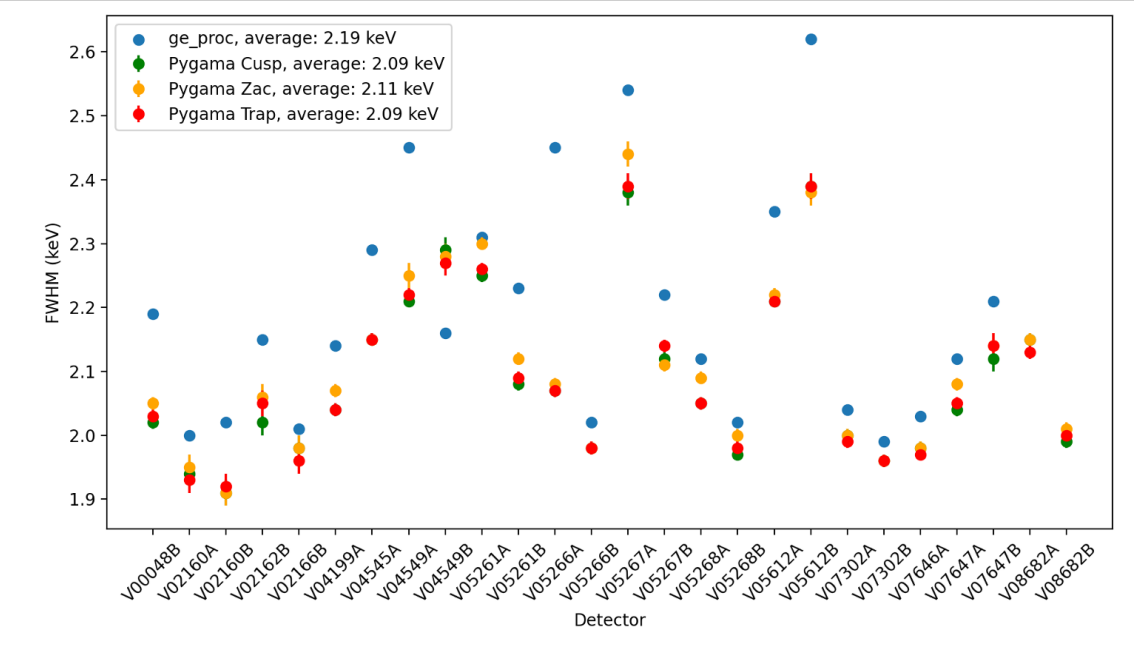

Figure 6.3: Energy resolution comparison for the old C++ processing (in blue) versus 3 different filters in the new python pygama processing. All show an improvement in the energy resolution over the old processing. For some detectors showed similar performance across all the filters and so the points may be overlapping.

their initial drift will be below the noise level, and therefore they will all appear to have similar drift times. An example of a 'good' detector with this structure is shown in Figure 6.4.

6.4.1 Batch 5 Detectors

One batch of detectors of interest in the PSD is the batch 5 detectors. These detectors were made with a slightly larger p+ electrode than previous detectors and in some cases a deep borehole was needed to achieve depletion (to maximise the n+ contact error and ensure all parts of the detector were close to it). However, these changes caused some issues in the charge collection. These issues are divided into 2 categories: named 3 blob and 4 blob detectors. In the 3 blob detectors the lower blob is extended out to higher A/E values (as shown in Figure 6.5), this is an issue as it hurts both the signal acceptance and the background rejection. The signal acceptance will be lower for these detectors if a high A/E cut is added as it will cut events in this tail. The background rejection will also be lower because background events in this region could have their A/E boosted into the signal band. Finally, this distribution causes an issue for the calibration routines as the peak is not the usual shape so either a drift time cut is needed to remove the tail before calibrating or a different peak shape definition. After this was found the collaboration went through a simulation campaign to try and understand the

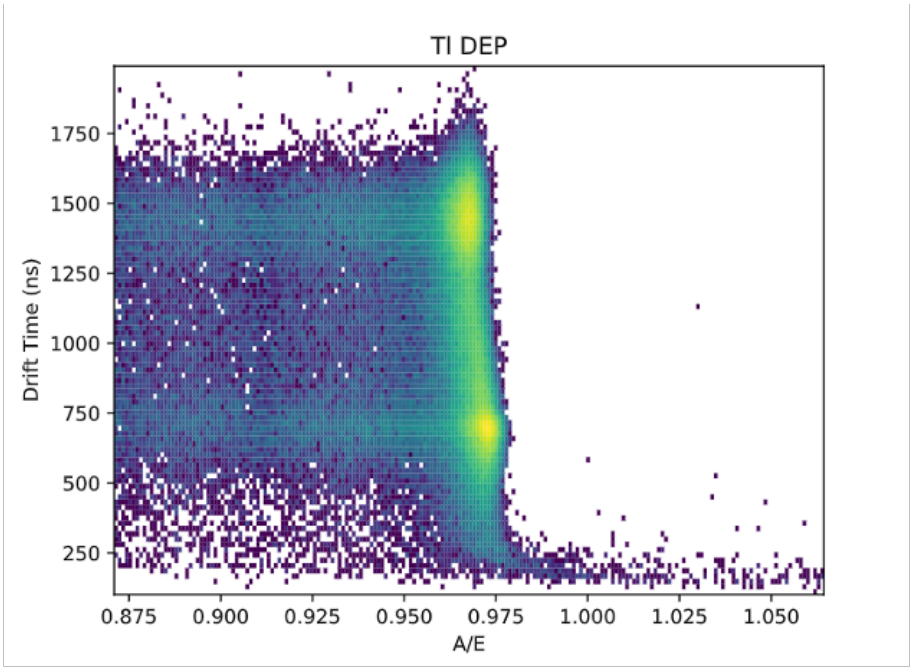

Figure 6.4: A/E vs drift time distribution for the DEP in a good ICPC detector for a Th-228 lateral flood measurement. Here DEP events are events within 10 keV of the DEP energy. The 2 blobs are clearly visible for the low and high drift time regions with the multi-site tail extending below. There is a slight shift between the upper and lower blobs due to collective effects which are more prevalent in the upper part due to the longer drift times giving more time for self repulsion and other effects.

issue which was traced back to the p+ electrode area. This was causing some charge paths in the lower parts of the detector to not follow the normal path and instead to ”hop” over the groove. These detectors are still usable for LEGEND analyses but with a lower efficiency than the other detectors.

The other class of detectors is the 4 blob detectors. In these detectors instead of just the lower blob being extended the entire detector spectrum is extended to higher A/E values. This means we have no way to calibrate these detectors using our standard methods as the peak is not well-defined. Currently, the collaboration has not managed to simulate why these detectors' behaviour is different, but it is suspected that as well as the p+ issue above the borehole depth is also a factor. Unfortunately these detectors are not usable for LEGEND $0\nu\beta\beta$ analyses due to the lack of calibration, although they can be used for background studies. This is an issue which could have been caught earlier in the process and possibly rectified but was only seen once I developed the production tools.

6.5 Conclusions

The characterisation of the detectors is crucial to the success of LEGEND. In this case it gave us a perfect springboard to develop, test and refine our analysis rou-

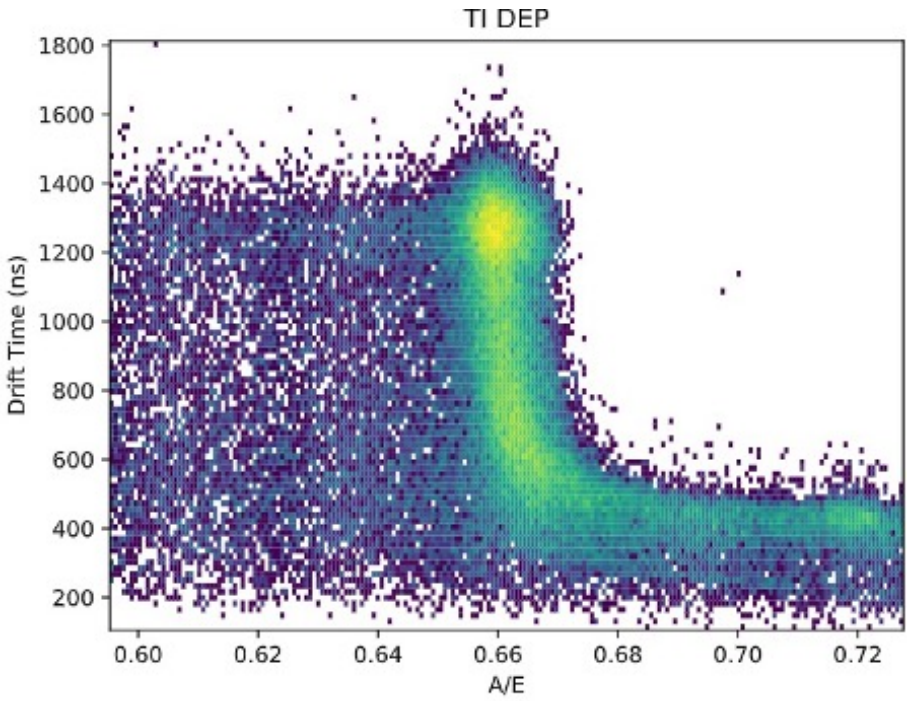

Figure 6.5: A/E vs drift time distribution for the DEP in a batch 5 detector. The extended tail for low drift time events can be seen. This hurts the background rejection for the low A/E cut for these detectors.

tines before the LEGEND data was available. It also informed the development of the dataflow software, metadata and production environment. The feedback loop from the characterisation to the production was crucial in the program to produce larger detectors which will be useful in reducing the background. Finally, the investigation of the batch 5 detectors showed the importance of the characterisation campaign in catching issues early, providing also insights on aspects of the geometry that weren't previously considered. In the end all these detectors were used in the first phase of LEGEND-200 although 2 have subsequently been removed for reworking (the "4 blob" detectors).

Chapter 7

Data Production Pipeline and Structure

In this chapter I will give an overview of the data production for the LEGEND-200 experiment. I led the development and implementation of this pipeline along with the supporting metadata packages. In some cases such as the dsp, hit, tcm and raw tiers the framework for running these was written by others, but I implemented the specific routines for LEGEND and integrated them into the pipeline. The general design of the data production is shown in Figure 7.1. This is managed in a production cycle whose inputs (software and metadata) is version controlled for preservation and replication. The software is run in a container, so it is portable with all processing being done on a dedicated server at LNGS with 200 CPUS and 200 Tb of storage. The production was validated by running on a subset of the data and a comparison to a secondary stack written in Julia is also in progress.

The challenge of the processing is managing the over 100 channels of germanium detectors all needing their own optimisation and calibrations. Additionally, there are 60 SiPM channels and 40 PMT channels, again all requiring their own processing. This necessitates a high level of automation with exceptionally reliable routines to minimise the number of failures. Typically, 1 Tb of data is taken per run and this is processed in as close to real time as possible to enable fast monitoring and feedback to the hardware team to maximise livetime and detector performance. Larger reprocessings are then run on analysis worthy datasets typically every few months, processing around 100 Tb of data in around 1 week. The whole analysis pipeline is able to run with no user interventions or issues validating the design and approach.

7.1 LEGEND Data

As described in previous chapters LEGEND consists of 3 different subsystems all of which have different types of detectors: germanium detectors (geds), the LAr

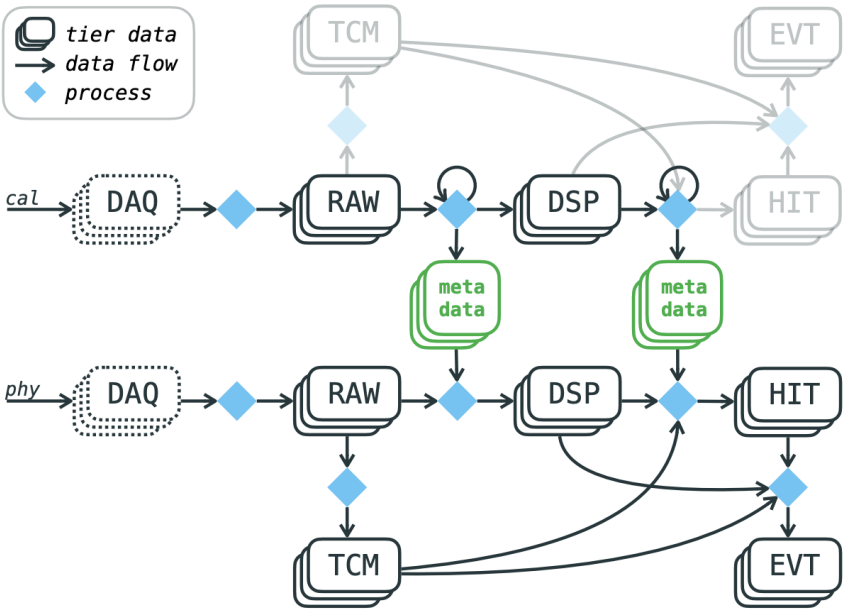

Figure 7.1: Diagram showing an overview of the data production steps. The data production follows a tiered structure reducing the over 100 Tb of raw data to a few Mb of final event data. The data is split into two streams calibration and physics data. Both run through the same pipeline steps but the calibrations are used to generate parameters/metadata which are then used in the processing of the physics data, and this is done on a channel by channel basis. DAQ is the data taken by the experiment, this is then converted to our lh5 format in the raw tier. The tcm (time coincidence map) tier groups the channels together in time, the DSP tier processes the data to extract parameters of interest, the hit tier calibrates these parameters and the evt tier puts all this together into our final events.

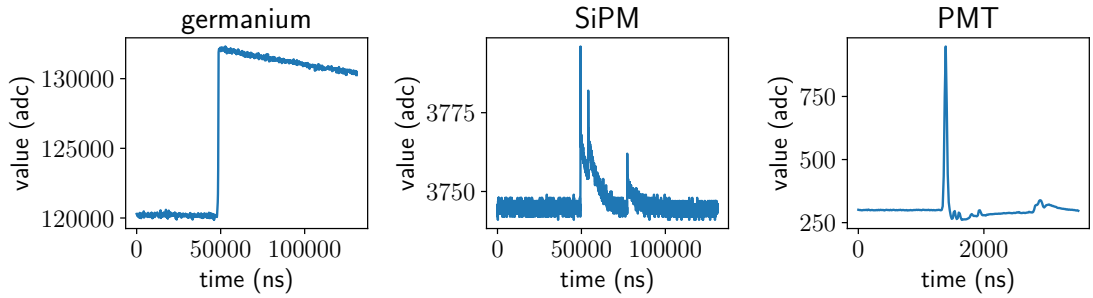

Figure 7.2: Example waveforms for the different subsytems in LEGEND-200. The germanium detectors have a slow signal with a typical rise time of around 1 us. By contrast, the SiPM signals are much faster with a rise time of around 100 ns. They also have multiple individual SiPMs per channel which are readout together, so often multiple signals are seen. Finally, PMT signals from the muon vetoare also very short with a fast spike signal.

system with its silicon photomultipliers (spms) and the muon veto system (muon) using photomultiplier tubes. These all have different signal structures and therefore different processing requirements as shown in Figure 7.2. The processing of the germanium detectors is the most complex and has already been discussed in a previous chapter.

For the LAr system, the SiPM signals are considerably faster (with rise times of around 100 ns) than the germanium signals and as each SiPM channel consists of multiple SiPMs we often see multiple signals. They are digitised in the same way as the germanium signals with the same length (128 us) and sampling rate (62.5 MHz). We have a mixture of SiPM decay shapes due to different manufacturers being used. Some have a single decay constant like the germanium channels while others have a fast and slow component. The two quantities of interest here are the energy and the time of the signal. The energy is the integral of the signal but due to the rate we use the maximum of the derivative of the signal instead which is also proportional to the energy due to the fixed shape of the SiPM signals. The timing information is very easy to extract with high resolution in comparison to the germanium detectors because of the fast nature of the signals.

The situation is similar for the muon system, but the integral can be used for the energy. The time resolution is of less importance here as a larger veto window is applied around muon events. These have their own DAQ system and therefore can have much shorter traces (4 us) with the same sampling (16 ns).

Finally, there is also a series of auxiliary (aux) channels, the most important of which are: the forced trigger channel, which is used to force a trigger to provide empty waveform data this is done every 20 seconds; the pulser channel, which pulses the front end of all detectors at an energy of roughly 1 MeV, and is used for monitoring the system this occurs every 20 seconds but out of phase with the forced triggers; and finally the muon veto channel, which provides a logic signal

for whether the muon veto system fired (this is no longer used in the analysis, as an offline reconstruction is preferred). All of these subsystems are stored in the same files and formats and are processed together. The requirements for these different processings drove the design of the data production pipeline with specialised routines for each subsystem.

A neutrinoless double beta decay event should consist of a single energy deposit in a single detector with no energy in any of the other subsystems. This drives the design of our triggering scheme which has 2 different modes: global and sparse. In global if any germanium channel is triggered then all channels for the germanium and LAr subsystems are read out. This is used for all physics data so for each event we can see all channels and check if events fulfil the requirement for a $0\nu\beta\beta$ event. We have recently also added the ability for the array to trigger also on the SiPM channels to widen the physics reach of this subsystem such looking for Bi-Po events(bismuth decays to polonium that then decays itself, due to the short half life of the polonium these decays are close together in time) and other LAr subsystem only events. The muon channels trigger independently of the other subsystems. We also have a sparse mode where only single channels are read out. This is used for calibration data to limit the data volume recorded as we do not need the full event information to calibrate a detector.

To store the data LEGEND uses a customised version of the hdf5 data format called lh5 [99]. This is a columnar format where all channels are stored in the same file but in separate tables. In global mode all of these tables are the same length whereas in sparse only the triggered channels are read out, so the tables are different lengths. The lengths of the tables will be the same across all our data production steps. It is only at the event tier where the format changes as all channels are combined to give the final event structure.

We also apply some 'trimming' to the data to reduce this size of the files. The germanium signals are split into two waveforms one full length but downsampled by a factor of 8, and one windowed on the leading edge. We also drop 2 bits on the SiPM to reduce their data volume. This combined with data compression routines reduces the data volume by around a factor of 5.

7.2 Data Taking

In LEGEND-200 we split our data taking into blocks called partitions defined as a continuous period in which the analysis parameters such as energy resolution are stable. At the hardware level we have a similar structure with the data being split into blocks called periods where our hardware is stable through the whole block. If any change is made for example to the voltage applied to detectors this triggers the start of a new period. A period can have multiple partitions if a detector is not

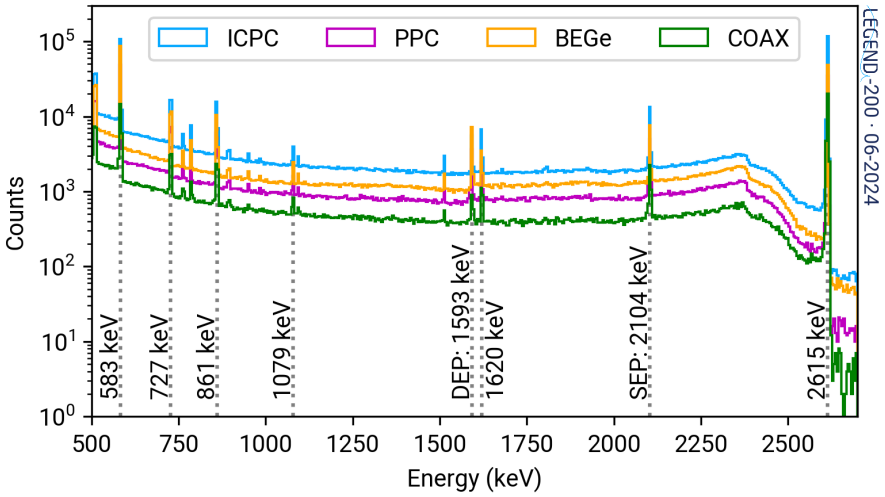

Figure 7.3: Example calibration spectrum from Th-228 showing major peaks split by detector type. The highest peak is the Tl-208 2.6 MeV gamma line with $Q_{\beta\beta}$ sitting just below the Tl-208 SEP at 2.1 MeV. The range of peaks allows for an interpolation of the calibration parameters to $Q_{\beta\beta}$.

stable for the whole period, but the change of a period nearly always necessitates a new partition. This is done to ensure that we can easily group this data together for calculating efficiencies and calibration parameters as described previously.

Furthermore, the physics data is split up into week long runs in order to monitor the stability needed for the $0\nu\beta\beta$ search. These are always bracketed at the beginning and end by a calibration. For this we use a Th-228 source due to the distribution of the gamma lines in its spectrum as shown in Figure 7.3. The highest full energy peak (FEP) line is at 2.6 MeV from Tl-208 (above the energy of $Q_{\beta\beta}$ allowing interpolation of calibration parameters), producing also 2 further lines of interest the Single Escape Peak (SEP) at 2.1 MeV and the Double Escape Peak (DEP) at 1.6 MeV for calibrating PSD. The weekly calibrations are used to generate parameters such as calibrations curves which are handled on a channel by channel basis before grouping into a single parameter file. This metadata can then simply be passed in at the appropriate tier for the physics data as shown in Figure 7.1.

Additionally, a number of plots are also generated for each channel at each tier of the production to monitor the data and provide easy checks for the analysts. For example plots showing the validity of the energy and PSD calibrations, the stability during the calibrations and the performance of the detectors are all created (see one below Figure 7.4).

7.2.1 Detector Status

One of the major inputs needed for the data production is the detector status. For this the weekly calibration data is combined with information from the physics data

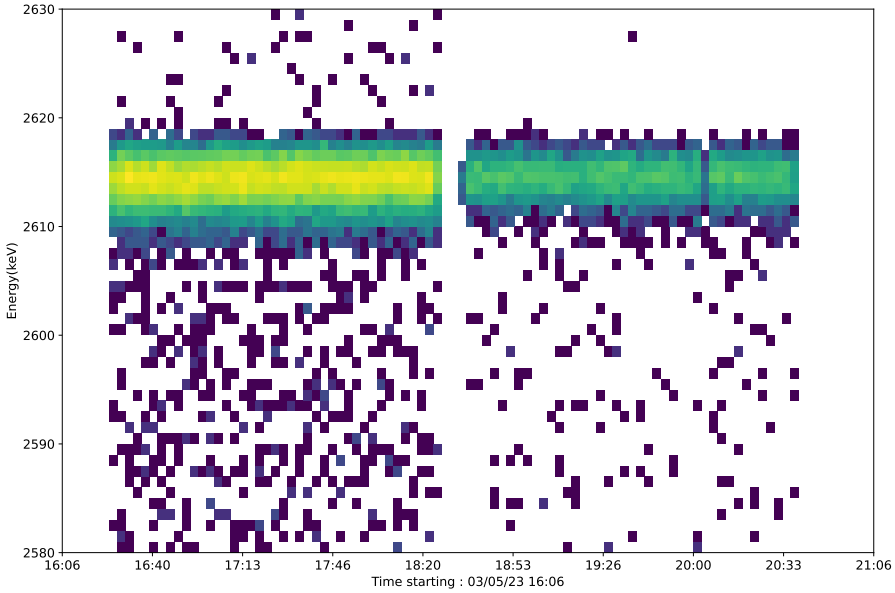

Figure 7.4: Example of a stability check in a single calibration run. Here events within 30 keV of the Tl-208 2.6 MeV gammma line are plotted in time to check for drifts or jumps.

such as the pulser data to check for gain drifts or PSD instability. For a detector to be included in the $0\nu\beta\beta$ analysis it must be stable in both energy and PSD. We require for the energy that the detector's gain drifts by no more than half its FWHM at $Q_{\beta\beta}$ over the course of a run, while for the PSD we require that the centroid of the A/E distribution drifts by no more than 40% of its width over the course of a partition and also that its width is stable within its error. However, even if a detector is not stable it could be still used for anti-coincidence purposes. The detectors are therefore split into three groups:

- off detectors - these detectors are either too unstable in energy, don't have a valid spectrum, or have other issues such as leakage current so can't be used.

- ac detectors - these detectors aren't stable enough in PSD or energy to be used in the main analysis but are stable enough to use as a veto, for example if a detector is not fully depleted it could still be usable as a veto

- on detectors - these detectors are stable in both energy and PSD and are used in the main analysis

An example of the detector statuses for the first run in the dataset is shown in Figure 7.5.

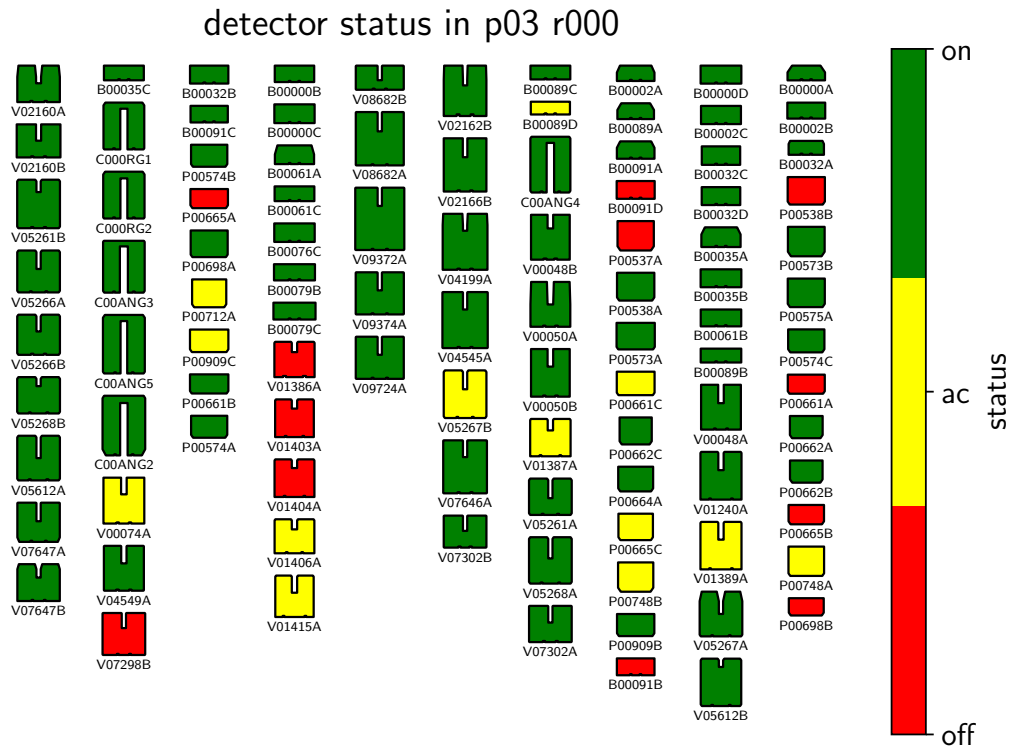

Figure 7.5: Example detector statuses for the first run in the dataset. 13 detectors were in ac, 12 detectors were off, leaving 76 detectors on

7.3 Data Pipeline

The data pipeline, depicted above, is common to all subsystems although the specifics of processing at each tier will differ. It was implemented in a custom package [100] using Snakemake, a python based Make [101]. This was chosen due its ability to run on a cluster, and its dependency handling so only steps after a change need to be re-run limiting the computational and time requirements for re processing data. All steps in the data processing output either a metadata file in JSON format or a data file in our lh5 format to allow Snakemake to track any changes.

The collaboration developed a python package: pygama to handle the routines specific to LEGEND [102]. For each tier, we use calibration data, which is taken weekly and described above, to generate parameters such as calibration curves. This is done separately for each channel and then grouped into a single file. This metadata is then passed to the data processing at the appropriate tier as shown in Figure 7.1.

At a high level we have 5 tiers of processing which in sequential order are: raw, tcm, dsp, hit and evt. The raw tier contains the waveforms without manipulations, the tcm the mapping of the channels in an event, the dsp tier the output of the signal processing, the hit tier the calibrated parameters, and the evt tier the final

event structure. Each successive tier represents a massive step in data reduction, with the data volume dropping by order of magnitudes. This allows for effective reprocessing of high-level tiers without the need to rerun the IO intensive raw-file generation and CPU/IO intensive DSP tier. A more detailed description of each tier is given below.

- Raw: The analysis group receives files from the DAQ in a format optimised for I/O performance compared to that used in the analysis which is optimised for processing, storage and usability. Therefore, the first step is simply to take these binary files and convert into the legend data format outlined above using the daq2lh5 package [103]. This step is performed continuously during data taking and only needs to be done once in a dedicated production cycle. We also split the waveform into 2 versions, a fully sampled waveform windowed on the leading edge and a downsampled version of the full waveform. This drops our data volume by around a factor of 5 when applied with data compression. Additionally, the blinding is applied for the physics data at this step. From here the higher level processing steps can be run as many times as needed.

- TCM: The second step is the time coincident map (tcm) tier, this uses the timestamp information in the files to group together the channels present in all subsystems in some time window using just the raw tier. This time window can be adjusted to suit the needs of the analysis. As we are in a global trigger mode for physics data, all channels are readout for each event and have the same timestamp so no time windowing is needed. However, for the calibration data we are in sparse mode, and so we need to group the channels if we want to look at multiplicity 2 and greater events or to help with pulser identification.

- DSP: The third step is the digital signal processing (DSP). It is handled by the dspeed package [104] and as can be seen in Figure 7.1 only relies on the raw tier. In this, we run a series of discrete processors over the data to extract quantities of interest. These are divided into 3 main areas: energy, data quality and shape/timing parameters which have been discussed in previous chapters. The processing is done on a channel by channel basis, so different processors can be run for different channels with many of the filters' parameters optimised for each channel. This means that the best possible energy resolution/PSD performance can be obtained for each channel. The optimisation steps are run prior to the DSP processing and use the weekly calibration data.

- Hit: At the fourth step, the hit tier, the quantities obtained in the dsp tier are calibrated to give the final variables. This is where the checks on the physical

nature of the waveform are calculated, the energy is calibrated and the PSD acceptance is decided. Similarly to the dsp, this is handled in blocks where each block is a simple mathematical operation such as applying a calibration curve or deciding if a hit passes a cut.

- Evt: Finally, the event tier is where individual hits in all the subsystems are combined into the final events using the tcm. The only information that changes between the hit and event tier is the energy estimate as we need the full event information in order to apply a crosstalk correction to the energy. This will not affect multiplicity one events but will affect higher multiplicity events. It is also crucial in deciding the multiplicity of the event. This is also where the muon veto and LAr cuts are calculated by combining all the detectors in these subsystems.

7.4 Blinding

All physics data in LEGEND is blinded for the $0\nu\beta\beta$ analysis. This is done by using the online DAQ energy estimator to remove all events within plus/minus 25 keV of the $Q_{\beta\beta}$ value. This DAQ energy estimator is calibrated at the start of each data taking period with the validity of these curves checked manually. The calibration curves are then stored in the metadata. For every calibration the validity of the blinding is checked to make sure no channel has shifted more than 5 keV from the blinding curve. If any channel has moved the data production is stopped until a manual check is performed.

For the physics data, if any channel has a hit in this window ($Q_{\beta\beta}\pm$25 keV) the event is removed for all channels to preserve the data structure. The non blinded data is stored on a separate area of the disk and is only accessible by a few analysts. This blinding is done to remove biases in the analysis that may come from, for example, tuning cuts to get the lowest number of surviving events in window. A larger window around the $Q_{\beta\beta}$ value is used as the fitting window for the analysis and can be used to estimate the background in the region. All analysis methods and procedures are reviewed and frozen before the data is unblinded which is done by running the same analysis chain but on the non-blinded files.

The spectrum after blinding can be seen in Figure 7.6. A small amount of leakage is expected at the edge of the blinding window as a different energy estimator is used for the blinding to the offline analysis. This is not a concern as the window is large enough in comparison to the difference between the filters that no event close to $Q_{\beta\beta}$ should pass unintentionally.

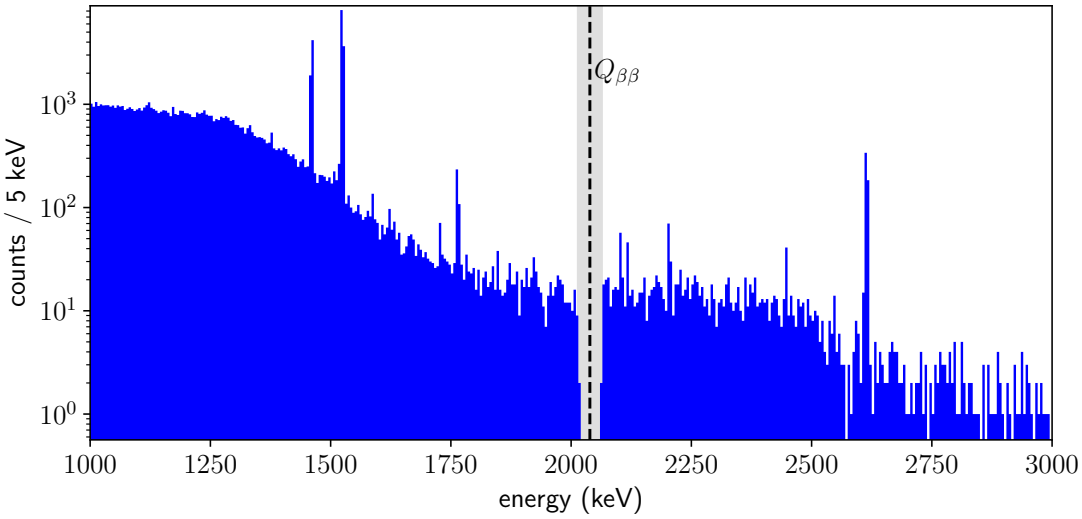

Figure 7.6: L200 physics spectrum with blinding applied. The effect can be clearly seen with all events in a 50 keV window around $Q_{\beta\beta}$ removed. A small amount of leakage is expected at the edge of the window due to the different energy filter used in the offline analysis to that used for the blinding.

Chapter 8

Dataset and Detector Performance

In this chapter I will introduce the dataset used for the first $0\nu\beta\beta$ analysis of LEGEND-200, how much exposure was accumulated and how the dataset was split up. I will then discuss the performance of the germanium detector array in this dataset in terms of the energy and PSD, along with the LAr subsystem performance.

In the first section I will discuss the work of the data quality team in deciding the set of usable detectors for the LEGEND-200 analysis. I contributed to this work by developing routines to check the performance of the detectors during calibrations as well as the baseline version of the detector usability from the calibrations before these were further refined by the team. The second section will discuss the work on the performance of the detectors, with my work looking at their stability in energy and PSD performance in calibrations, and the energy scale team also looking at physics performance as a cross-check. Finally, for completeness, I will present the collaboration's work on the performance of the LAr subsystem.

8.1 Neutrino 2024 Dataset

The LEGEND-200 array was deployed with 101 detectors corresponding to 142 kg in November 2022. It underwent a commissioning phase for 3 months and then started physics data taking in March 2023. It continued physics data-taking until February 2024, as shown in Figure 8.1. The exposure gain was fairly constant in time with a high duty cycle in stable periods, the two large gaps seen were periods of hardware operations where the array was not in a sufficiently stable state for physics data-taking. At the end of February 2024, it was decided to undergo campaign to characterise the background to gain experience for LEGEND-1000. In the next month the outer barrel was removed and data was taken. Finally, the mini shrouds and some calibration guide tubes were removed and another month of data was taken.

In total 242 days of physics data were taken. This data is split into 6 periods (numbered p03-p09 with p05 removed) with a change of period defined by

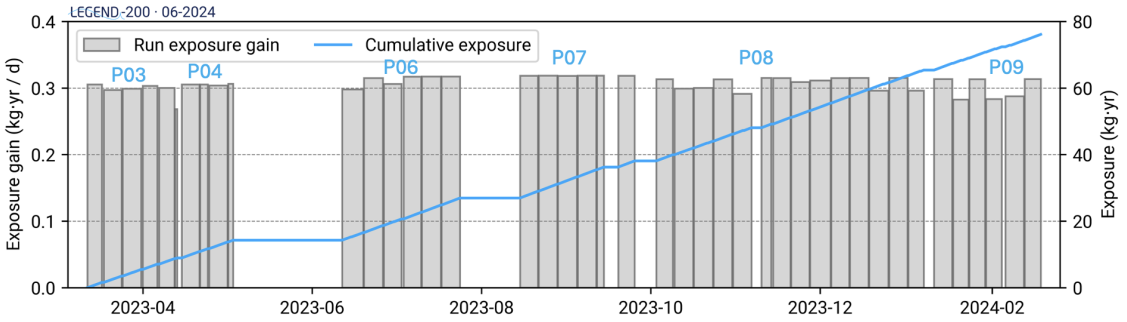

Figure 8.1: Plot of the data taking of LEGEND-200, showing the exposure accumulated in each run in the bars and the cumulative exposure of the dataset in time. 76.3 kg yr was accumulated in total.

a change in the hardware status most often due to an adjustment to the detector biasing voltages to optimise performance (for example a couple of detectors took much longer to ramp up to operational voltages due to leakage current issues that stabilised over the course of the data taking). The bias settings of the transistors in the front-end electronics were also retuned at the start of p08 to reduce harmonic oscillations (ringing) affecting the read-out signals of some PPC/Ortec detectors. Of the 142 kg deployed 12 kg was not usable due to either becoming disconnected when immersed in LAr or not achieving depletion because of leakage current issues.

For a detector to be usable for the $0\nu\beta\beta$ analysis it needs to have both a stable energy scale and a stable PSD at below the per-mille level. Once a week 3 sets of Th-228 calibration sources were lowered into the array and data taken at 2 different positions for 2-3 hours each. This procedure was chosen to ensure at least 1000 events were accumulated in the 2.6 MeV gamma-line. The energy scale is tracked using these calibrations by looking at the change in gain with the Tl-208 gamma line at 2.6 MeV and using the pulser injected into the front-end electronics. However, it was found that the pulser itself fluctuated more than the germanium detectors due to temperature dependencies in the fanout cards (this was reproduced using a heat gun on the cards) so only large movements could be tracked. Finally, the energy resolution is checked to make sure it is stable. A run is considered usable for a detector if it is stable (any fluctuations in gain less than half the FWHM at $Q_{\beta\beta}$ for that detector) between the calibrations at the start and end.

A major effect was seen when the temperature of the cleanroom changed because of failures in cooling, for example due to power outages. Time periods where the temperature was elevated were removed using the slow control. The PSD is tracked using just the calibrations and this was done using a sample of events from the low energy (1-1.3 MeV) Compton continuum of the Tl-208 gamma line at 2.6 MeV which consists about 50% of single Compton scattered events. We

fit these events in A/E to extract the centroid and width of the peak and track these over time to ensure the centroid is not moving by more than 25% of the width and the width is stable at the percent level. Some detectors were not usable from a PSD standpoint due to uncertainty in their performance, for example the ringing detectors mentioned above, or due to the routines not yet being developed, for example the COAX detectors need a specialised machine learning routine due to the different charge collection in these detectors. These can hopefully be used in the future when more studies are completed.

In the end the dataset was divided into 2 categories: firstly a background dataset which consists of all detectors with a stable energy scale for use in background studies and secondly a $0\nu\beta\beta$ dataset which consists of only detectors with full PSD. The background dataset has 76.2 kg yr of exposure and the $0\nu\beta\beta$ dataset 48.3 kg yr. Here we use the full background dataset for discussions until the PSD cut which can only be applied for the $0\nu\beta\beta$ dataset. The final exposures for each detector can be seen in Fig. 8.2 for the background dataset on the left and the $0\nu\beta\beta$ dataset on the right. The largest contribution came from the big detectors on string 5 in both cases. We hope to recover much of the exposure difference between the two plots with further studies on the PSD.

8.2 Detector Performance

First before looking at the physics data the results from the calibration runs will be discussed. In particular, we are interested in 3 main quantities that make up the energy systematics: the energy bias which is the difference between the reconstructed energy and the true energy, the energy resolution and the calibration stability which is how much the gain is changing over a run. For the energy bias the validity of the calibrations was first checked by looking at the peak residuals of 3 major gamma lines: the line at 583 keV, the Tl-208 SEP at 2.1 MeV and the Tl-208 2.6 MeV line for each individual calibration run, shown in Fig. 8.3.

The stability of the calibrations was looked at as shown in Figure 8.4. The stability is measured by taking the fitted centroid of the 3 gamma lines all from Tl-208: the 583 keV line the SEP at 2.1 MeV and the 2.6 MeV gamma line, for subsequent calibrations run and looking at the shift between these in raw energy values. This can then be converted to keV by applying the calibration constants from the first. As you can see from Figure 8.4 the array is very stable at well below the keV level. As the resolution of the detectors can be thought of as the combination of the intrinsic resolution and the additional smearing from any instability, the effect from this is subdominant to the energy resolution of the detector themselves which is above the 2 keV level.

Moving on to the energy resolution, which was also checked for each detector.

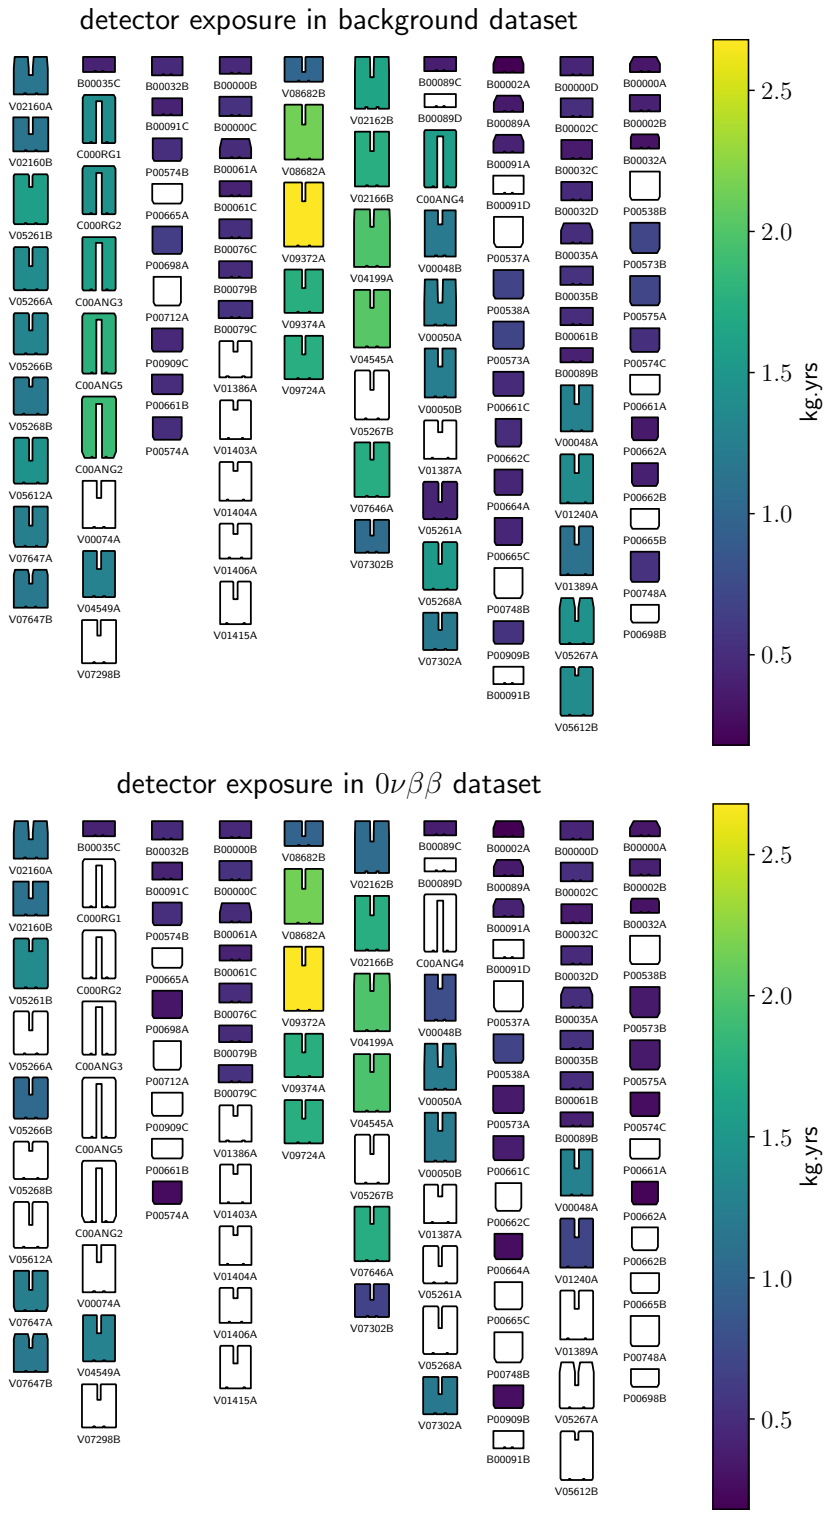

Figure 8.2: Exposures for each detector for the background dataset at the top and the $0\nu\beta\beta$ dataset on the bottom. Some types of detector could not be included in the $0\nu\beta\beta$ dataset due to their PSD still being under development, they will hopefully be included in the future.

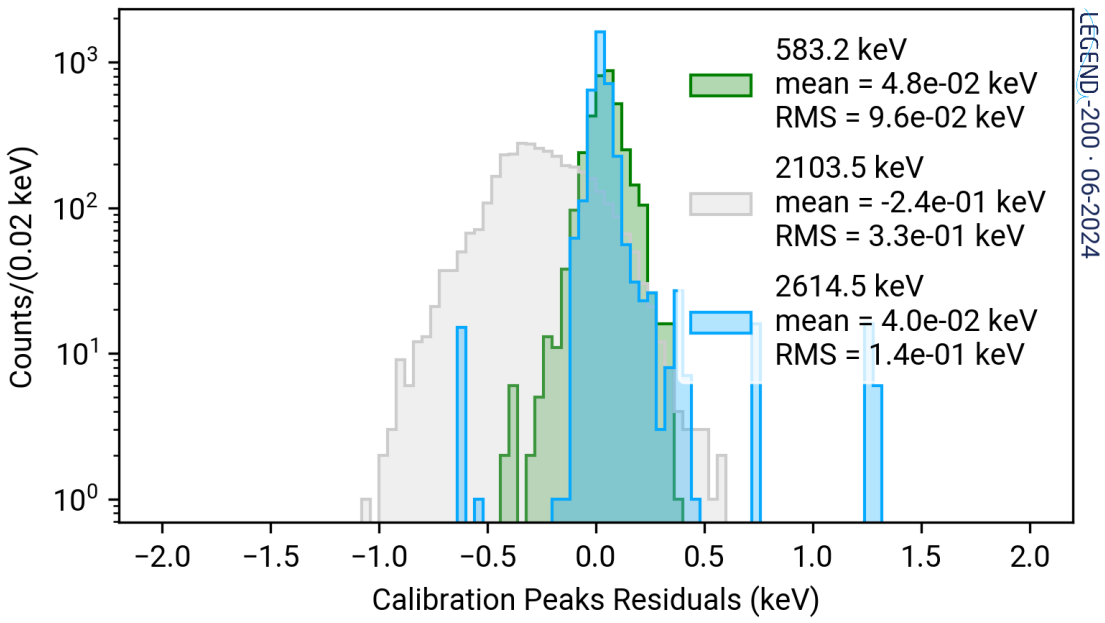

Figure 8.3: Residuals of the final calibration for 3 major peaks in the calibration data. This is the difference between the reconstructed position and the true value from literature, done for the Tl-208 gamma lines at 583 keV and 2.6 MeV and the SEP at 2.1 MeV. This is used to check the accuracy of the energy calibration. There is a slight bias in the SEP indicating we are under correcting the non linearities in the system, although this is well below the level of our energy resolution which is around 2 keV. This will be studied more to improve in the future.

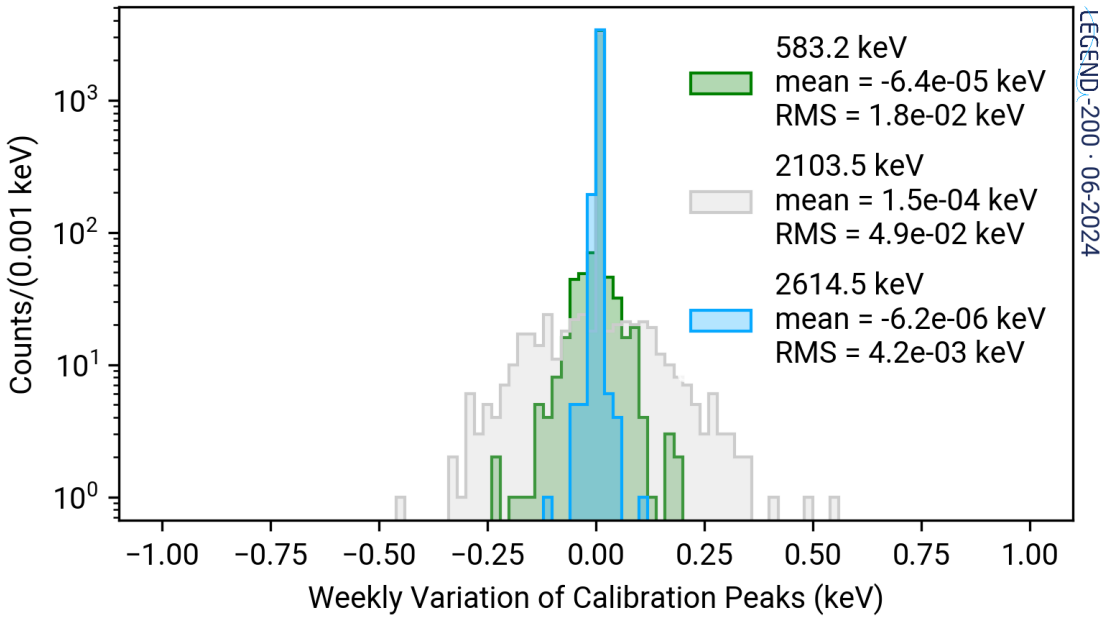

Figure 8.4: Stability of the array between calibrations. Measured by looking at the shift of three peaks between subsequent calibrations. The array is extremely stable at well below the keV level which means this effect is subdominant to the energy resolution of the detectors.

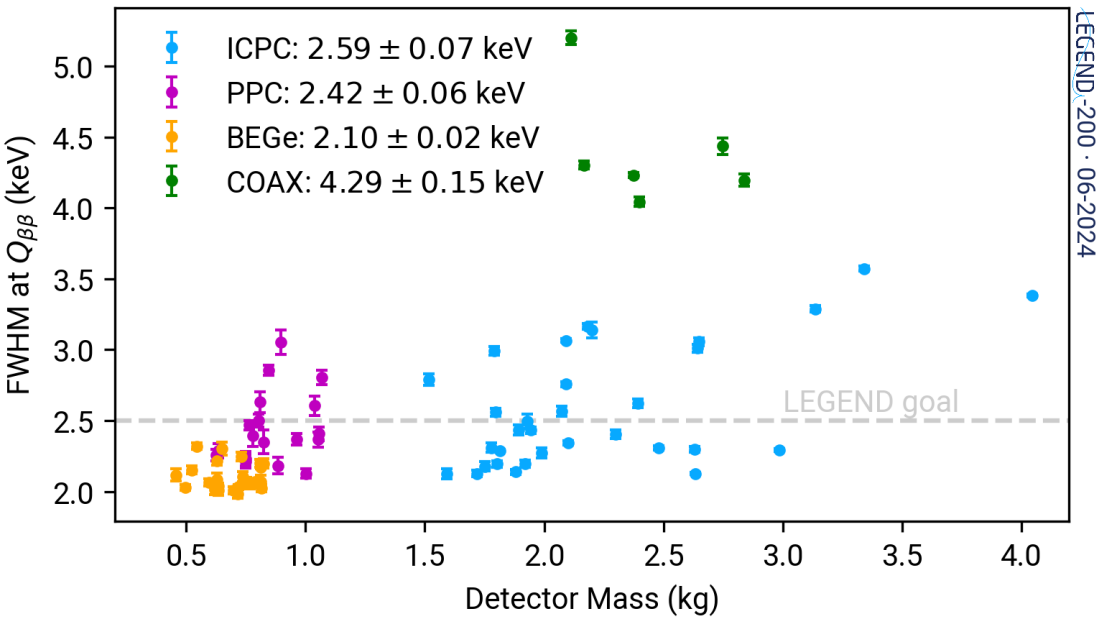

Figure 8.5: Energy resolution of the detectors against their mass. A line with the goal of 2.5 keV is shown. Two thirds of detectors are below this level, with the average compatible with this goal. The COAX detectors are around 1 keV higher, as expected, due to their higher noise. The slightly degraded performance of the largest detectors when compared to characterisation in Figure 6.2 indicates this can be improved further with optimised electronics.

The resolution is defined as the full width at half maximum of the gamma lines interpolated to $Q_{\beta\beta}$, this is shown against the mass of the detector in Figure 8.5. The BEGe detectors have the best resolution at a weighted average by mass of 2.1 keV, while the COAX detectors have the worst resolution at 4.29 keV (as seen in GERDA [105]). This is due to the high noise in these detectors which are all grouped together in the array. The performance of the ICPC detectors was in line with expectations at 2.6 keV.

Overall, LEGEND-200 achieved its goal of 2.5 keV for the average energy resolution which is the leading resolution in the field. This directly translates to an improvement in the sensitivity of the experiment as the energy region of interest shrinks, increasing the signal- to-background. The BEGe detectors saw a marked improvement of 0.5 keV in energy resolution over GERDA [105] due to the improved electronics and energy optimisation routines. The largest detectors still have a good energy resolution validating that we can produce these large detectors with good performance. In fact in commissioning deployments the resolution was even better due to lower noise, so an improvement in their resolution could be achieved with optimised electronics which will be pursued in the next data-taking phase.

The residuals and energy resolution were also checked for the physics data using the two K gamma lines at 1460 (from K-40) and 1525 keV (from K-42) to

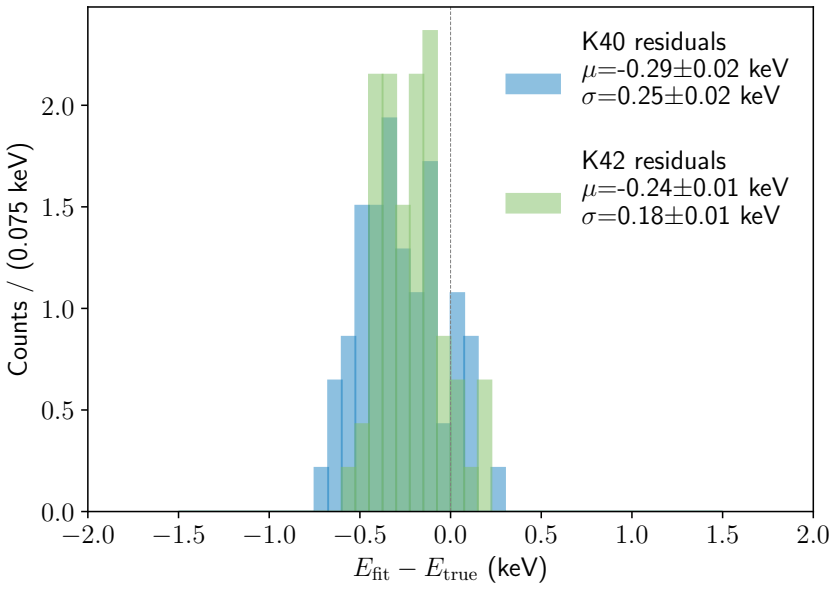

Figure 8.6: Residuals of the two potassium peaks for each partition for each detector. This was found by fitting the peaks in physics data with a gaussian for each detector and comparing the fitted centroid to the true value. The good agreement validates that our calibration is accurate.

ensure the validity of the calibrations. These were fit and the position and FWHM extracted to compare to the expected values from the calibrations. The bias is shown in Figure 8.6 and the comparison of expected to measured resolution in Figure 8.7. The good agreement in both cases between the measured value and the expected value from our calibration confirms that the energy calibration was accurate and applies well to the physics data. The collaboration found the bias at $Q_{\beta\beta}$ using an interpolation between the Tl-208 DEP at just below 1.6 MeV and the Tl-208 SEP at 2.1 MeV as these are the two closest gamma lines with sufficient statistics for interpolation.

Moving on to the PSD performance from calibrations. For all the high statistic gamma lines above 1 MeV survival fractions were calculated by fitting the peaks in energy for events passing and failing the PSD cut. This was done for all detectors with valid PSD, In addition, the survival fraction of the Compton Continuum after PSD application in the region of $Q_{\beta\beta}$ was also looked at and this is shown against the mass of the detector in Figure 8.8. All the detectors are in the expected region of <50% with the bigger ICPCs having a higher rejection due to the higher probability of multiple Compton scatters.

The collaboration calculated the efficiency of the PSD for each detector in each partition. This was done by taking the survival fraction of the Tl-208 DEP and adjusting for the difference between the DEP and $0\nu\beta\beta$ efficiency. This is related to both the different topologies with DEP events concentrated at the corners and edge of the detector while $0\nu\beta\beta$ events are homogenous and the energy dependence with the DEP at 1592 keV and $Q_{\beta\beta}$ at 2039 keV. For the energy depen-

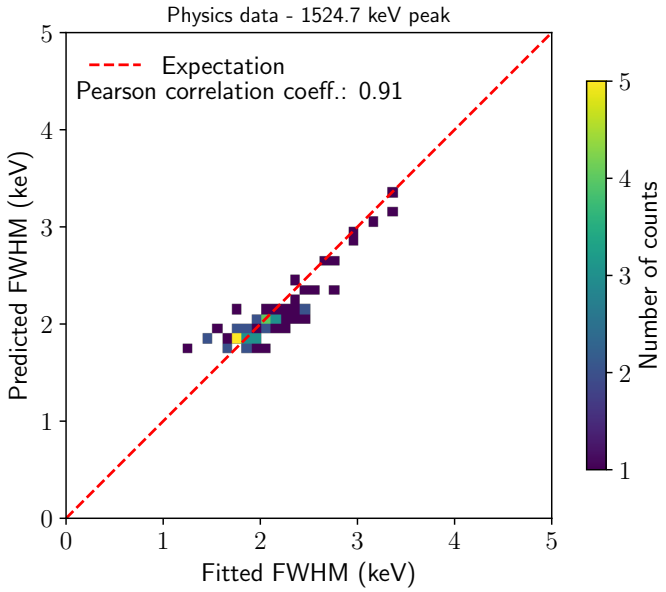

Figure 8.7: Comparison of measured to expected energy resolution for each partition for each detector. The measured energy resolution was extracted from a gaussian fit for each detector while the expected resolution is found from an interpolation of the resolution curves from calibrations. The correlation coefficient is also shown. The good agreement validates our calibrations.

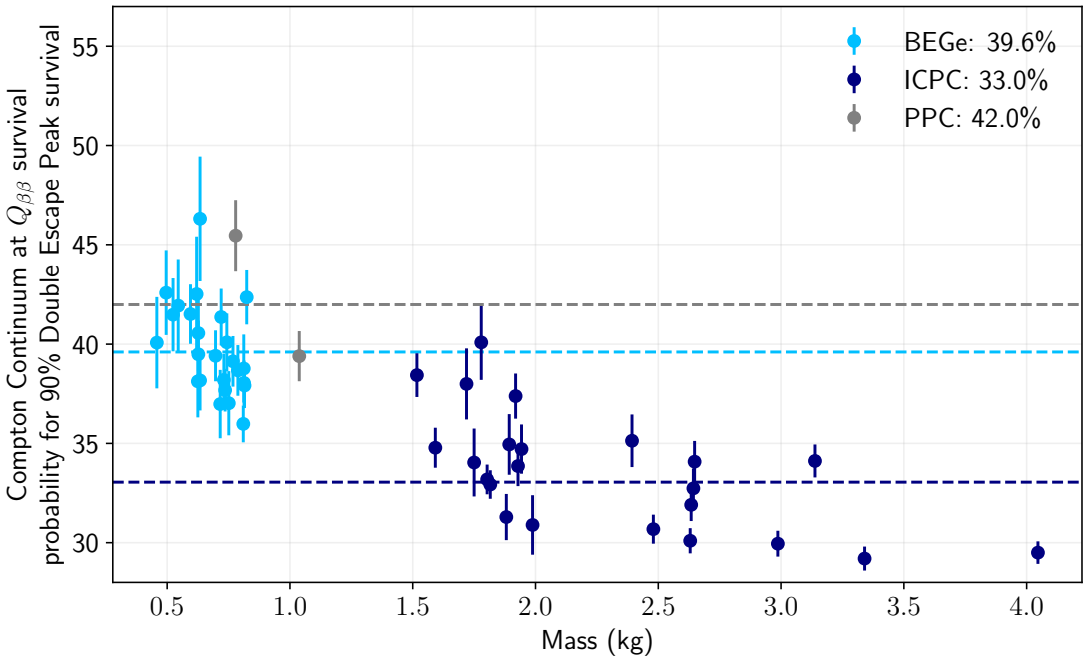

Figure 8.8: Plot of efficiency at $Q_{\beta\beta}$ vs detector mass for a single partition for each detector. Large mass detectors have a higher Compton rejection due to the higher probability of multiple scatters.

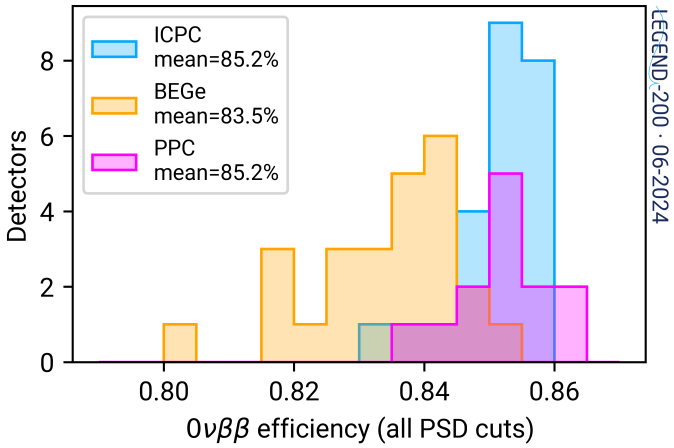

Figure 8.9: Plot of the efficiency of the PSD grouped by the different detector types. The COAX detectors do not feature in this plot as their PSD is not yet validated.

dence the best estimate is to use Co-68 data which has a range of DEP gamma lines both above and below $Q_{\beta\beta}$, while the topology difference is estimated using the pulse shape simulations. As we do not yet have fully processed Co-68 with LEGEND-200 test stand data was used instead. The final efficiency was calculated by multiplying the A/E low side efficiency with the chosen surface cuts efficiency after this adjustment. The final efficiencies are given in Figure 8.9 with the colours corresponding to the different detector types. As stated above the COAX detectors are not yet validated for PSD so are not included in this plot.

8.2.1 LAr Subsystem Performance

Away from the germanium detectors the performance of the LAr subsystem was also looked at. I include this work from the LAr group to give a complete overview of the performance of the detector with my contribution solely being on the data processing management tools and routines.

The LAr subsystem is upgraded considerably from GERDA as can clearly be seen in Figure 8.10 comparing the Ar-39 light yield in physics data between the two. In addition to using physics data, the performance of the LAr subsystem was checked using special calibration data. With the standard calibration sources the rate was too high for the SiPM arrays therefore low activity Th-228 sources were used as well as Cs-137 sources. The Cs-137 run is shown in Figure 8.11 and shows that we see 1 p.e. per 8 keV deposited in the LAr and that the system is highly linear with energy. The Th-228 low activity run also shows the improvement in light yield compared to GERDA Figure 8.12. Here, a comparison of the light yield of the Tl-208 DEP between GERDA and LEGEND-200 is shown. This gamma line occurs when both photons from the annihilation of the positron created in pair production escape the germanium detector. As described above this

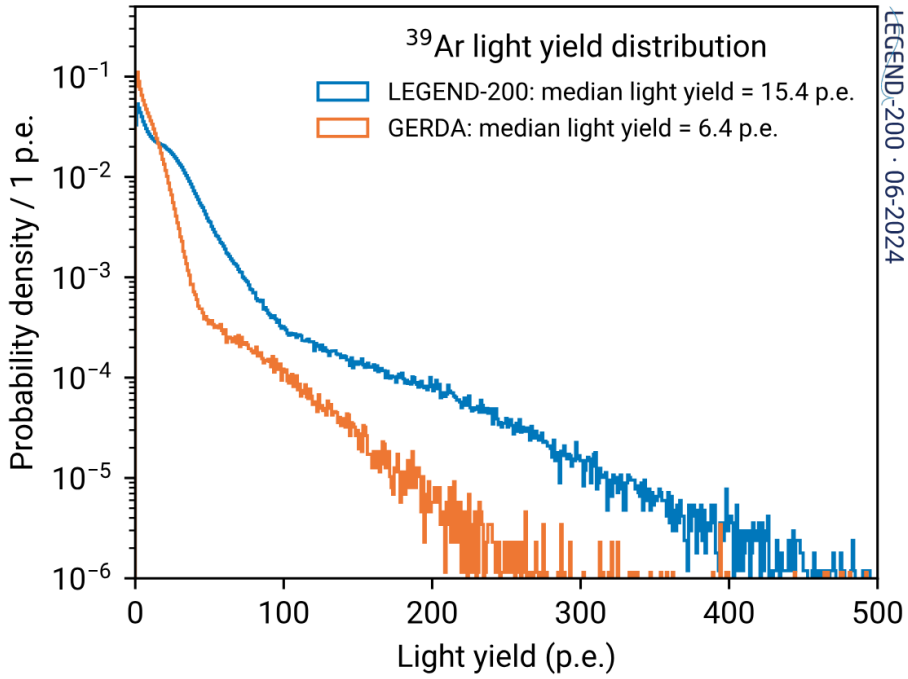

Figure 8.10: Comparison of the Ar-39 light yield between GERDA and LEGEND-200. LEGEND-200 has a higher median light yield than GERDA. The shape of the spectrum is determined by the Ar-39 beta spectrum and the geometry of the LAr subsystem. GERDA had only an outer barrel in contrast to LEGEND-200 which has optical fibres on either side of the Ge detectors. GERDA's detectors were also clustered more tightly together increasing the shadowing effects.

gamma line is preserved by the PSD cut, so the high light yield here illustrates the complementarity of the analysis cuts where this peak is preserved with the PSD cut but would be suppressed with the LAr anticoincidence cut.

One issue found was the presence of crosstalk from the germanium detectors to the SiPMs. This was particularly hard to tag and remove due to the variance in the signal shapes. It might that a simpler veto condition had to be used for the LAr anticoincidence cut than would have been used ideally. This is an area of active work to understand and improve for the future both at the hardware and software level.

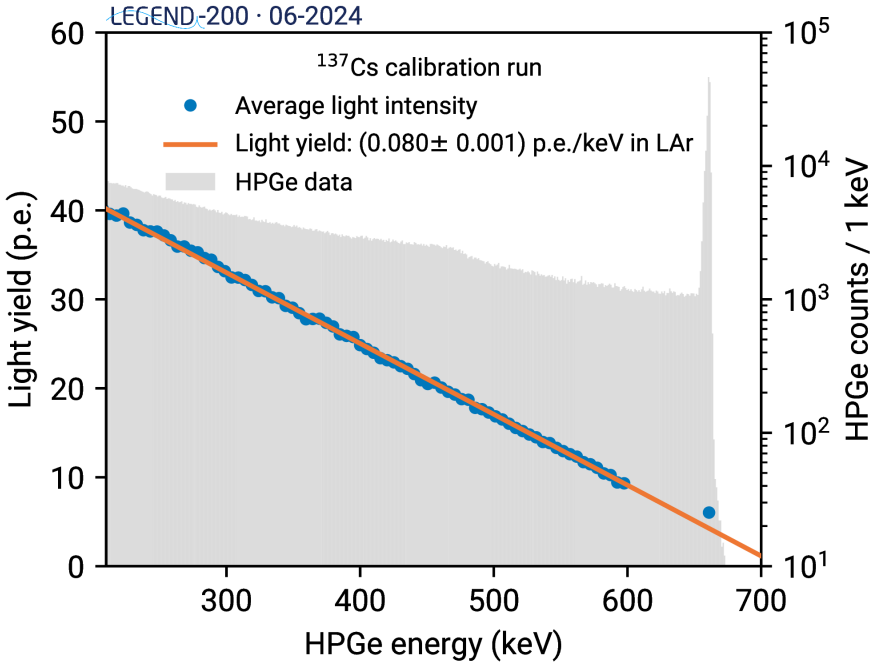

Figure 8.11: Light yield in the special calibration run with Cs-137 showing we see 1 p.e. per 8 keV deposited in the LAr and validating the linearity of the system. In grey is the light deposited in the Ge detectors for each event. The energy in the LAr will decrease as more energy is deposited in the Ge detectors.

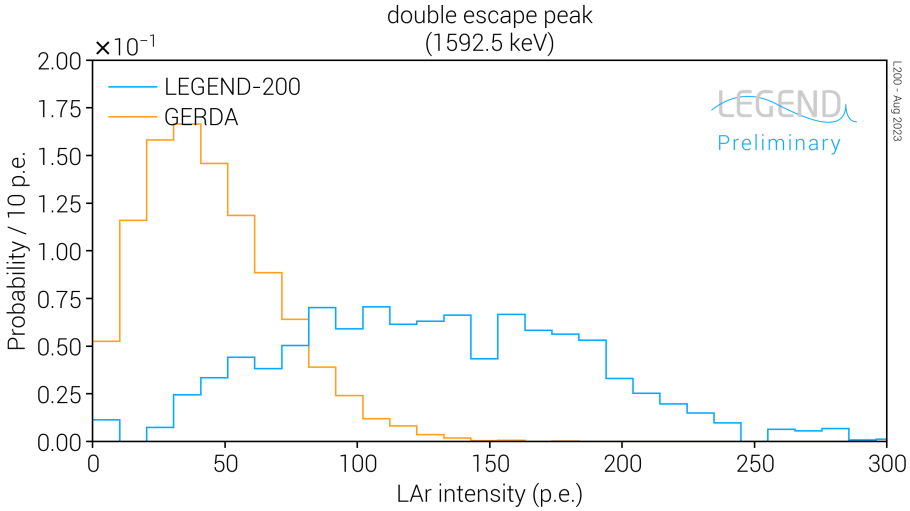

Figure 8.12: Light yield for the DEP in the special calibration run with a low activity Th-228 source. A comparison is shown to GERDA showing the higher light yield in LEGEND-200

Chapter 9

Physics Data Characterisation and Background Analysis

This chapter will discuss the work on characterising the background in LEGEND-200. An elevated background compared to expectations was found before analysis cuts were applied. My work focused on using the data to look at the homogeneity of this background which was studied by detector, string, detector type and in time to investigate the source. This work and the work on the production were built on by the background working group for more in depth studies such as the analysis of the counts in the major gamma lines and building a background model to identify the possible sources of the elevated rate which I will describe here for completeness. I will then show the effect of each of the analysis cuts on the data. Finally, the background in the region of interest (BI) before unblinding will be shown. The background rate in the ROI after all cuts was found to be still compatible with the goal of 0.2 counts / (keV ton yr).

9.1 Quality Cuts and Muon Veto

The first set of cuts applied to the data are the quality cuts discussed in earlier sections. These are used to remove non-physical events. In total 50.2% of the data was removed by these cuts with the majority removed by the discharge cut (95.7% of cut events) which simply flags any event which exceeds the DAQ range. The large number of discharges was likely due to the larger detectors being too close to the copper supports and due to the electronics. The effect of these are shown in Figure 9.1. In the top panel is the full spectrum before and after quality cuts, showing that these unphysical events can be reconstructed with energies throughout the spectrum. In the bottom panel is the effect on low energy where the effect is particularly apparent as we move from an irregular spectrum with spike features to a smooth spectrum dominated by Ar-39 (one of the isotopes present in the LAr).

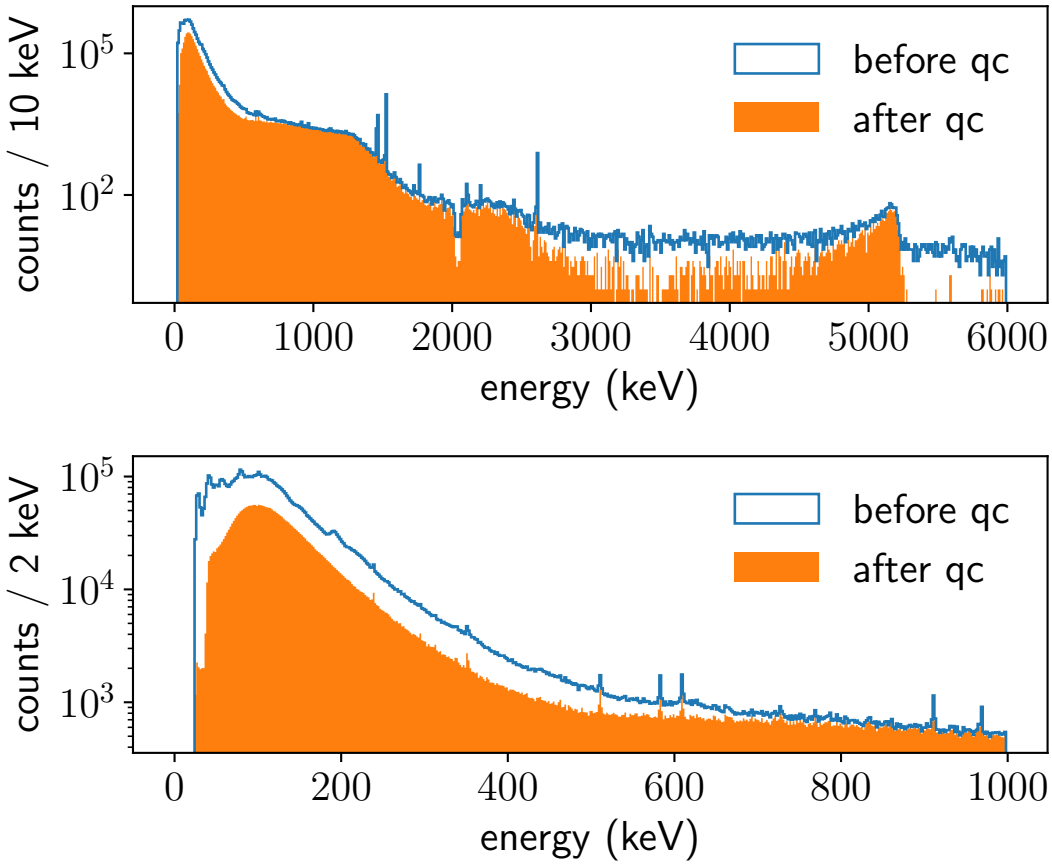

Figure 9.1: Effect of quality cuts on LEGEND-200 data. The top plot shows the full energy spectrum from 0-6 MeV. As discharges are reconstructed with an essentially random energy we see that the quality cuts have an effect across the full energy range. The quality cuts do, however, have a larger effect at lower energy as many unphysical events will be reconstructed to have a low energy. The bottom plot shows just the energy range from 0-1 MeV. Here the effect of the quality cuts is very clear as before we have spikes in the energy spectrum, and after we recover a smooth spectrum dominated by Ar-39.

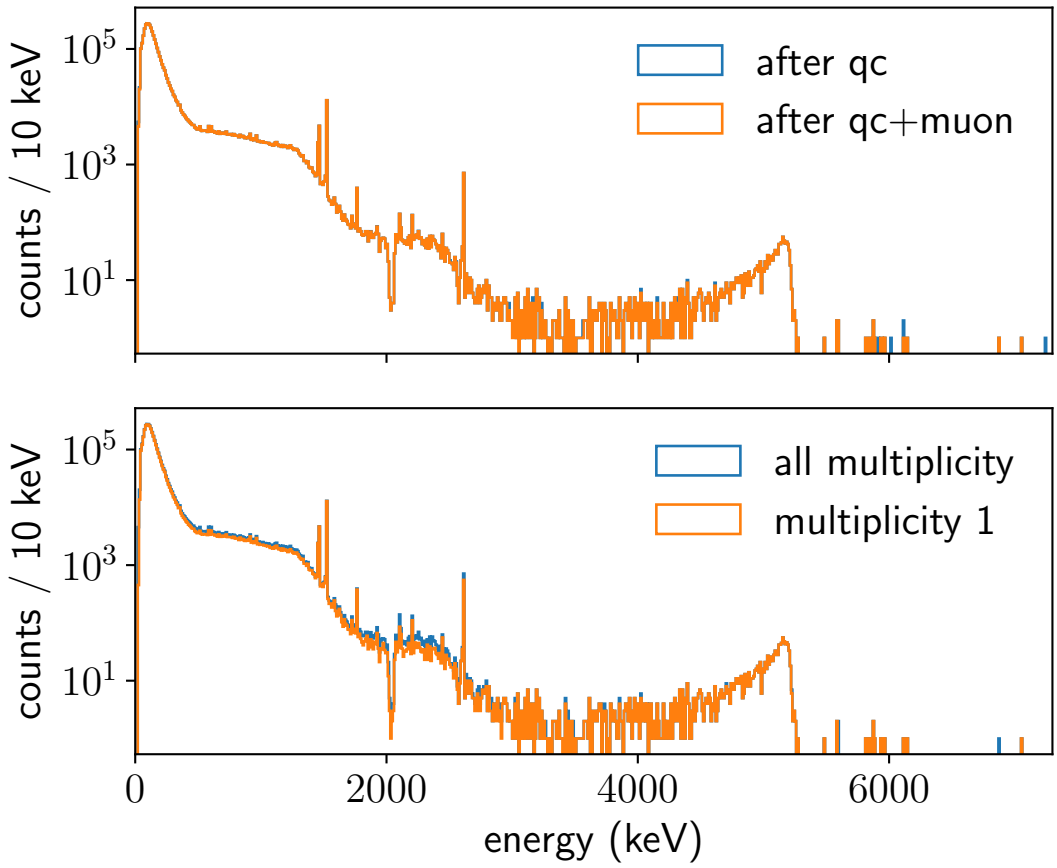

Figure 9.2: On the top panel is the effect of the muon cut on LEGEND-200 data after quality cuts. Most events removed are at very high energies. On the bottom panel is the effect of the multiplicity cut which is more even in energy. After applying the multiplicity cut we obtain the raw physics' spectrum for background analysis. At low energy around 100 keV the spectrum is dominated by Ar-39 beta decay, above this is the $2\nu\beta\beta$ region with the two prominent K-lines at around 1.5 MeV. Between the K-lines and the Tl-208 line at 2.6 MeV is the $Q_{\beta\beta}$ region. Above the Tl-208 line is the alpha region which is dominated by Po-210 at around 5 MeV.

As our rates are dominated in the germanium detectors by these discharges work went into understanding them. It was found that these rates were dominated by a few detectors. In particular two COAX detectors were found to be at the multiple mHz level possibly due to their lower capacitance and one of the new very large ICPC detectors which might be due to it being closer to the copper supports. This will be minimised in the next deployment.

The muon veto is applied to remove muon events which are primarily found at very high energies as shown in Figure 9.2. Then the final cut applied at this stage is the multiplicity one cut which means that each event only is in a single detector. Higher multiplicity spectra can also be analysed with multiplicity two of particular use in the background model (closer background sources are more likely to have higher multiplicity events) and even higher multiplicities used for BSM searches.

The multiplicity all spectrum is compared to the multiplicity one spectrum in the bottom panel of Figure 9.2. With these cuts applied we get what will be termed the raw physics' spectrum.

9.2 Prominent Backgrounds Before Analysis Cuts

The physics spectrum for LEGEND-200 is shown in Figure 9.2 along with a labelled version split down by region in Figure 9.3. At low energy (<0.5 MeV) the raw spectrum is dominated by the beta decay from Ar-39 which is abundant in atmospheric Ar. This is the limiting factor in the low energy physics of LEGEND-200 as many possible BSM searches are swamped by this background and is why LEGEND-1000 will be using underground LAr. Above this is the $2\nu\beta\beta$ spectrum which is mixed with the Compton spectrum from the K-lines, K-40 and K-42, at 1460 and 1525 keV respectively. K-42 is formed from the beta decay of Ar-42 but has a longer half life, forming charged ions that may drift in the liquid argon volume after formation through the convection currents and possibly under the influence of the electric fields. This causes it to collect on the outside of the nylon mini shrouds around the germanium detectors, or, for the small amount already inside, on the n+ electrode of the detector (where the high voltage is applied) and cabling. In LEGEND-200 we see that any changes in the electric field configuration, such as voltage changes on detectors, can cause spikes in the K-42 rate while the electric field stabilises. K-42 always undergoes beta decay followed by gamma decay which can be strongly suppressed by an anticoincidence cut on the LAr scintillation light as will be shown in future sections. K-40, meanwhile, is a primordial isotope and is present in many components such as the cables. Due to this, we see a strong position dependence in the K-40 rate. It can undergo a range of decays: $\beta+$, $\beta-$ and electron capture. The gamma line at 1461 keV is from the electron capture branch and therefore has no associated beta. For this reason it is not suppressed by the LAr but is suppressed instead by the PSD cut.

Above the K-lines the spectrum is dominated through the $Q_{\beta\beta}$ region around 2039 keV by Bi-214 from the U-238 decay chain and Tl-208 from the Th-232 decay chain as well as to a lower level beta decays from the K-42. Both decay chains are shown in Figure 9.4 with these two gamma lines towards the bottom of each chain. U-238 and Th-232 and their daughters are present in low activities in many components the other gamma lines from these chains extend all the way down to 200 keV, many of which are visible in LEGEND-200 and are labelled in 9.3. In general the Th-228 decay chain lines are suppressed more by the LAr, due to the presence of coincident gammas and betas, than the U-238 chain lines. For the two main background isotopes, it can be seen from Figure 9.4 that the Tl-208 2.6 MeV gamma line often comes (85% of the time) with an associated gamma

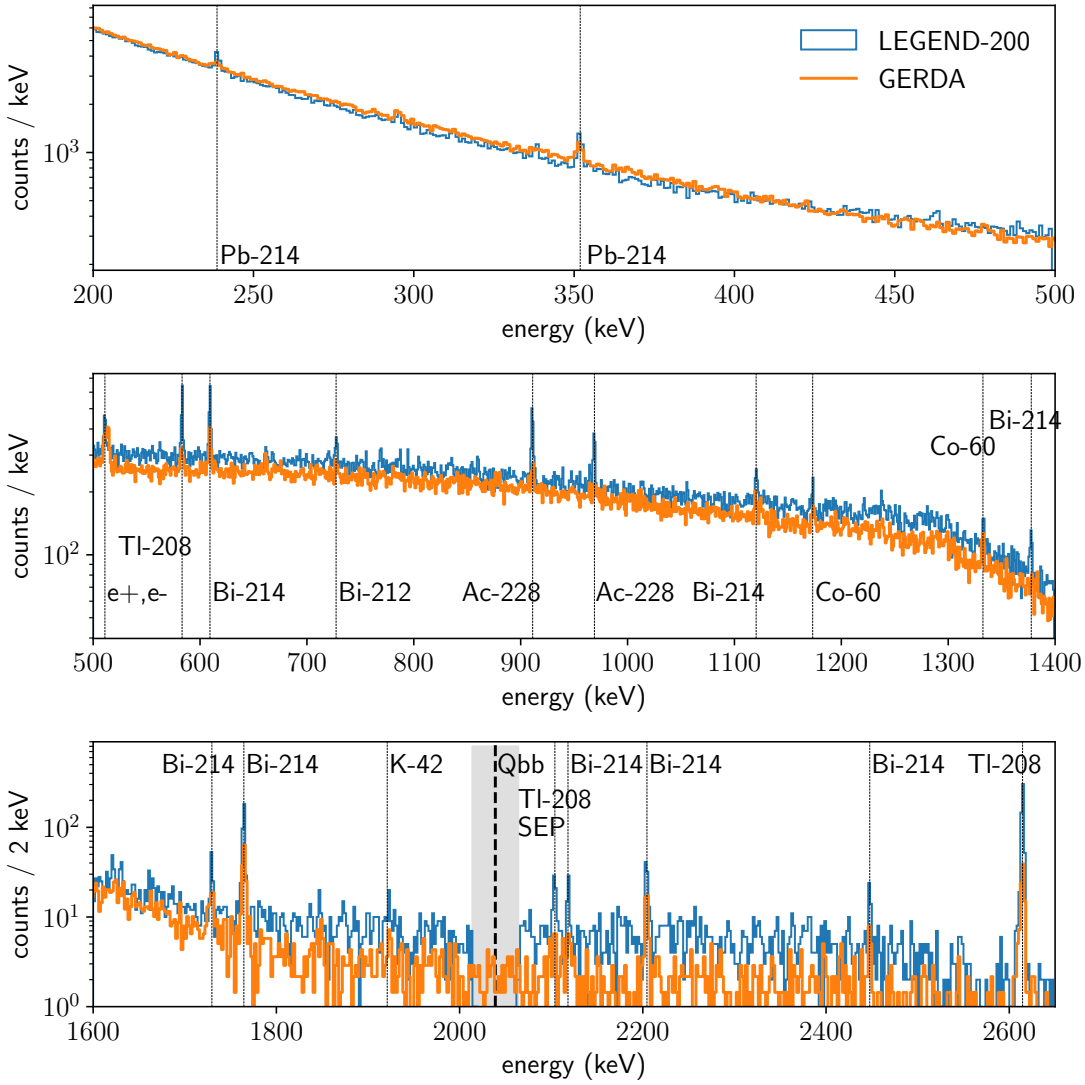

Figure 9.3: Raw physics' spectrum for LEGEND-200 split by energy range with the major gamma lines labelled. The top inset is 200-500 keV where the prominent background is the Ar-39 beta decay spectrum. The middle is the region below the two K-lines which is the $2\nu\beta\beta$ dominant region and finally the final plot is the region above the K-lines which includes $Q_{\beta\beta}$ and the surrounding region used to estimate the background index. A comparison with GERDA Phase II data is also shown.

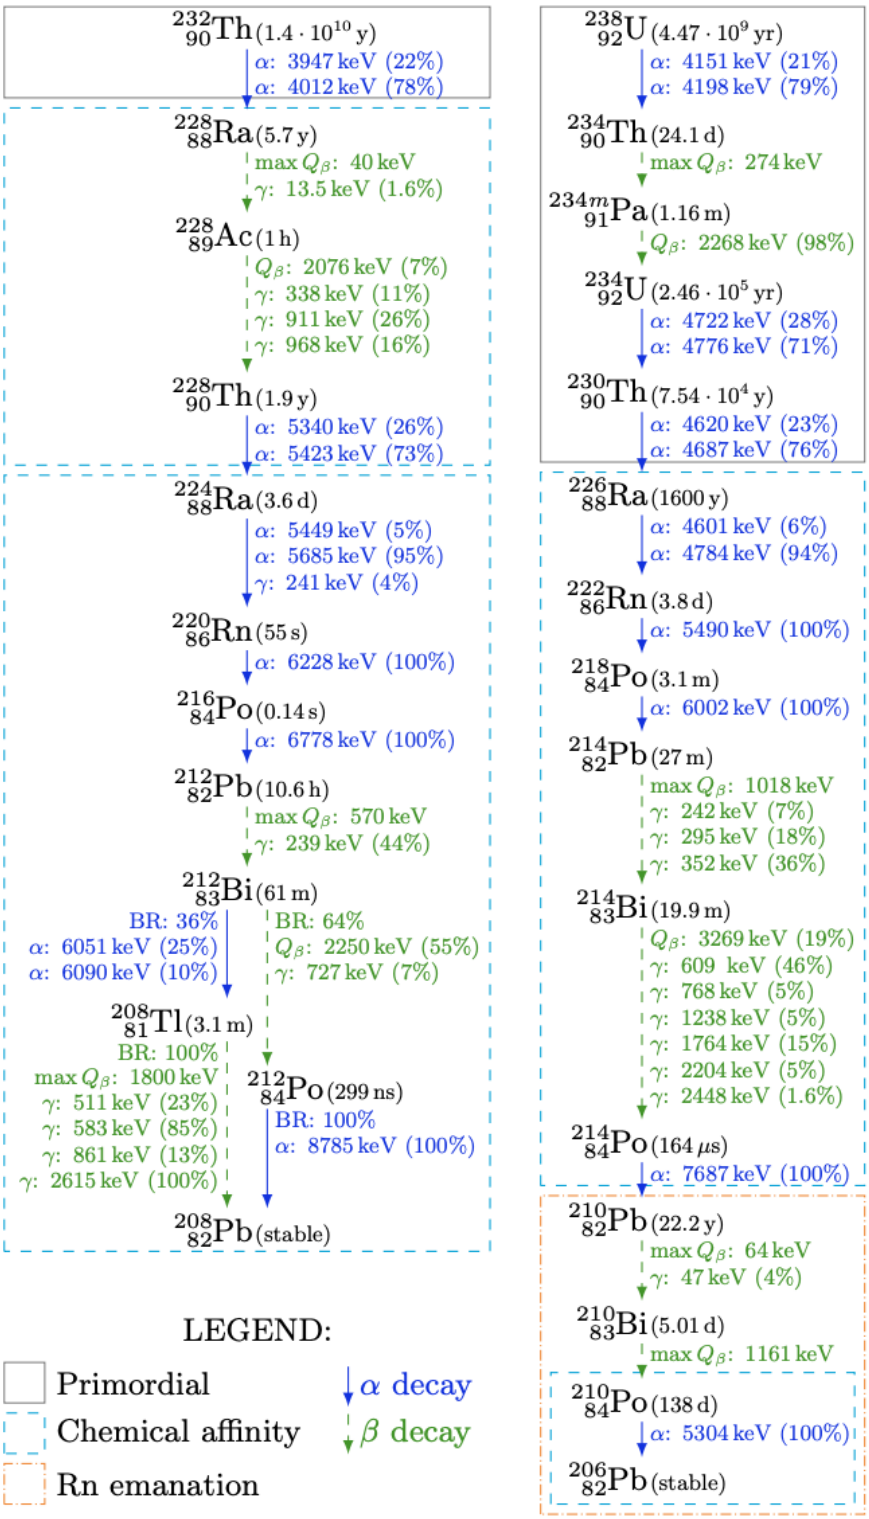

Figure 9.4: On the left is the decay scheme of Th-232 and on the right the decay scheme of U-238. For Legend-200 the major gamma lines of interest are towards the bottom of each chain with Tl-208 for Th-232 and Bi-214 for U-238. Apart from the gamma lines, the alpha decays of Po, Ra and Rn are also of interest. Alpha decays will only be detected when they occur near the p+ electrode and due to the dead layer will be at lower energies than the decay Q value. This means they can appear in the ROI if sufficiently depressed. From Ref. [2]

at 583 keV for example, while Bi-214 is a beta decay which can have a range of different coincident gammas but for LEGEND-200 the most important is the 2204 keV line and the 2448 keV line as these are above $Q_{\beta\beta}$. Therefore, how well these backgrounds are suppressed varies with source location, a further away Th-228 source will not be well suppressed as the coincident gamma is likely to be emitted in the opposite direction away from the detector. While for the Bi-214 if the beta can't escape the material or is in a close location then it is less likely to be suppressed.

Finally, above the Tl-208 line at 2614 keV the spectrum comprises mainly of alpha decays in particular from non lead supported Po-210 at around 5 MeV. These are events in the p+ electrode, which, produced by B implantation, is a conductive layer only a few tens of micrometer thick. It is much thinner than the conductive layer of the n+ electrode (the dead layer) which is produced by Li diffusion and can exceed 1 mm in depth. Electron hole pairs produced in this p+ electrode layer will recombine right away and are therefore not collected within the detector. This creates the distinctive shape we see in the spectrum of Figure 9.2 at around 5 MeV of a peak populated by the events in which alpha particles penetrate the detector at shallow angles (albeit still below the Q value of the decay as some energy will be lost in this dead layer), and a tail extended at lower energy for those at larger angles. They are suppressed almost totally by the PSD cut. The lack of events above 5.2 MeV implies low contamination of Ra-226 on the detector surfaces. Finally, the dip in the energy spectrum at $Q_{\beta\beta}$ is due to the blinding.

In Figure 9.3 a comparison between the LEGEND-200 and GERDA spectra is shown. We expect to see some increase over GERDA in the spectrum due to more supporting materials for the increased mass of detectors and the increase in individual detector's mass which results in high gamma-ray detection efficiency. In both cases the improved LAr anti coincidence in LEGEND-200 ensures we veto this background more efficiently. Looking at the spectrum it can be seen that many of the gamma lines are more prominent in LEGEND-200 than seen in GERDA, as well as the Compton continuum being systematically higher also. In particular, the U-238 decay chain lines such as Ac-228 and Bi-214 are much more prominent implying elevated levels of background. Most clear is the Tl-208 line from the Th-232 chain at 2614 keV which is more prominent. Interestingly the Kr-85 line at 514 keV (seen in the GERDA spectrum as a widening of the 511 annihilation peak) which is a possible impurity in the LAr does not appear to be present in LEGEND-200. This could be due to a different LAr source for LEGEND-200.

9.2.1 Background Homogeneity Studies

To understand whether the background is homogeneous or concentrated in parts of the array I broke the rates down by various selections including by: string,

position and detector. Additionally, I studied how the rates of specific samples of events changes across strings, positions in the array and detector type. The event samples include full-energy absorbed gamma-lines, which can be sharply tagged by their energy and provide direct and robust info on the presence of the gamma-ray emitters such as the K-peaks, Tl-208 and Bi-214. Other samples contain in large windows events which are expected to be dominated by specific background contributions such as $2\nu\beta\beta$ and alphas. A simple counting analysis was done in a 20 keV window around the peak for each of these regions. The $2\nu\beta\beta$ region in particular (1-1.3 MeV) was used both to find detector issues and to validate that detectors were working well as we know the expected rate in this region so lower rates could indicate a detector issue.

The region of interest (ROI) used in LEGEND-200 is from 1930-2190 keV which is limited by the gamma lines at 1921 keV from K-42 and 2204 keV from Bi-214. Additionally, a 10 keV window around the two gamma lines from the Tl-208 SEP at 2104 keV and the Bi-214 FEP at 2118 keV are excluded from the ROI. Finally, as this data is blinded we have a 50 keV window in the centre of this region which is excluded leaving a region of 190 keV in total.

First, breaking the rates down by detector. Figure 9.5 shows the rates in the Tl-208 gamma line. On the top of the panel is the background rate (in counts / (keV ton yr)) for each individual detector showing their positions in the array and on the bottom is the raw number of counts in each string (note strings 6 and 12 are empty strings). In general the rates are homogeneous within statistical uncertainties, although there is one possible hotspot at the bottom detector of string 1 in the Tl-208 peak. A similar picture is also seen in the Bi-214 gamma line at 2204 keV. Comparing this to the same break down in the ROI region (Figure 9.6) we do not see a hotspot, but this might simply be down to lower statistics. This hotspot is not significantly higher than the other detectors but after disassembly the components nearby have been kept separately for assay to see if there was any surface contamination (we do not expect a difference in bulk contamination as the components are identical to other detectors). Of the other detectors nothing conclusive can be said from the detector breakdown as some are higher in Tl-208 and others in Bi-214.

The next splittings were by time period and by detector type. The data in LEGEND is divided into periods where we have a consistent hardware status, this means they may be of different lengths. The background rates here, of course, take into account the exposure, but the statistical uncertainties will be larger for the shorter periods. For the breakdown by detector type we have 4.5 times more exposure in ICPC detectors than the other detector types which are all with 1kg yrof one another.

In Figure 9.7 is a look at the background rates for each period, the exposure

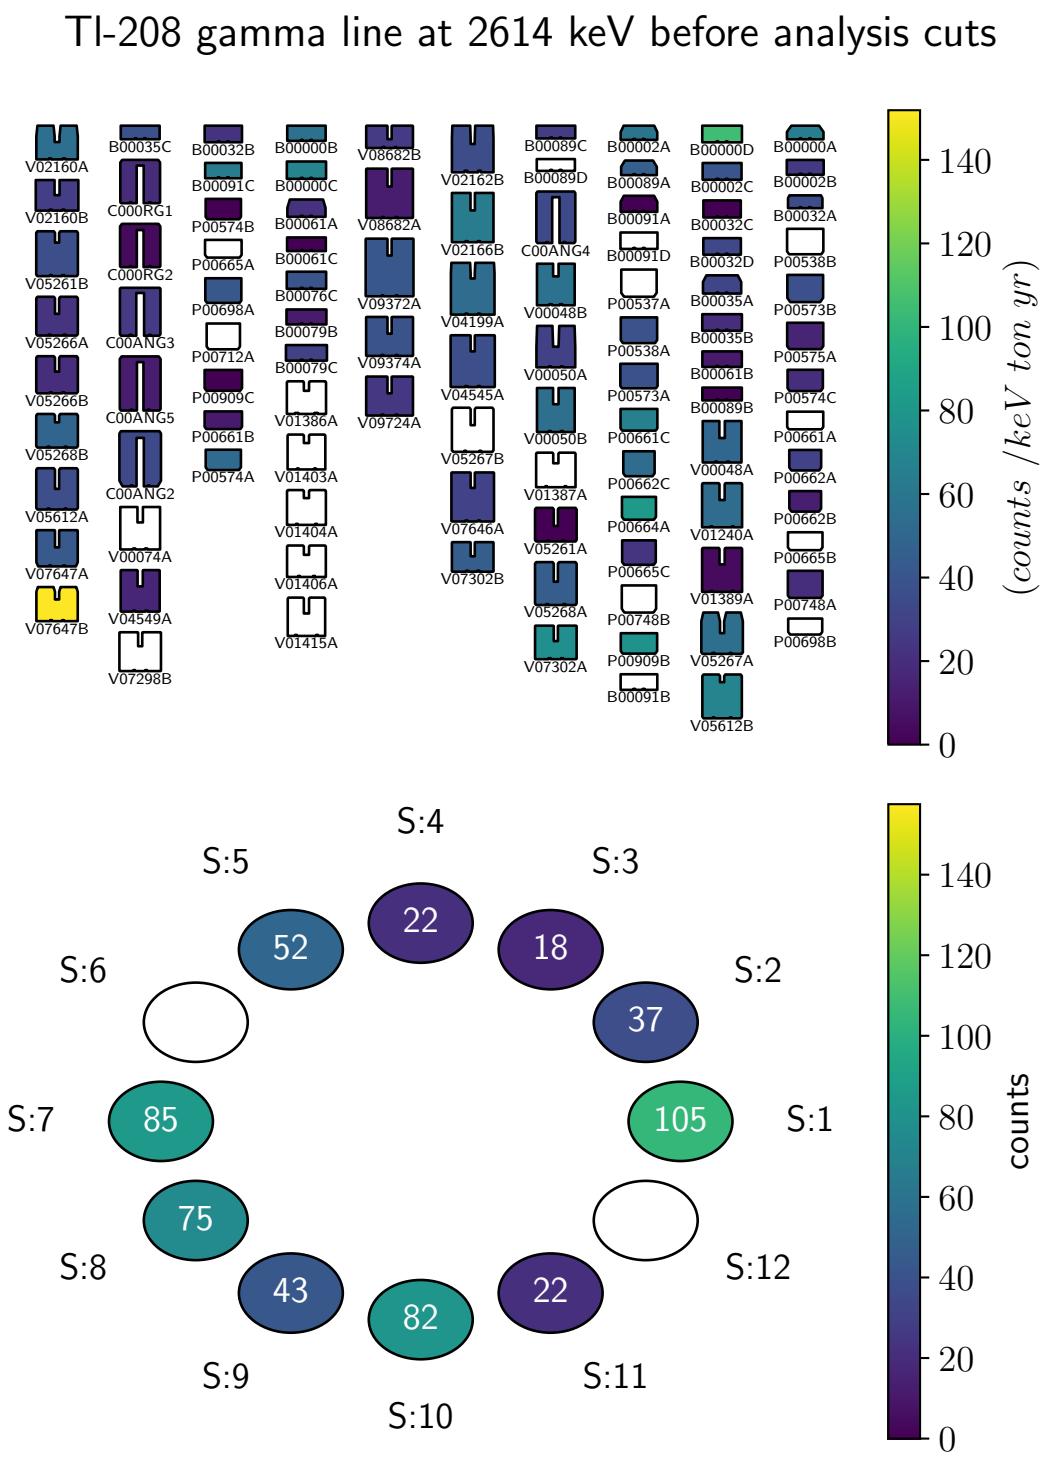

Figure 9.5: Background rates in the Tl-208 FEP at 2614.5 keV broken down by detector on the left. On the right is the number of counts in each string. There is a possible hotspot at the bottom detector of the first string.

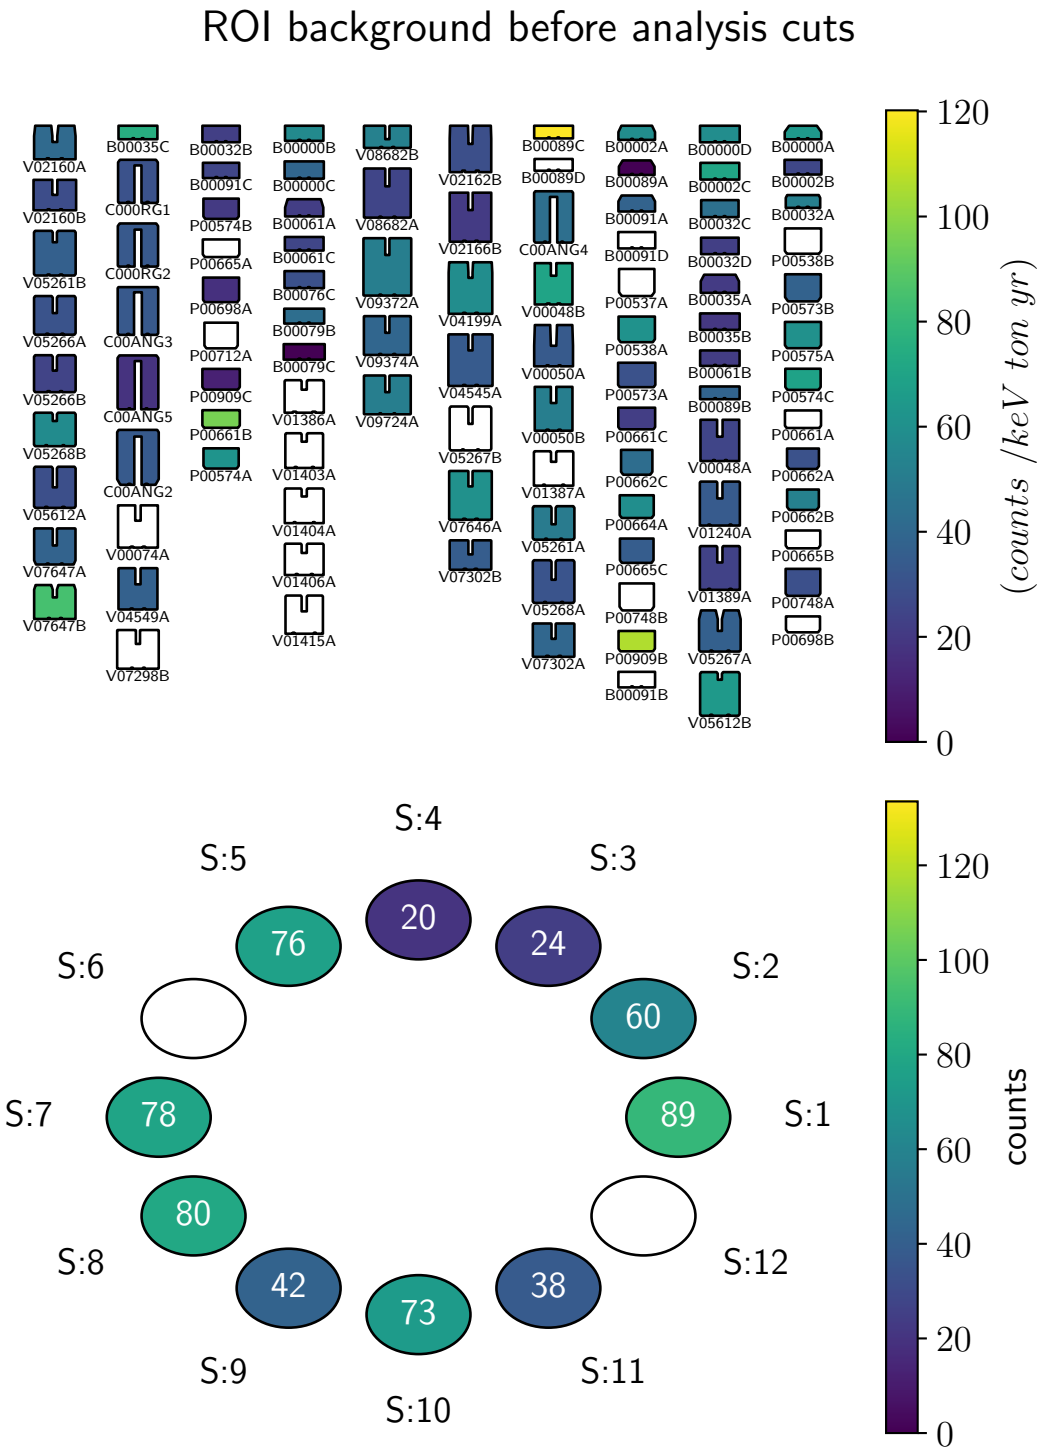

Figure 9.6: Background rates in the ROI from 1930 to 2190 keV, excluding the 2 gamma lines, broken down by detector on the top. On the bottom is the number of counts in each string. No particular hotspots are seen in the ROI, the highest detector B00089C had a low exposure, so its rate is compatible with the other detectors.

weighted average is also included as the orange band on the plot. On the left side of the plot is the BI which is decreasing in time, this is expected as the tail of alpha-induced events with degraded energy in the ROI decays away. This was found to be consistent with the rate of events in the alpha-dominated part of the spectrum between 3-7 MeV (shown on the right side of the plot) where we see the same effect by period. This implies that to a large extent the BI is dominated by Po-210.

Figure 9.8 instead shows the same plot but by detector type, on the left-hand side it can be seen that the PPCs have a slightly higher BI but due to the lower statistics this is not significant. This may be down to degraded alpha events on the passivated surface of the PPC detectors as we do not see any alphas for these detectors at higher energies. Otherwise, all detector types are consistent within uncertainties. The difference between the detector types which are higher in the alpha region and those higher in the ROI seems inconsistent with Po-210 being the dominant background in the ROI but may be explained by the location of the events on the detector surfaces and how this impacts the ratio of the tail to peak in the alpha spectrum.

Looking at the right-hand side for the alpha region, we see a strong detector type dependence. COAX detectors with their p+ electrodes covering a larger area have the highest rate, while ICPC detectors and BEGes are lower with both having similar electrode areas but the ICPCs having large masses. Finally, the PPC rate is very low due to having the smallest p+ electrode but also possibly due to alphas in these detectors being degraded to lower energies due to the passivated surface. This shows that having a smaller p+ contact is beneficial for reducing the alpha background, but we expect to remove these events with a very high efficiency using PSD cuts.

Finally, the data was broken down by position. We split the array into 5 regions, the top, top middle, middle, bottom middle and bottom. This position dependence is shown in Figure 9.9 for K-40 which has particularly clear gradient down the array with the top detectors almost a factor of 2 higher than the bottom detectors. This is likely due to the K-40 being present in the cables. No significant position dependence was seen in the other gamma lines except those due to the possible hotspot detectors.

In summary, the BI was found to be homogeneous across the array both by detector, by position and by detector type albeit within the limited statistics. A decrease in time was seen which was consistent with the alpha background decaying away. One possible hotspot was identified in both Tl-208 and Bi-214 gamma lines at the bottom detector of the first string.

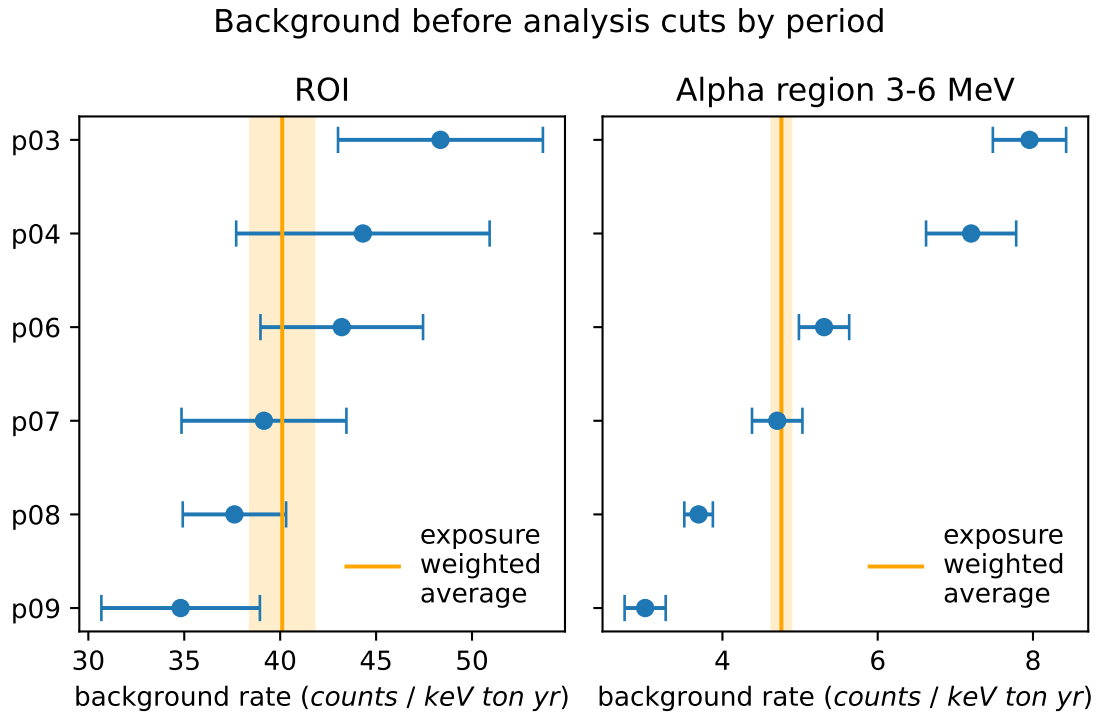

Figure 9.7: Breakdown of the background rate in the ROI by period on the left. On the right is the equivalent for the alpha region from 3-7 MeV. The background in the ROI is decaying in time as expected which coincides with the decrease in the ROI background. This is likely due to the alpha background decaying away as seen in the right plot. The rates in the alpha region were fitted to an exponential decay and found to be consistent with the half life of Po-210.

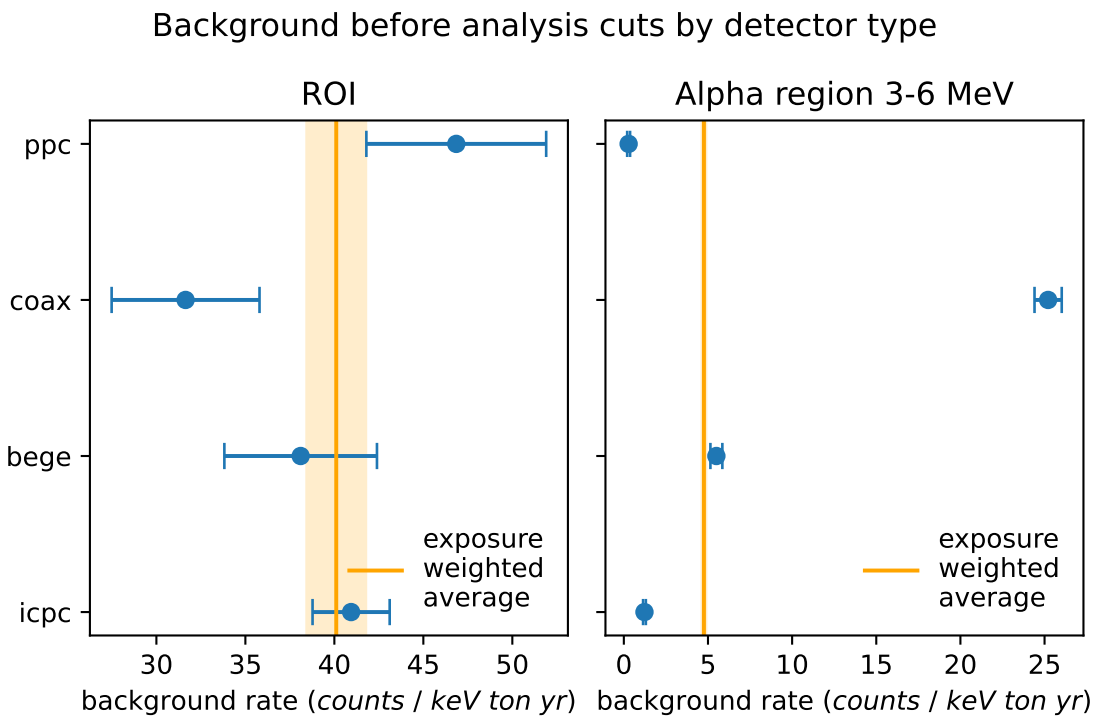

Figure 9.8: Breakdown of the background rate in the ROI by detector type on the left. On the right is the equivalent for the alpha region from 3-7 MeV. There is a mild hint that the BI for PPC detectors might be higher than for BEGes and ICPCs, but this is not statistically striking as it is still compatible within about 2σ. The COAX detectors have a higher rate of alpha events due to the larger p+ electrode while the PPCs have a very low alpha rate in this window. This is likely due to the smaller p+ electrode and alphas being degraded to lower energies due to the passivated surface.

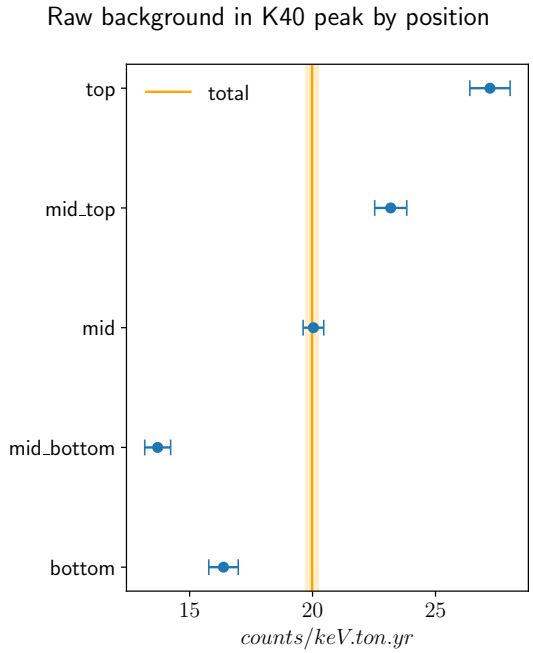

Figure 9.9: Breakdown of the rate in the K-40 per position in the array. A clear trend can be seen from the top to the bottom, implying the source is likely to be the cabling or another source that decreases through the array.

9.2.2 Background Model

Building on the work described above, a background model was built based on Ref. [106] by the background team. This also used data from special background runs as well as inputs from the radioassay campaign. A comparison was first made between the expectations just from the radioassay measurements and the data as shown in the top panel of Figure 9.10. The data showed a significant excess over the assay results.

Monte Carlo simulations were run for various components of the array and a binned likelihood fit to the data performed. Multiplicity two data was also used to better constrain the background location as closer components are more likely to have events in multiple detectors. This model works mainly on distance from detectors where close sources will have a different ratio of peak to Compton and different ratios of the various peaks compared to farther sources. This means that it is ambiguous for sources that are at a similar distance, i.e. it will not separate the detectors support plates from their front end electronics. The full background model fit is shown in the bottom panel of Figure 9.10.

The model's validity was confirmed using both the goodness of fit and projections, for example the reconstructed counts for each string. These confirmed the results of the homogeneity studies described above indicating a uniform background and disfavouring having hotspots in the array. The background model favours this increased activity being at a mixture of close and far distances from the detectors, and it disfavours all of this excess being on a single component.

9.3 Liquid Argon Anti-coincidence

From this point onward the dataset discussed will be the smaller $0\nu\beta\beta$ dataset which has a lower exposure (48.3 kg yr). It also means some detectors are no longer used such as the COAX detectors for which PSD routines are still being developed. The raw physics spectrum for this dataset is shown in Figure 9.11.

The first requirement of a $0\nu\beta\beta$ event is that all the energy is contained within a single detector. We have two analysis cuts to remove events where this is not the case. The first was described and applied above which is the multiplicity cut where energy is split between detectors, the other case is a coincidence between the LAr scintillation and a hit in a germanium detector which we remove by applying a threshold on the light detected by the LAr instrumentation. For this data release it was decided to use a simple condition of either 4 SiPMs above the noise threshold or a reconstructed energy of more than 4 photoelectrons in total. A more complex likelihood based classifier was also developed but not used in this release due to issues with crosstalk between germanium and SiPM channels. It is expected that

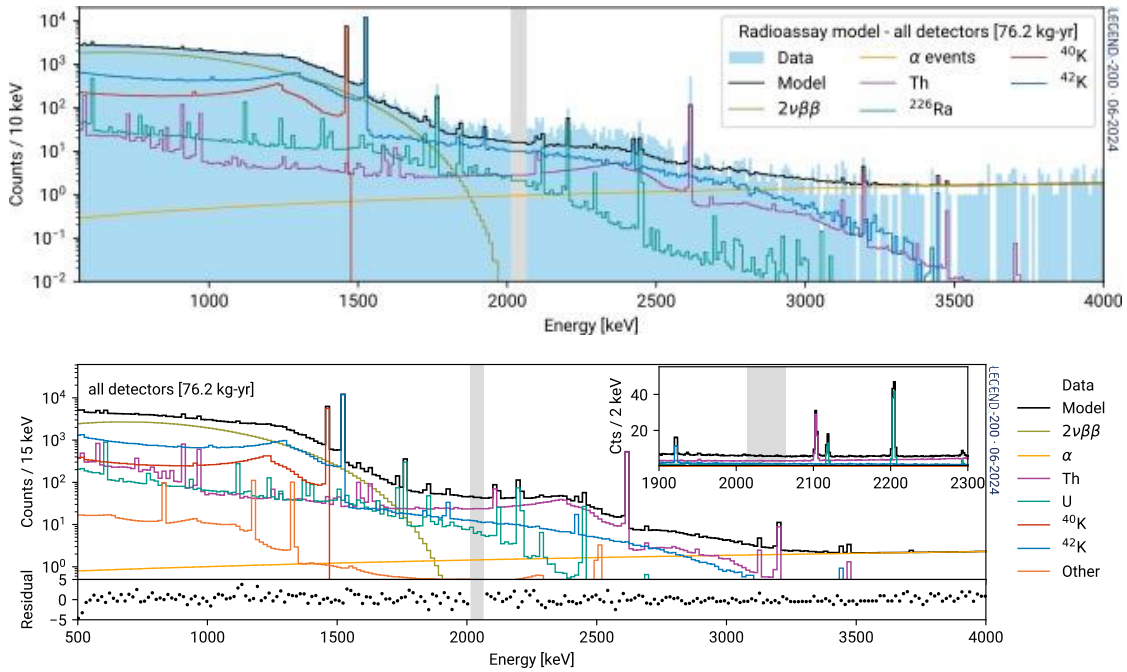

Figure 9.10: In the top panel is the priors from assay for the background model fit compared to data. It is seen that the data has an excess over the expectations from assay. In the bottom panel is the background model fit for LEGEND-200. The inset shows the residuals of the fit showing it fits the data well.

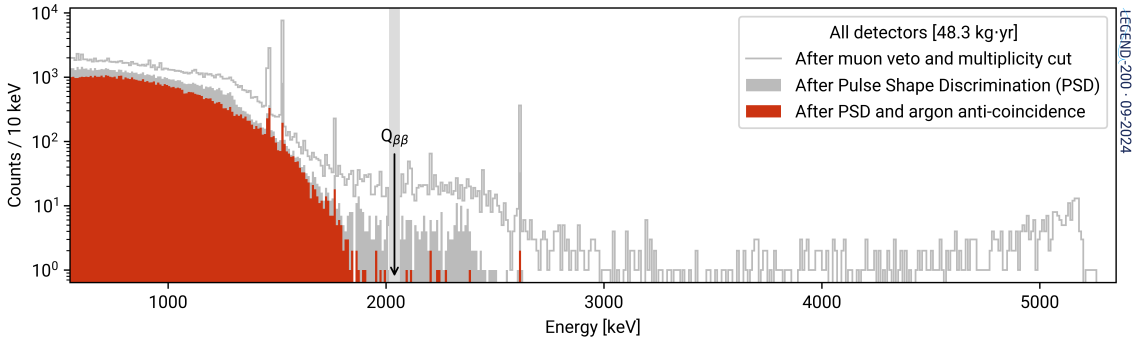

Figure 9.11: The full physics spectrum of LEGEND-200 for the $0\nu\beta\beta$ dataset. The effect of the PSD and LAr anticoincidence cuts are shown. After the LAr anticoincidence we obtain a clean $2\nu\beta\beta$ spectrum. The combination of the two cuts suppress the gamma lines and their Compton continuum by more than an order of magnitude.

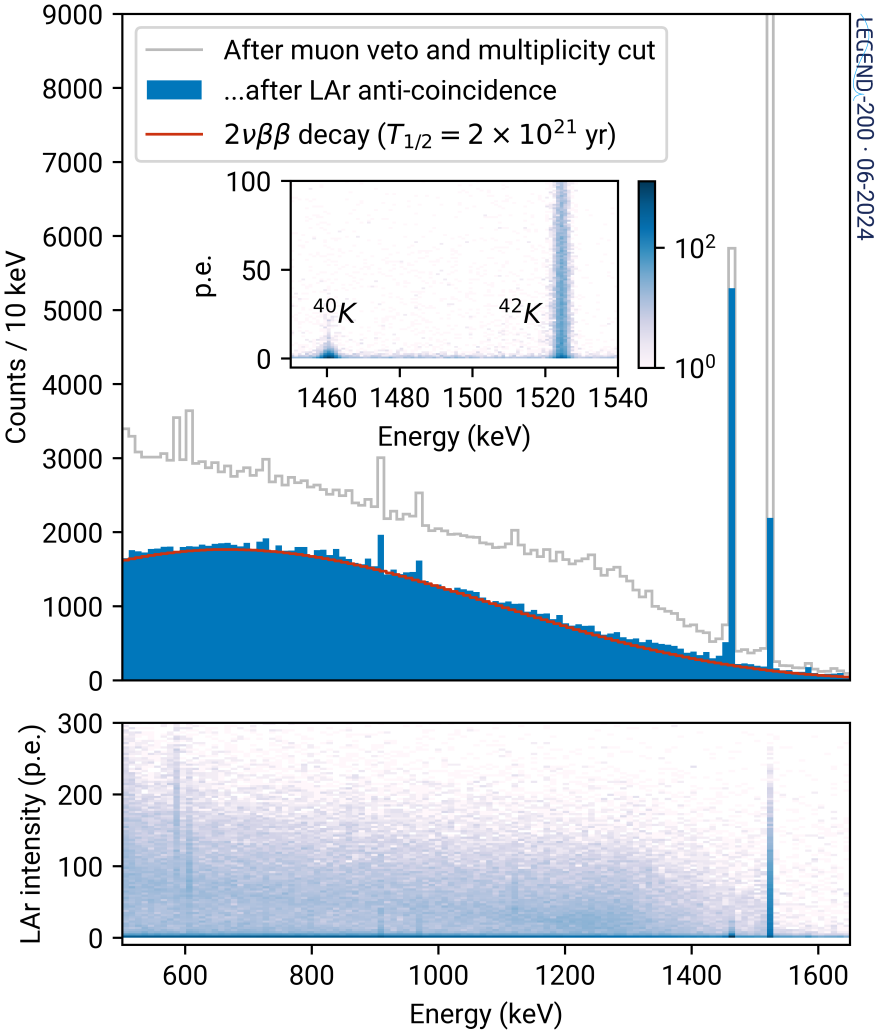

Figure 9.12: Spectrum for LEGEND-200 after LAr cut. Here the suppression of the K-42 line (1525 keV) which comes with an accompanying beta decay is clearly seen. While the K-40 line (1460 keV) is not suppressed as it has no coincident decay. The inset shows the distribution of the total energy deposited in the LAr in photoelectrons. A fit to the $2\nu\beta\beta$ region is also shown to illustrate that after the LAr cut we obtain a very clean spectrum.

this will improve the performance of the LAr anticoincidence cut in future releases.

Shown in Figure 9.12 is a zoomed in plot of the effect of the LAr cut on the data from 500-1700 keV with the inset showing the distribution of the energy in photoelectrons produced by the LAr group. The cut removes almost all the background in the $2\nu\beta\beta$ region (0.5-1.3 MeV) as expected giving us an almost clean spectrum. The K-42 peak is strongly suppressed as this comes with a corresponding beta decay which is easily detected in the LAr. In contrast, the K-40 peak is conserved as it has no coincident decay. The LAr cut removes 74% of the events in the ROI and, as shown in Figure 9.13, efficiently suppresses the Tl-208 SEP which comprises events where one photon from the pair production escapes the detector.

After the LAr anticoincidence is applied we obtain a very clean spectrum in the

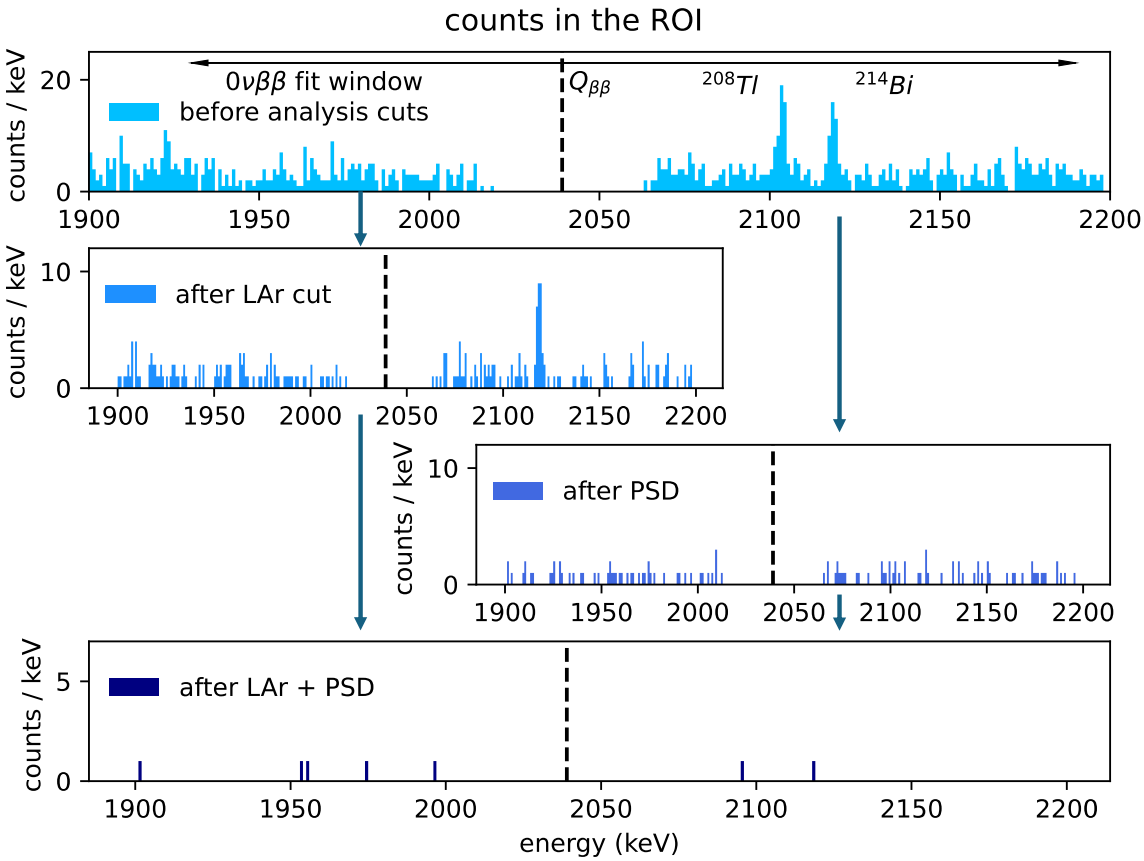

Figure 9.13: Events in ROI after each analysis cut is applied. In the top panel is the raw physics' spectrum after the multiplicity one cut, muon cut and quality cuts, second is the application of the LAr anticoincidence cut then the PSD cut (without the LAr anticoincidence cut) and finally the applying both the LAr and PSD cuts.

$2\nu\beta\beta$ region. The collaboration used this as a cross-check of the analysis and in particular of the active volume calculations as we know the expected rate for this decay. This cross-check proved particularly useful in identifying detectors which were not fully depleted. The rates for each detector are shown in Figure 9.14 along with the average for each detector type. This is expressed as a ratio of the rate to the expected rate. The ratios are all greater than 1 as we do still have some background in this region even after the cut. The COAX detectors have a higher rate of background as seen in other regions. Interestingly the PPC detectors are quite a bit lower implying a lower background in this region than the other detectors or possibly some issue in their active volume calculation such as the enrichment fraction or dead layer growth as these detectors have had their dead layers remeasured.

9.4 PSD

The second analysis cut applied are the PSD cuts which is applied to remove multi-site and surface events in the detectors. First to check the validity of the PSD cut we look at the survival fraction in the $2\nu\beta\beta$ region after the LAr anticoincidence is applied to remove most of the background. This is expected to be around 90% (as this is what we expect from the PSD calibration using the DEP) depending on the detector. Of course due to the presence of background this number is almost always slightly lower. This region is also used as an input for the efficiency of the PSD cut where unfolding this background is the major source of uncertainty. We found that all detectors are compatible with the expected value and therefore the PSD is performing well. The BEGe detectors are slightly (around 10%) lower than the ICPC detectors, although compatible due to the lower statistics. The PPCs were also slightly (around 5%) lower despite having a lower background here which could mean this background estimate is low.

Looking at the energy spectrum in the ROI in Figure 9.13, we see PSD removes around 78% of the data. This is higher than would be expected from calibrations where the efficiency in the Compton region is around 40% and could hint at a higher multi-site rate than expected due to a close source of background which would be more likely to have multiple scatters in a single detector than a source further away. The two gamma lines from the Tl-208 SEP and Bi-214 are both suppressed by more than 90% as expected especially for the SEP which is highly multi-site.

Figure 9.11 shows the energy spectrum after PSD is applied in the third panel. It can be seen that the whole spectrum is suppressed with the alphas being completely removed. All the gamma lines are also heavily suppressed with the Tl-208 decreased by a factor of 10 and similarly for the two K-peaks. No events survive

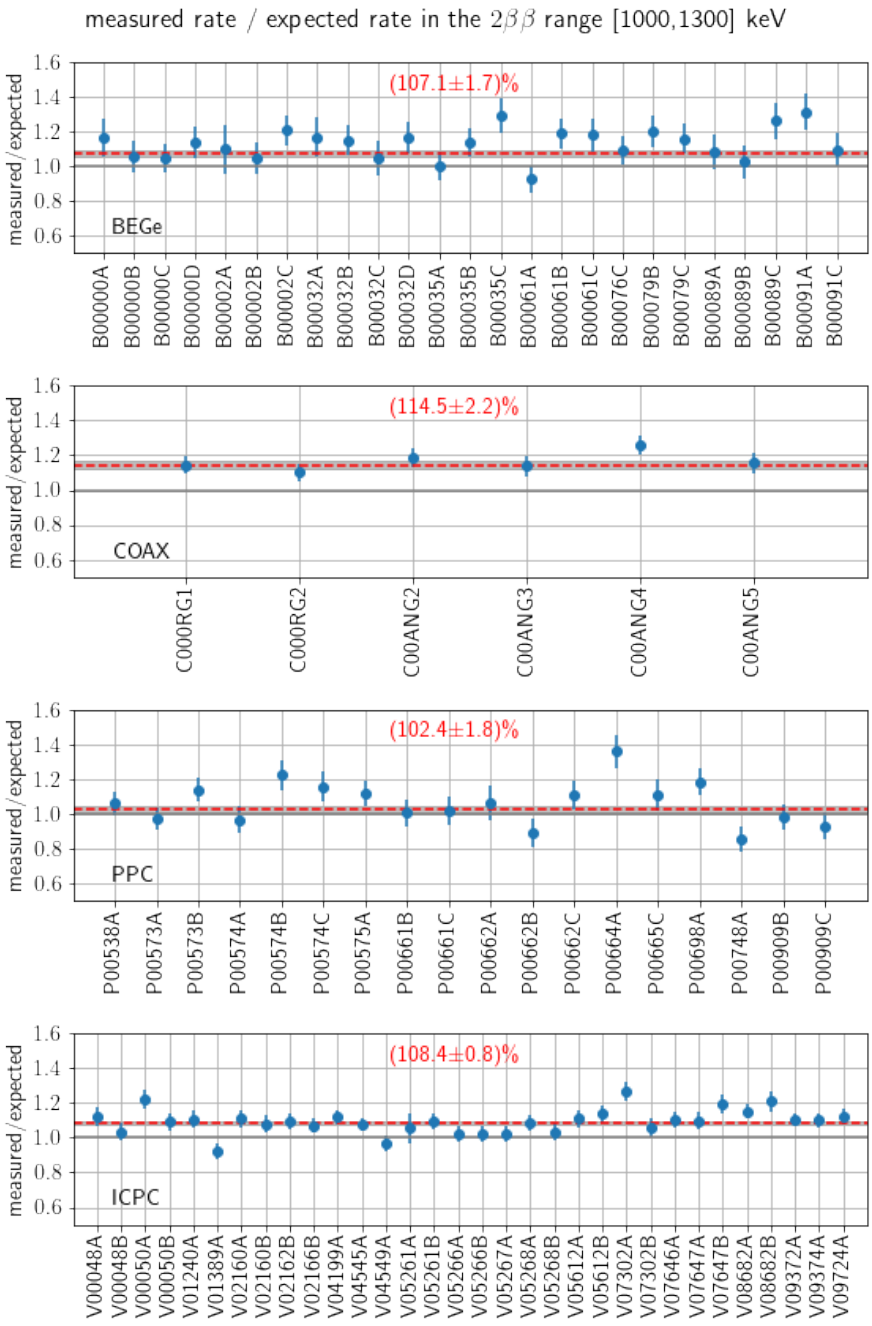

Figure 9.14: Ratio of observed to expected rate for the $2\nu\beta\beta$ region after LAr anticoincidence cut made by the collaboration. The average for each of the detector types is also shown. The COAX detectors come out slightly higher which is expected due to the higher background while the PPC detectors are lower which could imply a lower background in this region or an issue with the active volume calculation estimated during the detector characterisation.

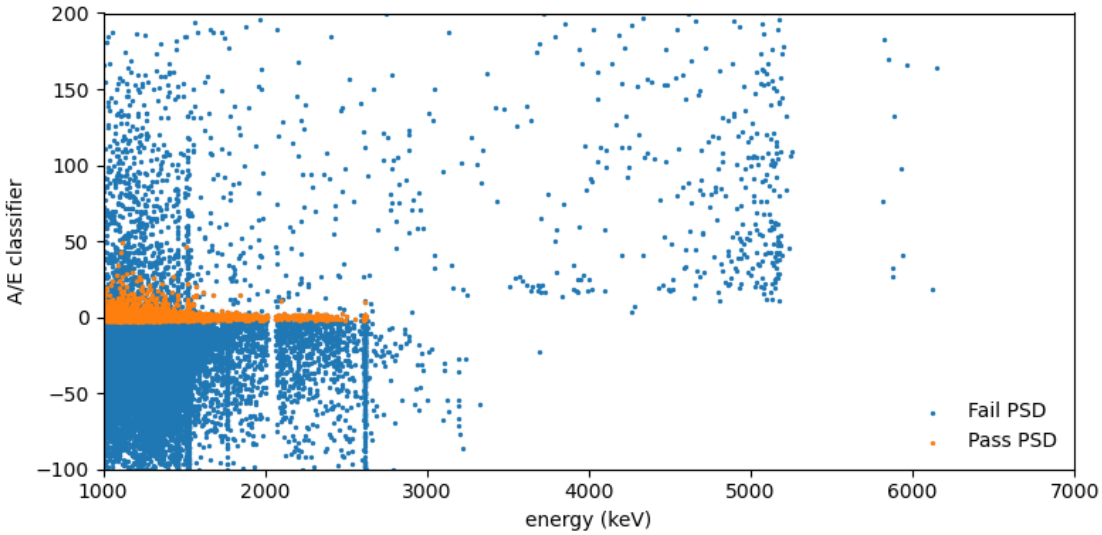

Figure 9.15: Scatter plot of the A/E Classifier for background data. In orange are the events passing the PSD cut and in blue are the events failing. The single site band can thus be clearly seen. At lower energy the background is mainly multi-site events which extend to lower A/E classifier values. With the two K-lines (1460 keV and 1520 keV) and the Tl-208 line (2.6 MeV) are clearly visible and also stretching up to higher values in the case of the K-lines due to events near the p+ electrode. The background in the ROI is mainly multisite but with some surface events. At higher energies the background is mainly surface alpha events at higher A/E values.

above the Tl-208 line which is a mixture of highly multi-site events and betas just above until the summation line at 3175 keV and alphas above this.

In Figure 9.15 is the A/E classifier for the background data with the events surviving in orange. The single site band can be clearly seen at 0 as these are the events surviving and the multi-site events extending below it. The two K-lines and the Tl-208 line are clearly visible stretching into this band with a few events also at higher values due to events near the p+ electrode. At higher energies above the Tl-208 line we are dominated by alpha events which are at high A/E values more than 10σ from the single site band. The background in the ROI is mainly multi-site events but with some surface events. One interesting feature is the presence of events at high A/E values surviving the cut, these are events in PPC detectors where instead of applying a high A/E cut we apply the LQ cut. This is a PSD cut looking at the end of the leading edge of the waveform which was found to work well in MAJORANA, more details can be found in [85]. The decision was made to only apply LQ and not both LQ and high A/E after a sensitivity study showed no improvement with both. This may change when higher statistics are available and will be discussed further in the next chapter.

The effectiveness of the PSD at removing surface events can be clearly seen with no events surviving above the Tl-208 line. Of course this only applies to the BEGe and ICPC detectors as we do not see many events in this region for the

PPC detectors.

9.5 LAr and PSD

Figure 9.11 shows the full energy spectrum before and after applying both the LAr and PSD cuts together, again only for the $0\nu\beta\beta$ dataset. The effect of the analysis cuts in reducing the background is clear with an almost complete suppression of the gamma lines, a very clean $2\nu\beta\beta$ region and no high energy alpha events.

Figure 9.16 displays a scatter plot of the A/E classifier against the reconstructed amount of LAr scintillation light in the ROI. This illustrates the complementarity of the two cuts with the LAr anticoincidence removing events at high reconstructed light intensity and the PSD cut removing events at low reconstructed light intensity. For example, a clear band of surface events can be seen with very low reconstructed light intensity in the LAr so are kept by this cut but very high A/E values so are removed by the PSD cut.

This brings us to the energy spectrum in the ROI after both cuts which is shown in Figure 9.13. In total only 5 events survive both cuts with more details on the counts in Table 9.1. No detector had more than one event and no time based correlations were found. Additionally, all waveforms surviving both cuts were checked by eye for any issues. One event was close to the LAr threshold at more than 3.5 p.e. with another just under 3 p.e. If a more sophisticated classifier was used, this event might have been rejected.

Taking these 5 events we can calculate an estimate for the final BI. This is shown in Figure 9.17 with the left panel showing the BI for each period and the right for each detector type. The orange band is the exposure weighted average for LEGEND-200 and the green is the GERDA result from [64] for comparison. After PSD and LAr anticoincidence the statistics are obviously very low, but we don't see any significant differences between periods or detector types. The final background estimate before unblinding was $0.54^{+0.30}_{-0.24}$ counts / (keV ton yr) which is compatible with the GERDA result of $0.50^{+0.16}_{-0.13}$ counts / (keV ton yr) and also with the LEGEND-200 goal of 0.2 counts / (keV ton yr) at 1.5σ. More exposure would be needed to confirm this result. The errors here are calculated using the Feldman-Cousins method [107].

time	Period	Detector	Energy (keV)	A/E value	Drift Time	LAr Energy	LAr Multiplicity
17-06-23 07:32:02	p06	V09374A	1974.73	-1.04	1296.0	3.66	3
22-08-23 04:53:22	p07	V04199A	1955.22	0.56	876.0	2.04	2
23-08-23 10:23:20	p07	B00061A	2095.72	-0.48	407.0	0.77	1
05-11-23 19:17:40	p08	V08682B	1996.49	-1.19	756.0	2.99	2
18-12-23 02:58:15	p08	V07647A	1953.14	-0.46	451.0	0.0	0

Table 9.1: Counts in the ROI after LAr and PSD cuts.

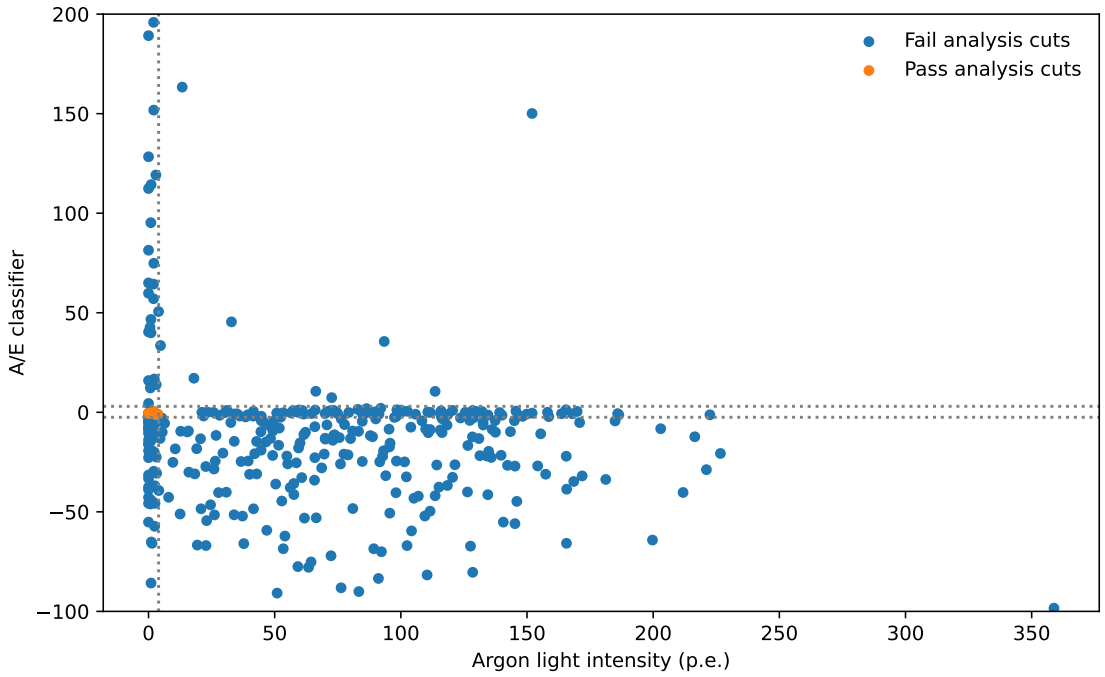

Figure 9.16: Scatter of the A/E Classifier against the light in LAr in the ROI energy window. This illustrates the complementarity of the two cuts. The grey dotted lines indicate the cuts on the two parameters. For the LAr energy this corresponds to 4 p.e. while for the PSD cuts these correspond to a high cut of 3 and a low cut of -2 (the actual cut value varies from detector to detector, but all are between -2.5 and -1).

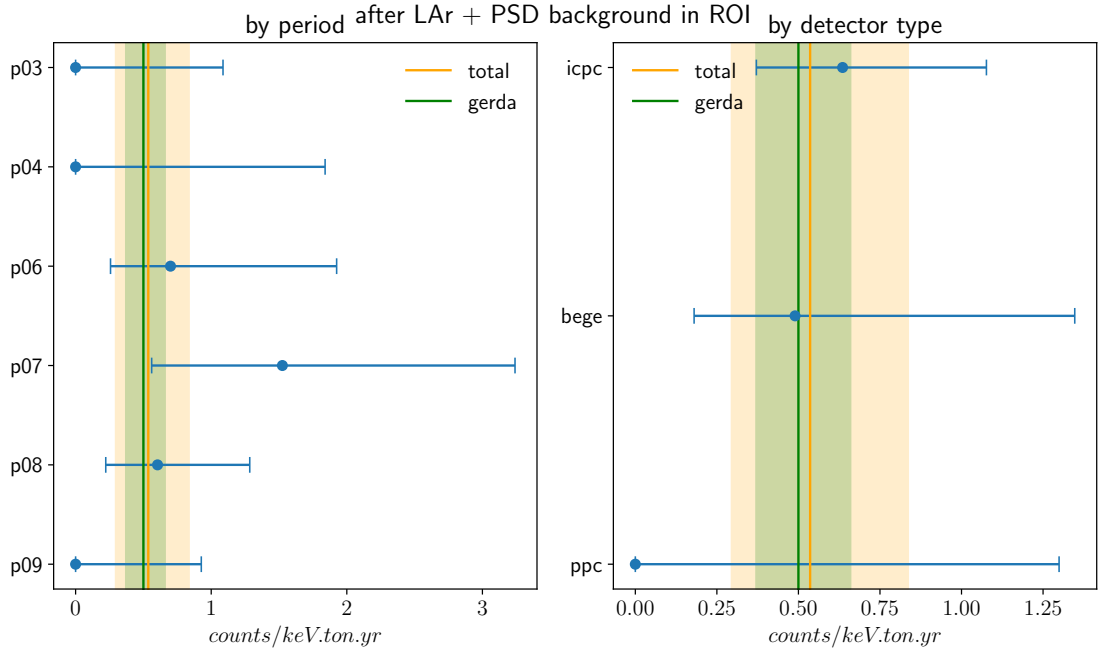

Figure 9.17: Background rate per period and per detector type in the ROI. The total background is shown in orange with a band for the 1 sigma uncertainties. Additionally, the GERDA background is included for comparison.

9.6 Summary

In this chapter the background in the LEGEND-200 dataset was studied. An elevated background compared to assay expectations was found in the raw dataset. This background was studied for its homogeneity, with one possible hotspot and some time dependence found. The background model built to explain this excess rules out a single source and instead favours a mixture of close and far sources. Additionally, data is being taken to investigate further. The LAr anticoincidence cut and PSD cuts were applied and the effects on the background shown with a nice suppression of the gamma lines and alphas from the PSD and a clean $2\nu\beta\beta$ region achieved by the LAr cut. Finally, both cuts were applied together with 5 events surviving in the ROI. The BI was measured and found to be compatible with the GERDA result and also with the LEGEND-200 goal at 1.5σ.

Chapter 10

Unblinding and Final Result

In this chapter the unblinding of the $0\nu\beta\beta$ dataset is described. This is then converted into a half life limit for the $0\nu\beta\beta$ decay of Ge-76 when combined with MAJORANA and GERDA. I wrote the blinding scheme for the data using the DAQ energy to remove events within 25 keV of $Q_{\beta\beta}$. The unblinded data was stored on a separate area of the disk that only a few people had access to. As described in Sec. 7.4 it was checked automatically in each run to ensure the validity of the blinding and I designed it so if any channel failed the data production would stop. As the data production lead I led the unblinding of the data, running the data production on the unblinded data and checking its validity.

The collaboration performed a statistical analysis using the systematic uncertainties and performance parameters derived in earlier chapters to extract the background index of LEGEND-200 and, when combined with GERDA and MAJORANA, to set a lower limit on the half life of the $0\nu\beta\beta$ decay of Ge-76.

10.1 Unblinding

With the routines frozen, the unblinding of the ROI proceeded in two stages. Firstly, one fifth of the data (which had been analysed previously before the blinding was applied) was unblinded to check the validity of the analysis. Then the rest of the data was unblinded to produce the final result.

After unblinding 2 more events were seen in the ROI bringing the total to 7. These events are shown in Table 10.1. As the data production lead I checked the events visually for any obvious issues.

One event was observed with an energy of 2040.26 keV, which corresponds to 1.6 keV above the $Q_{\beta\beta}$ value, which converts to 1.4σ when considering the specific detector energy resolution of 2.9 keV (FWHM). This event is in a PPC type detector with the pulse shape suggesting an event close to the p+ surface as seen in Figure 10.1, but probably still in the bulk-volume. Its fully compatible with our expectations and detailed simulations will be performed to pin down the

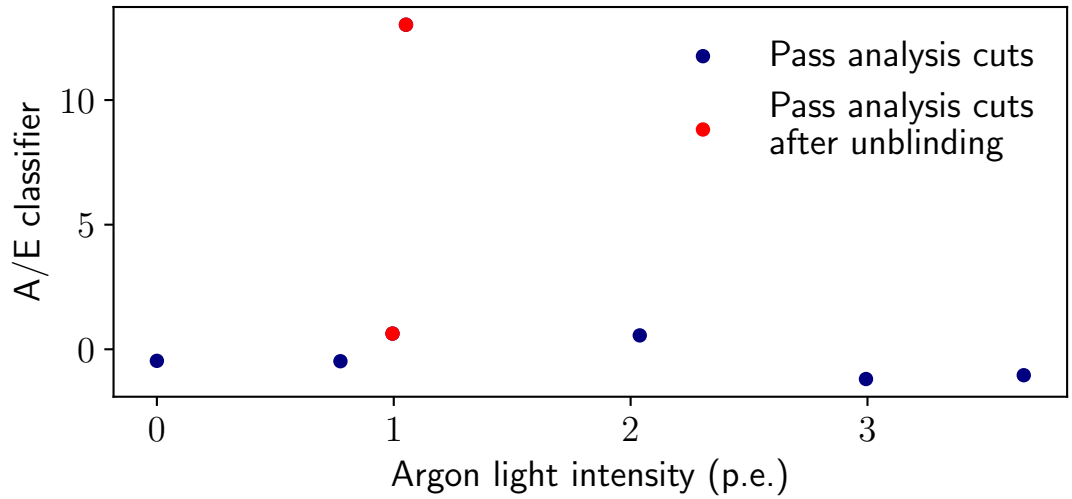

Figure 10.1: Scatter of the A/E Classifier against the light in LAr in the ROI energy window after unblinding. The two new events are shown in red. Both have only around 1 photoelecton worth of energy in the LAr. One has a very high A/E value while the other is more typical of a bulk event.

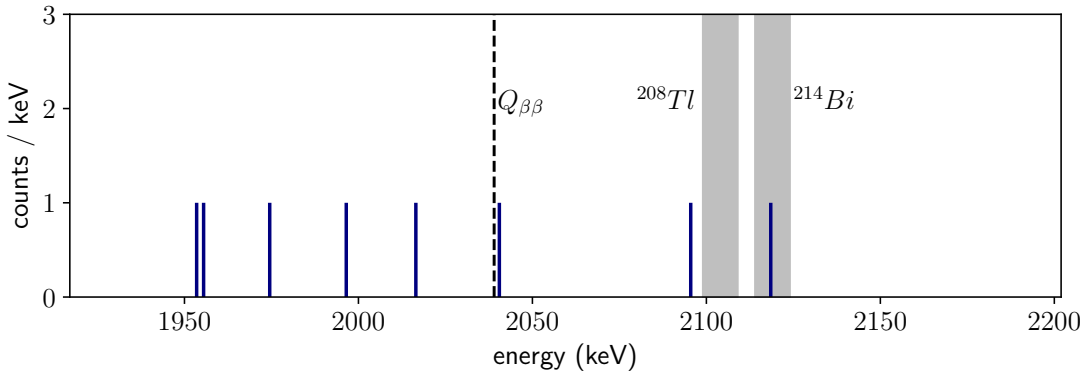

Figure 10.2: Events in unblinded background region after LAr and PSD cuts. 2 new events are added after unblinding, one at 2017 keV and one at 2040 keV. This gave a final number of 7 events in the ROI. The regions within 5 keV of the Tl-208 SEP and Bi-214 gamma line are shown in grey. These regions are removed when calculating the background index and performing the final fit.

time	Period	Detector	Energy (keV)	A/E value	Drift Time	LAr Energy	LAr Multiplicity
17-06-23 07:32:02	p06	V09374A	1974.73	-1.04	1296.0	3.66	3
22-08-23 04:53:22	p07	V04199A	1955.22	0.56	876.0	2.04	2
23-08-23 10:23:20	p07	B00061A	2095.72	-0.48	407.0	0.77	1
05-09-23 03:31:03	p07	P00661C	2040.26	13.02	374.0	1.05	1
29-10-23 05:16:24	p08	V00048A	2016.76	0.63	1367.0	0.99	1
05-11-23 19:17:40	p08	V08682B	1996.49	-1.19	756.0	2.99	2
18-12-23 02:58:15	p08	V07647A	1953.14	-0.46	451.0	0.0	0

Table 10.1: Counts in the ROI after unblinding.

exact location inside the detector where energy was deposited.

For these detectors it was decided to apply the LQ cut instead of high A/E due to the superior efficiency in removing surface events. However, in a very small number of cases, this means that events survive which would be cut by high A/E (our cut removes all events with an A/E value of more than 3). More data is needed to understand the interplay of these 2 cuts, which hopefully some of the background runs will help with as these are likely to have more surface events especially from K-42 betas. The other event is at lower energy and is in an ICPC type detector. It is a standard event with no particularly notable features. Both events have very little light in the LAr and are not near the LAr anticoincidence cut threshold of 4 p.e. or LAr multiplicity 4.

With this unblinding we can now show the full ROI which is presented in Figure 10.2. The events are, to first approximation, uniformly distributed in the ROI, although we do see more events below $Q_{\beta\beta}$ than above and there is no obvious clustering of events at a particular energy. The events are also uniformly distributed in time, again with no obvious clustering.

10.2 Final Result

Using the unblinded data the collaboration performed a frequentist analysis to extract the background index of LEGEND-200 and, when combined with the data of GERDA and MAJORANA, to set a limit on the half life of the $0\nu\beta\beta$ decay of Ge-76. As described previously the data in LEGEND-200 is grouped into partitions for each detector where for each partition we have: an exposure, an active volume estimate and containment efficiency, efficiencies on the PSD and data quality, efficiency of the muon and LAr anticoincidence cuts, and finally the energy bias and resolution at $Q_{\beta\beta}$ all with their associated uncertainties. The background index was taken to be common across all partitions as following the findings discussed in the last chapter we could not find any evidence of inhomogeneities in the background.

A simple gaussian was used for the signal energy distribution with the full likelihood function given our events ($\vec{\theta}$) as:

$$\mathcal{L}(\Gamma_{1/2}|\vec{\theta}) = \prod_k \left[\frac{\mu_k^{N_k} e^{-\mu_k}}{N_k!} \prod_{i=1}^{N_k} \frac{1}{\mu_k} \left(\frac{\mu_{b,k}}{\Delta E} + \frac{\mu_{s,k}}{\sqrt{2\pi\sigma_k^2}} e^{-\frac{(E_i - Q_{\beta\beta} - \Delta_k)^2}{2\sigma_k^2}} \right) \right. \\ \left. \times \frac{1}{\sqrt{2\pi\sigma_{\Delta_k}^2}} e^{-\frac{(\Delta_k - \hat{\Delta}_k)^2}{2\sigma_{\Delta_k}^2}} \frac{1}{\sqrt{2\pi\sigma_{\sigma_k}^2}} e^{-\frac{(\sigma_k - \hat{\sigma}_k)^2}{2\sigma_{\sigma_k}^2}} \right] e^{-\alpha^2/2} \tag{10.2.1}$$

where E_i are the energies of each of the N_k events within a partition k. Each event has an energy bias Δ_k, a resolution σ_k, which are both constrained by $\hat{\Delta}_k \pm \sigma_{\Delta_k}$ and $\hat{\sigma}_k \pm \sigma_{\Delta_k}$ respectively, and an efficiency ϵ_k. The number of background counts in the fit window of width ΔE is given by $\mu_{b,k} = \mathcal{B}\ \Delta E\ \mathcal{E}_k$ where $\mathcal{B}$ is the background index. Finally $\mu_{s,k}$ is the number of signal counts given by :

$$\mu_{s,k} = \frac{ln\ 2\ \mathcal{N}_A}{m_{76}} (\epsilon_k + \alpha\sigma_{\epsilon_k}) \mathcal{E}_k \Gamma_{1/2} \tag{10.2.2}$$

where $\mathcal{N}_A$ is Avogadro's number, m_{76} is the molar mass of Ge-76, ϵ_k is the efficiency of the partition, $\mathcal{E}_k$ is the exposure of the partition, $\Gamma_{1/2}$ is the half life and α is a nuisance parameter to correlate all of the efficiency uncertainties together.

Put simply this likelihood takes in each event seen in the ROI which each come with a corresponding energy bias and resolution, various efficiencies from e.g. the PSD and containment, and the exposure of the partition. It calculates for each of these the likelihood that that event is a signal ($0\nu\beta\beta$) event and multiplies these together.

A frequentist analysis to constrain $T_{1/2}$ was run using a profile-likelihood two-sided test statistic calibrated using a parameteric bootstrap based on the best fit parameters. The BI of LEGEND-200 extracted from the best fit was $5.3 \pm 2.2 \times 10^{-4}$ counts / (keV kg yr) .

GERDA Phase II and MAJORANA were then included as extra partitions based on their final publications [64, 85]. MAJORANA, GERDA and LEGEND's weight in the combined analysis scales to first order with their exposure of 64.5, 103.7 and 48.2 kg yr respectively. The events in each dataset are shown in Figure 10.3.

No statistically significant signal was found and a 90% CL lower limit of $> 1.9 \times 10^{26}$ yr was set, with LEGEND-200 contributing 30% to the limit's median expectation. The limit is considerably lower than the median value of 2.8×10^{26} yr expected assuming there is no signal (i.e. the median sensitivity). This is primarily due to an event seen close to $Q_{\beta\beta}$ and be attributed by the statistical analysis to a potential signal. However, this results was not intepreted as a hint for a signal by the collaboration, as the combined p-value of 26% for the no-signal hypothesis is still very reasonable As shown in Figure 10.4, this result is now the leading limit in Ge-76, and it strongly motivates the pursuing of the LEGEND program with the expectation of exploring in discovery-mode the entire parameter space allow by the inverted ordering by 2040.

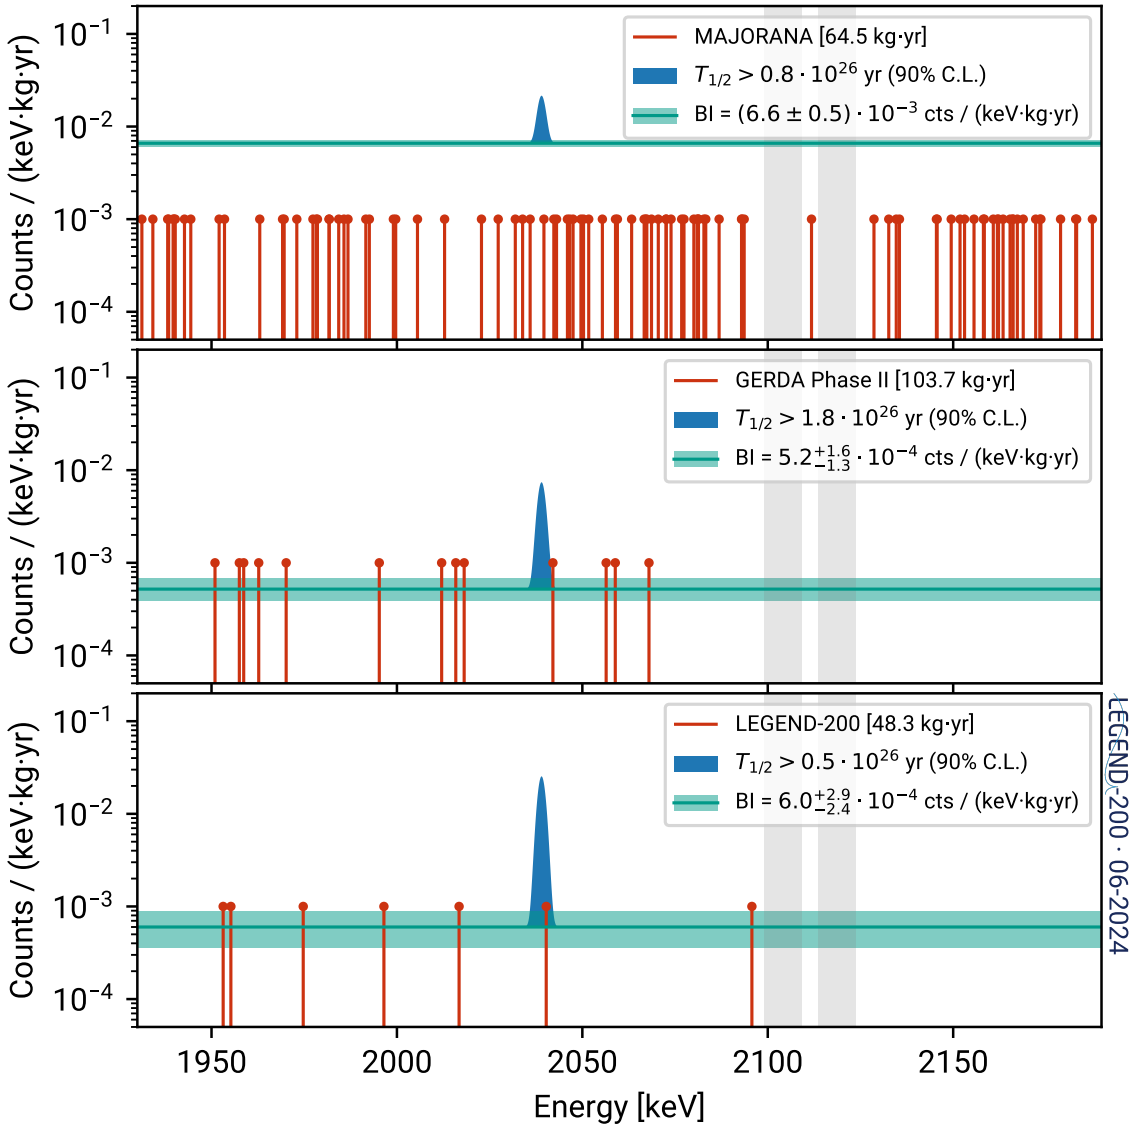

Figure 10.3: Events in unblinded background region after all analysis cuts from GERDA, the MAJORANA Demonstrator and LEGEND-200. Along with the relevant background index and lower limits on the half life.

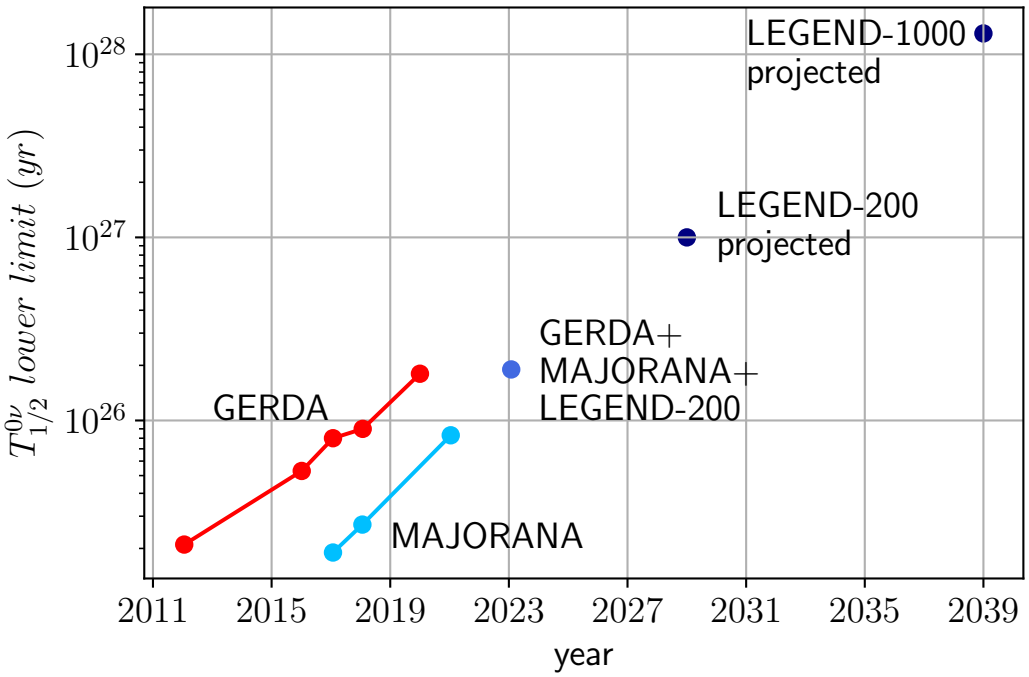

Figure 10.4: The evolution of the Ge-76 half life lower limit over time. The newest result presented here is shown in blue while the projected sensitivity of LEGEND-200 and LEGEND-1000 are shown in navy.

Chapter 11

Summary and Outlook

This thesis focused on the development of the analysis routines analysis framework and automated data processing pipeline for the LEGEND-200 experiment. I have presented these routines and their performance from the low level data processing, to the higher level energy and PSD calibrations and the overall data processing design. These routines were first applied to characterisation data enabling the development of novel detectors and then to LEGEND-200 enabling the first result of the experiment.

In total, 76.2 kg yr of exposure was accumulated and processed with 48.3 kg yr so far used ofr the $0\nu\beta\beta$ search. I showed that both performance of the detectors and the analysis routines were excellent with the energy resolution at $Q_{\beta\beta}$ achieving our goal of 2.5 keV FWHM. The PSD performance was also fantastic with an efficiency of greater than 80% in all partitions. The analysis routines proved to be robust and reliable with all data (over 100 Tb) being processed with no issues on the LNGS server, in particular I would highlight the performance of the innovative calibration routines I developed which ran more than 5000 times without a miscalibration.

On analysing the physics data, it was found that the background was higher than expected from radioassay measurements. An extensive campaign was undertaken to understand the background better. This included my studies of the background homogeneity before and after cuts, and the building of a background model by the collaboration. The conclusion of this work was that the background showed high homogeneity and a single source was ruled out. The current hypothesis is that it is formed of a mixture of sources both near and far from the germanium detectors. Further, background runs were taken with components removed to break degeneracies and rule out components. These runs are still being analysed.

The background after analysis cuts was then presented with the expected high level of suppression at more than 2 orders on magnitude in the ROI. I showed that only five events remained in the ROI after all cuts giving a background index

estimate compatible with the LEGEND-200 goal and with GERDA. These studies gave strong confidence in the analysis methods and led to the collaboration approving the unblinding and first release of the LEGEND-200 flagship analysis results. As the data production lead, I led the unblinding process with 2 further events seen in the ROI after unblinding. One of these events was close to $Q_{\beta\beta}$, in fact only 1.4σ away. More work and data is needed to understand whether it should be considered a likely signal or background event as this is a region of the parameter space which is not yet well understood.

A frequentist analysis was performed by the collaboration after the unblinding on this data. A final background index of $5.3 \pm 2.2 \times 10^{-4}$ counts / (keV ton yr) was obtained for LEGEND-200. The data from LEGEND-200 was then combined with the GERDA and MAJORANA data to produce a combined lower limit on the half life of $> 1.9 \times 10^{26}$ yr at 90% CL. This result is now the leading limit in Ge-76.

LEGEND-200 will be redeployed later this year with a higher mass of working detectors and additional improvements to the array. With these improvements we can expect to obtain an even better energy resolution, an improved LAr anticoincidence performance and an improved sensitivity. LEGEND-200 will be the leading experiment in the search for $0\nu\beta\beta$ for the next few years, with KamLAND-ZEN and CUPID currently undergoing upgrades, as it moves towards the top of the inverted hierarchy.

This will set the stage for LEGEND-1000 later this decade which will use 1 tonne of enriched germanium. With an eventual sensitivity of 10^{28} yr, LEGEND-1000 will be able to probe the entire inverted hierarchy and half of the normal hierarchy. The work outlined here on the routines developed for LEGEND-200 can scale up to fulfil the needs of LEGEND-1000 and already constitutes the baseline design. With a number of other experiments also moving towards the tonne scale, the next decade will be an exciting time for the field of $0\nu\beta\beta$ with the potential of a discovery that will revolutionise the field of particle physics and cosmology as we know it.

Bibliography

[1] A. S. Barabash. “Experiment double beta decay: Historical review of 75 years of research”. In: *Physics of Atomic Nuclei* 74.4 (Apr. 2011), pp. 603–613. issn: 1562-692X. doi: 10.1134/s1063778811030070. url: http://dx.doi.org/10.1134/S1063778811030070.

[2] Matteo Agostini et al. “Toward the discovery of matter creation with neutrinoless double-beta decay”. In: *Reviews of Modern Physics* 95.2 (May 2023). issn: 1539-0756. doi: 10.1103/revmodphys.95.025002. url: http://dx.doi.org/10.1103/RevModPhys.95.025002.

[3] James Chadwick. “Intensitätsverteilung im magnetischen Spectrum der β -Strahlen von radium B+ C”. In: *Verhandl. Dtsc. Phys. Ges.* 16 (1914), p. 383.

[4] Wolfgang Pauli. “Pauli letter collection: letter to Lise Meitner”. Typed copy. url: https://cds.cern.ch/record/83282.

[5] Herbert Pietschmann. *Neutrino: Past, Present and Future*. 2006. arXiv: physics/0603106 [physics.hist-ph]. url: https://arxiv.org/abs/physics/0603106.

[6] E. Fermi. “An attempt of a theory of beta radiation. 1.” In: *Z. Phys.* 88 (1934), pp. 161–177. doi: 10.1007/BF01351864.

[7] E. Fermi and F. Wilson. “Fermi’s Theory of Beta Decay”. In: *American Journal of Physics* 36 (Dec. 1968), p. 1150. doi: 10.1119/1.1974382.

[8] M. Goeppert-Mayer. “Double Beta-Disintegration”. In: *Phys. Rev.* 48 (6 Sept. 1935), pp. 512–516. doi: 10.1103/PhysRev.48.512. url: https://link.aps.org/doi/10.1103/PhysRev.48.512.

[9] Ruben Saakyan. “Two-Neutrino Double-Beta Decay”. In: *Ann. Rev. Nucl. Part. Sci.* 63 (2013), pp. 503–529. doi: 10.1146/annurev-nucl-102711-094904.

[10] Ettore Majorana. “Teoria simmetrica dell’elettrone e del positrone”. In: *Nuovo Cim.* 14 (1937), pp. 171–184. doi: 10.1007/BF02961314.

[11] G. Racah. “On the symmetry of particle and antiparticle”. In: *Nuovo Cim.* 14 (1937), pp. 322–328. doi: 10.1007/BF02961321.

[12] F. Perrin. “Possibilité d’émission de particules neutres de masse intrinsèque nulle dans les radioactivités β”. In: *C. R. Acad. Sci. (Paris)* 197 (1933), pp. 1625–1627.

[13] F. Perrin and D. H. Delphenich. *Possibility of the emission of neutral particles of zero intrinsic mass during β-radioactivity*. https://www.neo-classical-physics.info/uploads/3/4/3/6/34363841/perrin_-_neutrino.pdf. Accessed: 2024-09-09.

[14] Joseph A. Formaggio, André Luiz C. de Gouvêa, and R.G. Hamish Robertson. “Direct measurements of neutrino mass”. In: *Physics Reports* 914 (2021). Direct measurements of neutrino mass, pp. 1–54. issn: 0370-1573. doi: https://doi.org/10.1016/j.physrep.2021.02.002. url: https://www.sciencedirect.com/science/article/pii/S0370157321000636.

[15] E. L. Fireman. “A Measurement of the Half-Life of Double Beta-Decay from $_{50}\mathrm{Sn}^{124}$”. In: *Phys. Rev.* 75 (2 Jan. 1949), pp. 323–324. doi: 10.1103/PhysRev.75.323. url: https://link.aps.org/doi/10.1103/PhysRev.75.323.

[16] E. L. Fireman and D. Schwarzer. “A Re-Investigation of the Double Beta-Decay from Sn^{124}”. In: *Phys. Rev.* 86 (4 May 1952), pp. 451–453. doi: 10.1103/PhysRev.86.451. url: https://link.aps.org/doi/10.1103/PhysRev.86.451.

[17] Mark G. Inghram and John H. Reynolds. “Double Beta-Decay of Te^{130}”. In: *Phys. Rev.* 78 (6 June 1950), pp. 822–823. doi: 10.1103/PhysRev.78.822.2. url: https://link.aps.org/doi/10.1103/PhysRev.78.822.2.

[18] S. R. Elliott, A. A. Hahn, and M. K. Moe. “Direct evidence for two-neutrino double-beta decay in $^{82}\mathrm{Se}$”. In: *Phys. Rev. Lett.* 59 (18 Nov. 1987), pp. 2020–2023. doi: 10.1103/PhysRevLett.59.2020. url: https://link.aps.org/doi/10.1103/PhysRevLett.59.2020.

[19] F. Reines and C. L. Cowan. “Detection of the free neutrino”. In: *Phys. Rev.* 92 (1953), pp. 830–831. doi: 10.1103/PhysRev.92.830.

[20] C. S. Wu et al. “Experimental Test of Parity Conservation in Beta Decay”. In: *Phys. Rev.* 105 (4 Feb. 1957), pp. 1413–1415. doi: 10.1103/PhysRev.105.1413. url: https://link.aps.org/doi/10.1103/PhysRev.105.1413.

[21] M. Goldhaber, L. Grodzins, and A. W. Sunyar. “Helicity of Neutrinos”. In: *Phys. Rev.* 109 (3 Feb. 1958), pp. 1015–1017. doi: 10.1103/PhysRev.109.1015. url: https://link.aps.org/doi/10.1103/PhysRev.109.1015.

[22] B. T. Cleveland et al. “Measurement of the solar electron neutrino flux with the Homestake chlorine detector”. In: *Astrophys. J.* 496 (1998), pp. 505–526. doi: 10.1086/305343.

[23] J. C. Street and E. C. Stevenson. “New Evidence for the Existence of a Particle of Mass Intermediate Between the Proton and Electron”. In: *Physical Review* 52.9 (Nov. 1937), pp. 1003–1004. doi: 10.1103/PhysRev.52.1003.

[24] M. L. Perl et al. “Evidence for anomalous lepton production in e^+e^- annihilation”. In: *PRL* 35.22 (Dec. 1975), pp. 1489–1492. doi: 10.1103/PhysRevLett.35.1489.

[25] G. Danby et al. “Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos”. In: *Phys. Rev. Lett.* 9 (1 July 1962), pp. 36–44. doi: 10.1103/PhysRevLett.9.36. url: https://link.aps.org/doi/10.1103/PhysRevLett.9.36.

[26] K. Kodama et al. “Observation of tau neutrino interactions”. In: *Physics Letters B* 504.3 (Apr. 2001), pp. 218–224. issn: 0370-2693. doi: 10.1016/s0370-2693(01)00307-0. url: http://dx.doi.org/10.1016/S0370-2693(01)00307-0.

[27] B. Pontecorvo. “Mesonium and Antimesonium”. In: *Soviet Journal of Experimental and Theoretical Physics* 6 (Jan. 1958), p. 429.

[28] B. Pontecorvo. “Neutrino Experiments and the Problem of Conservation of Leptonic Charge”. In: *Soviet Journal of Experimental and Theoretical Physics* 26 (May 1968), p. 984.

[29] B. Pontecorvo. “Inverse beta process”. In: *Camb. Monogr. Part. Phys. Nucl. Phys. Cosmol.* 1 (1991), pp. 25–31.

[30] Steven Weinberg. “A Model of Leptons”. In: *Phys. Rev. Lett.* 19 (1967), pp. 1264–1266. doi: 10.1103/PhysRevLett.19.1264.

[31] S. L. Glashow. “Partial Symmetries of Weak Interactions”. In: *Nucl. Phys.* 22 (1961), pp. 579–588. doi: 10.1016/0029-5582(61)90469-2.

[32] Abdus Salam. “Weak and Electromagnetic Interactions”. In: *Conf. Proc. C* 680519 (1968), pp. 367–377. doi: 10.1142/9789812795915_0034.

[33] Y. Fukuda et al. “Evidence for oscillation of atmospheric neutrinos”. In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. doi: 10.1103/PhysRevLett.81.1562. arXiv: hep-ex/9807003.

[34] Q.R. Ahmad et al. “Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by ^8B Solar Neutrinos at the Sudbury Neutrino Observatory”. In: *Phys. Rev. Lett.* 87 (7 July 2001), p. 071301. doi: 10.1103/PhysRevLett.87.071301. url: https://link.aps.org/doi/10.1103/PhysRevLett.87.071301.

[35] Matthew D. Schwartz. *Quantum Field Theory and the Standard Model*. Cambridge University Press, 2013.

[36] Alessandro Strumia and Francesco Vissani. *Neutrino masses and mixings and...* 2010. arXiv: hep-ph/0606054 [hep-ph]. url: https://arxiv.org/abs/hep-ph/0606054.

[37] Osamu Sawada and Akio Sugamoto, eds. *Proceedings: Workshop on the Unified Theories and the Baryon Number in the Universe: Tsukuba, Japan, February 13-14, 1979*. Tsukuba, Japan: Natl.Lab.High Energy Phys., 1979.

[38] Yoichiro Suzuki. "The Super-Kamiokande experiment". In: *The European Physical Journal C* 79.4 (2019), p. 298. doi: 10.1140/epjc/s10052-019-6796-2.

[39] André Schöning, Sebastian Bachmann, and Rohin Narayan. "A Novel Experiment to Search for the Decay μ→ eee". In: *Physics Procedia* 17 (2011). 2nd International Workshop on the Physics of fundamental Symmetries and Interactions - PSI2010, pp. 181–190. issn: 1875-3892. doi: https://doi.org/10.1016/j.phpro.2011.06.035. url: https://www.sciencedirect.com/science/article/pii/S187538921100366X.

[40] F. R. Klinkhamer and N. S. Manton. "A saddle-point solution in the Weinberg-Salam theory". In: *Phys. Rev. D* 30 (10 Nov. 1984), pp. 2212–2220. doi: 10.1103/PhysRevD.30.2212. url: https://link.aps.org/doi/10.1103/PhysRevD.30.2212.

[41] M. Fukugita and T. Yanagida. "Barygenesis without grand unification". In: *Physics Letters B* 174.1 (1986), pp. 45–47.

[42] Andre de Gouvea. "Neutrino Masses and Mixing - Theory". In: *HQL 2008 - 9th International Conference on Heavy Quarks and Leptons* (Mar. 2009).

[43] Peter Minkowski. "μ→eγ at a rate of one out of 109 muon decays?" In: *Physics Letters B* 67.4 (1977), pp. 421–428. doi: https://doi.org/10.1016/0370-2693(77)90435-X.

[44] R. Foot et al. "See-saw neutrino masses induced by a triplet of leptons". In: *Zeitschrift für Physik C Particles and Fields* 44.3 (1989), pp. 441–444. doi: https://doi.org/10.1007/BF01548579.

[45] G.B. Gelmini and M. Roncadelli. In: *Physics Letters B* 99.5 (1981), pp. 411–415. doi: https://doi.org/10.1016/0370-2693(81)90559-1.

[46] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. "Remarks on the Unified Model of Elementary Particles". In: *Progress of Theoretical Physics* 28.5 (Nov. 1962), pp. 870–880. issn: 0033-068X. doi: 10.1143/PTP.28.870. eprint: https://academic.oup.com/ptp/article-pdf/28/5/870/5258750/28-5-870.pdf. url: https://doi.org/10.1143/PTP.28.870.

[47] J W F Valle. “Neutrino physics overview”. In: *Journal of Physics: Conference Series* 53 (Nov. 2006), pp. 473–505. issn: 1742-6596. doi: 10.1088/1742-6596/53/1/031. url: http://dx.doi.org/10.1088/1742-6596/53/1/031.

[48] DUNE Collaboration. *Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report Volume 2: The Physics Program for DUNE at LBNF*. 2016. arXiv: 1512.06148 [physics.ins-det]. url: https://arxiv.org/abs/1512.06148.

[49] Fengpeng. An et al. “Neutrino physics with JUNO”. In: *Journal of Physics G: Nuclear and Particle Physics* 43.3 (Feb. 2016), p. 030401. issn: 1361-6471. doi: 10.1088/0954-3899/43/3/030401. url: http://dx.doi.org/10.1088/0954-3899/43/3/030401.

[50] Ko Abe and Hide-Kazu Tanaka. “Hyper-Kamiokande construction status and prospects”. In: *Frontiers in Physics* 12 (2024). issn: 2296-424X. doi: 10.3389/fphy.2024.1378254. url: https://www.frontiersin.org/journals/physics/articles/10.3389/fphy.2024.1378254.

[51] Andre de Gouvea. *2004 TASI Lectures on Neutrino Physics*. 2004. arXiv: hep-ph/0411274 [hep-ph]. url: https://arxiv.org/abs/hep-ph/0411274.

[52] J. Schechter and J. W. F. Valle. “Neutrino decay and spontaneous violation of lepton number”. In: *Phys. Rev. D* 25 (3 Feb. 1982), pp. 774–783. doi: 10.1103/PhysRevD.25.774. url: https://link.aps.org/doi/10.1103/PhysRevD.25.774.

[53] Boris Kayser. “Are neutrinos their own antiparticles?” In: *J. Phys.: Conf. Ser.* 173 (June 2009), pp. 012, 013. doi: 10.1088/1742-6596/173/1/012013. url: https://doi.org/10.1088/1742-6596/173/1/012013.

[54] S. Dell’Oro et al. “Neutrinoless Double Beta Decay: 2015 Review”. In: *Advances in High Energy Physics* (2016). doi: 10.1155/2016/2162659.

[55] Matteo Agostini, Giovanni Benato, and Jason A. Detwiler. “Discovery probability of next-generation neutrinoless double-β decay experiments”. In: *Phys. Rev. D* 96 (5 Sept. 2017), p. 053001. doi: 10.1103/PhysRevD.96.053001. url: https://link.aps.org/doi/10.1103/PhysRevD.96.053001.

[56] M. Aker et al. *Direct neutrino-mass measurement based on 259 days of KATRIN data*. 2024. arXiv: 2406.13516 [nucl-ex]. url: https://arxiv.org/abs/2406.13516.

[57] A. Ashtari Esfahani et al. “Tritium Beta Spectrum Measurement and Neutrino Mass Limit from Cyclotron Radiation Emission Spectroscopy”. In: *Phys. Rev. Lett.* 131 (10 Sept. 2023), p. 102502. doi: 10.1103/PhysRevLett.131.102502. url: https://link.aps.org/doi/10.1103/PhysRevLett.131.102502.

[58] B. Alpert et al. “HOLMES: The electron capture decay of ^{163}Ho to measure the electron neutrino mass with sub-eV sensitivity”. In: *The European Physical Journal C* 75.3 (Mar. 2015). issn: 1434-6052. doi: 10.1140/epjc/s10052-015-3329-5. url: http://dx.doi.org/10.1140/epjc/s10052-015-3329-5.

[59] DESI Collaboration et al. *DESI 2024 VI: Cosmological Constraints from the Measurements of Baryon Acoustic Oscillations*. 2024. arXiv: 2404.03002 [astro-ph.CO]. url: https://arxiv.org/abs/2404.03002.

[60] Nathaniel Craig et al. *No νs is Good News*. 2024. arXiv: 2405.00836 [astro-ph.CO]. url: https://arxiv.org/abs/2405.00836.

[61] Itamar J. Allali and Alessio Notari. *Neutrino mass bounds from DESI 2024 are relaxed by Planck PR4 and cosmological supernovae*. 2024. arXiv: 2406.14554 [astro-ph.CO]. url: https://arxiv.org/abs/2406.14554.

[62] LEGEND Collaboration. *LEGEND-1000 Preconceptual Design Report*. 2021. arXiv: 2107.11462 [physics.ins-det]. url: https://arxiv.org/abs/2107.11462.

[63] Steven R. Elliott and Petr Vogel. “DOUBLE BETA DECAY”. In: *Annual Review of Nuclear and Particle Science* 52.1 (Dec. 2002), pp. 115–151. issn: 1545-4134. doi: 10.1146/annurev.nucl.52.050102.090641. url: http://dx.doi.org/10.1146/annurev.nucl.52.050102.090641.

[64] M. Agostini et al. “Final Results of GERDA on the Search for Neutrinoless Double-β Decay”. In: *Physical Review Letters* 125.25 (Dec. 2020). issn: 1079-7114. doi: 10.1103/physrevlett.125.252502. url: http://dx.doi.org/10.1103/PhysRevLett.125.252502.

[65] O. Azzolini et al. “Final Result on the Neutrinoless Double Beta Decay of ^{82}Se with CUPID-0”. In: *Physical Review Letters* 129.11 (Sept. 2022). issn: 1079-7114. doi: 10.1103/physrevlett.129.111801. url: http://dx.doi.org/10.1103/PhysRevLett.129.111801.

[66] A. Agrawal et al. *Improved limit on neutrinoless double beta decay of ^{100}Mo from AMoRE-I*. 2024. arXiv: 2407.05618 [nucl-ex]. url: https://arxiv.org/abs/2407.05618.

[67] CUORE Collaboration. “With or without ν? Hunting for the seed of the matter-antimatter asymmetry”. In: (2024). arXiv: 2404.04453 [nucl-ex]. url: https://arxiv.org/abs/2404.04453.

[68] S. Abe et al. *Search for Majorana Neutrinos with the Complete KamLAND-Zen Dataset*. 2024. arXiv: 2406.11438 [hep-ex]. url: https://arxiv.org/abs/2406.11438.

[69] K. Eguchi et al. “First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance”. In: *Physical Review Letters* 90.2 (Jan. 2003). issn: 1079-7114. doi: 10.1103/physrevlett.90.021802. url: http://dx.doi.org/10.1103/PhysRevLett.90.021802.

[70] CUORE Collaboration. “Data-driven background model for the CUORE experiment”. In: (2024). arXiv: 2405.17937 [nucl-ex]. url: https://arxiv.org/abs/2405.17937.

[71] V. Albanese et al. “The SNO+ experiment”. In: *Journal of Instrumentation* 16.08 (Aug. 2021), P08059. issn: 1748-0221. doi: 10.1088/1748-0221/16/08/p08059. url: http://dx.doi.org/10.1088/1748-0221/16/08/P08059.

[72] W. Quinn. “The sensitivity of the NEMO technique to neutrinoless double beta decay and the commissioning of the SuperNEMO demonstrator module.” PhD thesis. University College London, 2023. url: https://discovery.ucl.ac.uk/id/eprint/10169597.

[73] Matteo Agostini et al. “Testing the inverted neutrino mass ordering with neutrinoless double-β decay”. In: *Phys. Rev. C* 104.4 (2021), p. L042501. doi: 10.1103/PhysRevC.104.L042501. arXiv: 2107.09104 [hep-ph].

[74] The CUPID Interest Group. *CUPID pre-CDR*. 2019. arXiv: 1907.09376 [physics.ins-det]. url: https://arxiv.org/abs/1907.09376.

[75] Junpei Shirai and for the KamLAND-Zen Collaboration. “Results and future plans for the KamLAND-Zen experiment”. In: *Journal of Physics: Conference Series* 888.1 (Sept. 2017), p. 012031.

[76] G. Adhikari et al. “nEXO: neutrinoless double beta decay search beyond 1028 year half-life sensitivity”. In: *Journal of Physics G: Nuclear and Particle Physics* 49.1 (Dec. 2021), p. 015104. issn: 1361-6471. doi: 10.1088/1361-6471/ac3631. url: http://dx.doi.org/10.1088/1361-6471/ac3631.

[77] Xun. Chen et al. *PandaX-III: Searching for Neutrinoless Double Beta Decay with High Pressure ^{136}Xe Gas Time Projection Chambers*. 2016. arXiv: 1610.08883 [physics.ins-det]. url: https://arxiv.org/abs/1610.08883.

[78] J. Martín-Albo et al. “Sensitivity of NEXT-100 to neutrinoless double beta decay”. In: *Journal of High Energy Physics* 2016.5 (May 2016). issn: 1029-8479. doi: 10.1007/jhep05(2016)159. url: http://dx.doi.org/10.1007/JHEP05(2016)159.

[79] The LZ Collaboration. *LUX-ZEPLIN (LZ) Conceptual Design Report*. 2015. arXiv: 1509.02910 [physics.ins-det]. url: https://arxiv.org/abs/1509.02910.

[80] J. Aalbers et al. “DARWIN: towards the ultimate dark matter detector”. In: *Journal of Cosmology and Astroparticle Physics* 2016.11 (Nov. 2016), pp. 017–017. issn: 1475-7516. doi: 10.1088/1475-7516/2016/11/017. url: http://dx.doi.org/10.1088/1475-7516/2016/11/017.

[81] J. Aalbers et al. “A next-generation liquid xenon observatory for dark matter and neutrino physics”. In: *Journal of Physics G: Nuclear and Particle Physics* 50.1 (Dec. 2022), p. 013001. issn: 1361-6471. doi: 10.1088/1361-6471/ac841a. url: http://dx.doi.org/10.1088/1361-6471/ac841a.

[82] L. Pertoldi. *0vbb-ge76-history*. https://github.com/gipert/0vbb-ge76-history.

[83] E. Fiorini et al. “A search for lepton non-conservation in double beta decay with a germanium detector”. In: *Physics Letters B* 25.10 (1967), pp. 602–603. issn: 0370-2693. doi: https://doi.org/10.1016/0370-2693(67)90127-X. url: https://www.sciencedirect.com/science/article/pii/037026936790127X.

[84] H. V. Klapdor-Kleingrothaus et al. “EVIDENCE FOR NEUTRINOLESS DOUBLE BETA DECAY”. In: *Modern Physics Letters A* 16.37 (Dec. 2001), pp. 2409–2420. issn: 1793-6632. doi: 10.1142/s0217732301005825. url: http://dx.doi.org/10.1142/S0217732301005825.

[85] I. J. Arnquist et al. “Final Result of the Majorana Demonstrator’s Search for Neutrinoless Double-β Decay in ^{76}Ge”. In: *Phys. Rev. Lett.* 130 (6 Feb. 2023), p. 062501. doi: 10.1103/PhysRevLett.130.062501. url: https://link.aps.org/doi/10.1103/PhysRevLett.130.062501.

[86] Nicola Nosengo. “Gran Sasso: Chamber of physics”. In: *Nature* 485.7399 (2012), pp. 435–438.

[87] G. Pühlhofer et al. *FlashCam: a fully-digital camera for the medium-sized telescopes of the Cherenkov Telescope Array*. 2015. arXiv: 1509.02434 [astro-ph.IM]. url: https://arxiv.org/abs/1509.02434.

[88] M. Agostini et al. “Pulse shape analysis in Gerda Phase II”. In: *The European Physical Journal C* 82.4 (2022), p. 284.

[89] G.F. Knoll. *Radiation Detection and Measurement*. Wiley, 2010. isbn: 9780470131480. url: https://books.google.co.uk/books?id=4vTJ7UDel5IC.

[90] B.G. Streetman and S. Banerjee. *Solid State Electronic Devices*. Always Learning. Pearson, 2015. isbn: 9781292060552. url: https://books.google.co.uk/books?id=oaXuoAEACAAJ.

[91] S. Ramo. "Currents Induced by Electron Motion". In: *Proceedings of the IRE* 27.9 (1939), pp. 584–585. doi: 10.1109/JRPROC.1939.228757.

[92] W. Shockley. "Currents to Conductors Induced by a Moving Point Charge". In: *Journal of Applied Physics* 9.10 (Oct. 1938), pp. 635–636. issn: 0021-8979. doi: 10.1063/1.1710367. eprint: https://pubs.aip.org/aip/jap/article-pdf/9/10/635/18304047/635_1_online.pdf. url: https://doi.org/10.1063/1.1710367.

[93] G. Benato et al. "Improvement of the GERDA Ge Detectors Energy Resolution by an Optimized Digital Signal Processing". In: *Physics Procedia* 61 (2015). 13th International Conference on Topics in Astroparticle and Underground Physics, TAUP 2013, pp. 673–682. issn: 1875-3892. doi: https://doi.org/10.1016/j.phpro.2014.12.069. url: https://www.sciencedirect.com/science/article/pii/S1875389214006828.

[94] Tommaso Comellato, Matteo Agostini, and Stefan Schönert. "Charge-carrier collective motion in germanium detectors for $\beta\beta$-decay searches". In: *The European Physical Journal C* 81.1 (Jan. 2021), p. 76. doi: 10.1140/epjc/s10052-021-08889-0. url: https://doi.org/10.1140/epjc/s10052-021-08889-0.

[95] Dušan Budjáš et al. "Pulse shape discrimination studies with a Broad-Energy Germanium detector for signal identification and background suppression in the GERDA double beta decay experiment". In: *Journal of Instrumentation* 4.10 (Oct. 2009), P10007. doi: 10.1088/1748-0221/4/10/P10007. url: https://dx.doi.org/10.1088/1748-0221/4/10/P10007.

[96] M. Bogovac and C. Csato. "Implementation of a truncated cusp filter for real-time digital pulse processing in nuclear spectrometry". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 694 (2012), pp. 101–106. issn: 0168-9002. doi: https://doi.org/10.1016/j.nima.2012.07.042. url: https://www.sciencedirect.com/science/article/pii/S0168900212008248.

[97] T. Wester. "Characterisation of Coincidence Data of the Gerda Experiment to Search for Double Beta Decays to Excited States". PhD thesis. Technische Universität Dresden, 2020. url: https://nbn-resolving.org/urn:nbn:de:bsz:14-qucosa2-380327.

[98] Tommaso Comellato, Matteo Agostini, and Stefan Schönert. “Charge-carrier collective motion in germanium detectors for $\beta\beta$-decay searches”. In: *Eur. Phys. J. C* 81.1 (2021), p. 76. doi: 10.1140/epjc/s10052-021-08889-0. arXiv: 2007.12910 [physics.ins-det].

[99] The LEGEND Collaboration. *legend-pydataobj*. https://github.com/legend-exp/legend-pydataobj.git.

[100] The LEGEND Collaboration. *legend-dataflow*. https://github.com/legend-exp/legend-dataflow.git.

[101] J. Köster et al. “Sustainable data analysis with Snakemake”. In: *F1000Research* 10.33 (2021). doi: 10.12688/f1000research.29032.2.

[102] The LEGEND Collaboration. *legend-pygama*. https://github.com/legend-exp/pygama.

[103] The LEGEND Collaboration. *legend-daq2lh5*. https://github.com/legend-exp/legend-daq2lh5.git.

[104] The LEGEND Collaboration. *legend-dspeed*. https://github.com/legend-exp/dspeed.git.

[105] M. Agostini et al. “Calibration of the Gerda experiment”. In: *Eur. Phys. J. C* 81.8 (2021), p. 682. doi: 10.1140/epjc/s10052-021-09403-2. arXiv: 2103.13777 [physics.ins-det].

[106] M. Agostini et al. “Modeling of GERDA Phase II data”. In: *Journal of High Energy Physics* 2020.3 (Mar. 2020). issn: 1029-8479. doi: 10.1007/jhep03(2020)139. url: http://dx.doi.org/10.1007/JHEP03(2020)139.

[107] Gary J. Feldman and Robert D. Cousins. “Unified approach to the classical statistical analysis of small signals”. In: *Physical Review D* 57.7 (Apr. 1998), pp. 3873–3889. issn: 1089-4918. doi: 10.1103/physrevd.57.3873. url: http://dx.doi.org/10.1103/PhysRevD.57.3873.

[108] Hans Dembinski, Piti Ongmongkolkul, et al. “scikit-hep/iminuit”. In: (Dec. 2020). doi: 10.5281/zenodo.3949207. url: https://doi.org/10.5281/zenodo.3949207.

List of Figures

List of Tables

Appendix A

Routines

A.1 Data Cleaning

As laid out in Sec. 5.2, the data cleaning is split into 3 types of cuts: the baseline cuts, tail cuts and leading edge cuts. For the baseline and tail cuts we want to check that these are regular and conforming to the expected shape of flat for the baseline and exponential for the tail.

Starting with the baseline cuts, two parameters are used to check for the flatness of the tail. The first is the slope of a linear fit to the baseline and the second is the rms of the baseline relative to this linear fit. In calibrations, the slope cut will remove waveforms on the tail of previous events and the rms cut will remove pileup of low energy events on the baseline. For both these cuts the peak of the distribution is fit with a Gaussian and the mode and sigma extracted, a double-sided cut is then applied to remove events 4σoutside of the peak as shown in Figure A.1.

The tail is simpler with only a single cut. However, as the tail is not a single exponential but a mixture of exponentials of different time constants finding a parameter that was independent of the energy was challenging. In the end the procedure found was to log the decay tail and then fit with a quartic polynomial, the rms was then calculated relative to this fit after exponentiating to return to the original space (otherwise the energy information is stored in the level of the noise). Similar to the baseline cuts a double-sided cut was applied to remove events 4σoutside of the peak. In calibrations this cut removed events with energy on the decay tail. A quartic polynomial was chosen to give plenty of freedom in the fit to the tail and so the cut was not too tight as it is possible to have a delayed charge collection in the tail which could affect the shape.

Finally, the leading edge cuts. These are split into 4 categories, depending on if the waveform is : empty, low energy, higher energy and crosstalk. For the empty waveforms we perform our standard procedure of fitting the distribution of in this case the maximum and minimum of our energy filter and performing a 4 sigma cut.

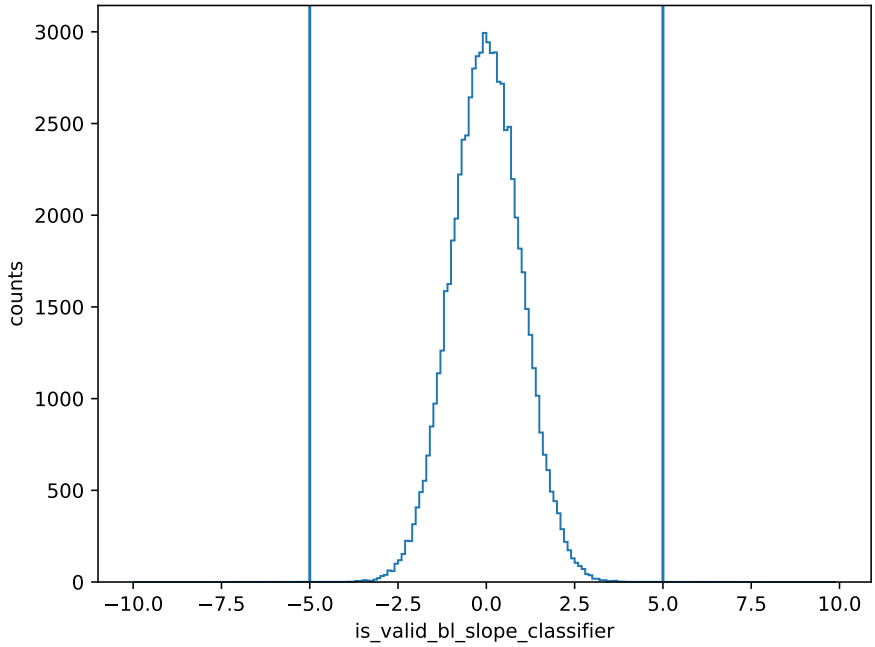

Figure A.1: Example of the 4 sigma double-sided cut for the baseline slope on 1 detector.

At low energies our time resolution is poor due to the low signal-to-noise of the signals. We therefore need to apply smoothing to the waveform to get accurate values. A simple trap filter is applied and the time of the max of this filter found, the waveform is cut if this falls outside the expected region. The same check is done for the crosstalk but with the minimum of the filter instead. Finally, for higher energy waveforms the same check is applied but in addition we can use more timing information. Checks are done on the t0 (start point of signal), rise time and the time of the maximum.

A.2 Energy

A.2.1 Pole Zero Correction

As already described in LEGEND-200, we do not have a single pole zero exponential decay but instead an admixture of two or more exponentials. In the analysis currently we only correct for a single long time constant although, work is ongoing to include a full treatment of the pole zero correction.

A routine was developed for deriving the pole zero constants for each detector. This was done by taking a few thousand waveforms, baseline subtracting and then apply a linear fit to the logged decay tail. This gives a distribution of slopes which should be Gaussian with some tail due to pileup events. The top of this Gaussian is fit and the mode extracted. The pole zero constant is then simply $\tau = -1/s$ where s is the slope.

A.2.2 Peak Fitting

A peak fitting routine was developed to give the best performance in terms the least number of failed fits. This was implemented as an extended maximum likelihood fit using the iminuit package [108]. The basic peak shape to fit was a Gaussian on a step background.

$$g(E) = \frac{n}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(E-\mu)^2}{2\sigma^2}\right) \tag{A.2.1}$$

$$f_{step} = \frac{a}{2}\text{erfc}\left(\frac{E-\mu}{\sqrt{2}\sigma}\right) \tag{A.2.2}$$

For high statistics peaks a low energy tail is also included. If the tail was used then the fit was performed in two stages. First the Gaussian and step were fit, and then the tail was added, and the fit performed again. The fit with and without the tail was compared and the best fit chosen based on the higher p-value. Additionally, if the error on the tail fraction was greater than the tail fraction itself then the fit without the tail was chosen.

$$f_{tail} = \frac{b}{2c}\exp(\frac{E-\mu}{c} + \frac{\sigma^2}{2c^2})\text{erfc}\left(\frac{E-\mu}{\sqrt{2}\sigma} + \frac{\sigma}{\sqrt{2}c}\right) \tag{A.2.3}$$

A.2.3 Energy Optimisation

The energy optimisation was performed for all the 3 possible energy filters in the analysis: the trap filter, cusp filter and zac filter. Each of these have 2 parameters: the flat top time and a shaping parameter. The flat top time is essentially a delay in the filter where the samples are not used in the energy calculation. It should be longer than the leading edge section of the signal to avoid ballistic deficit effects.

For LEGEND, the flat top time was chosen for all three filters by taking events in the 2.6 MeV gamma line. A full drift time estimate ($t_{100} - t_0$) was calculated for all events in the peak and the 99th percentile of this distribution found to remove any extreme outliers. The flat time was then taken to be this value times a factor of 1.5 as a buffer. This method was found to be simple and robust. Once this was done the shaping parameter could be optimised for each filter. A selection of events in the 2.6 MeV gamma line were taken and an initial set of filter values were tested. For each filter value the energy was calculated and then different charge trapping values (alphas) were swept through implemented using:

$$E_{CTC} = E(1 + \alpha dt_{eff}) \tag{A.2.4}$$

For each value of alpha, a fit was performed as above with the FWHM and the maximum extracted. This gave a distribution of the FWHM/max versus alpha

which was then minimised to find the optimal value of alpha. Finally, the full width at 20% max was returned for that value of the shaping parameter. This was chosen to minimise the tail of the distribution as it is more sensitive than the FWHM After the initial values were calculated a full Bayesian optimisation was performed to find the global minimum. This gave an optimal shaping parameter and charge trapping value for each filter.

A.2.4 Energy Calibration

As described in Sec. 5.3 the energy calibration was performed in two steps. For the run to run calibration, the 2.6 MeV peak was first found by binning the 95^{th} to 99.9^{th} percentile of the energy spectrum and taking the maximum as above the FEP very few events are expected in single calibrations. The peak fitting algorithm above was then run on the 2.6 MeV peak to extract the mode of this peak and then the energy spectrum was scaled using:

$$E_{cal1} = \frac{mode}{2614.511}.E_{uncal} \tag{A.2.5}$$

.

For the second step the calibrations were then grouped together and the peak fitting routine above was run for each expected peak. In total 26 peaks were searched for from the low energy gamma at 277 keV to the high energy summation line at 3475 keV. A couple of prominent peaks such as the 511 electron annihilation line could not be used as these are double lines. In the future however with an improved treatment these could be included.

For each fit the output was tested for validity with the following criteria:

- iminuit status was valid
- the p-value of the fit was greater than $10 * * - 20$
- the covariance matrix was finite and non-zero
- the errors were finite, non-zero and the fractional error was greater than $10 * * - 7$

to ensure only good fits were used in the calibration. Additionally, we required that the number of events in the peak was greater than 100 to remove cases where the peak was not present and instead a statistical fluctuation in the background was fit. The calibration parameters were then found by fitting the peak position (the mode of the fit) versus the known energy of the peak. Additionally, the FWHM was also fit from these fits as shown in Figure A.2. The DEP and SEP were first excluded as these lines are Doppler broadened then both a quadratic and linear fit was performed for the FWHM and the difference between the two used as a

Energy	Isotope	Peak shape
238.632	Pb-212	excluded double line
240.986	Ra-224	excluded double line
277.371	Tl-208	gauss on linear
288.2	Bi-212	gauss on linear
300.087	Pb-212	gauss on linear
452.98	Bi-212	gauss on linear
510.77	Tl-208	excluded double line
511	e^+e^-	excluded double line
549.73	Rn-220	gauss on linear
583.187	Tl-208	gauss with tail on step
727.330	Bi-212	gauss with tail on step
763.13	Tl-208	gauss on linear
785.37	Bi-212	gauss on linear
860.557	Tl-208	gauss with tail on step
893.408	Bi-212	gauss on linear
927.6	Tl-208	gauss on linear
952.120	Bi-212	gauss on linear
982.7	Tl-208	gauss on linear
1078.62	Bi-212	gauss on linear
1093.9	Tl-208	gauss on linear
1512.7	Bi-212	gauss on linear
1592.511	Tl-208 DEP	gauss with tail on step
1620.50	Bi-212	gauss with tail on step
1679.7	Bi-212	gauss on linear
1806.0	Bi-212	gauss on linear
2103.511	Tl-208 SEP	gauss with tail on step
2614.511	Tl-208	gauss with tail on step
3125.511	summation	gauss on linear
3197.7	summation	gauss on linear
3475.1	summation	gauss on linear

Table A.1: Table of peaks used in partition level calibration with the isotope and peak shape used.

systematic on the energy scale. The quadratic term is generally only needed for modelling incomplete charge collection effects such as charge trapping. If our charge trapping correction is working well then this term should be small.

$$\text{FWHM}(E) = \sqrt{a + bE + cE^2} \tag{A.2.6}$$

$$\text{FWHM}(E) = \sqrt{a + bE} \tag{A.2.7}$$

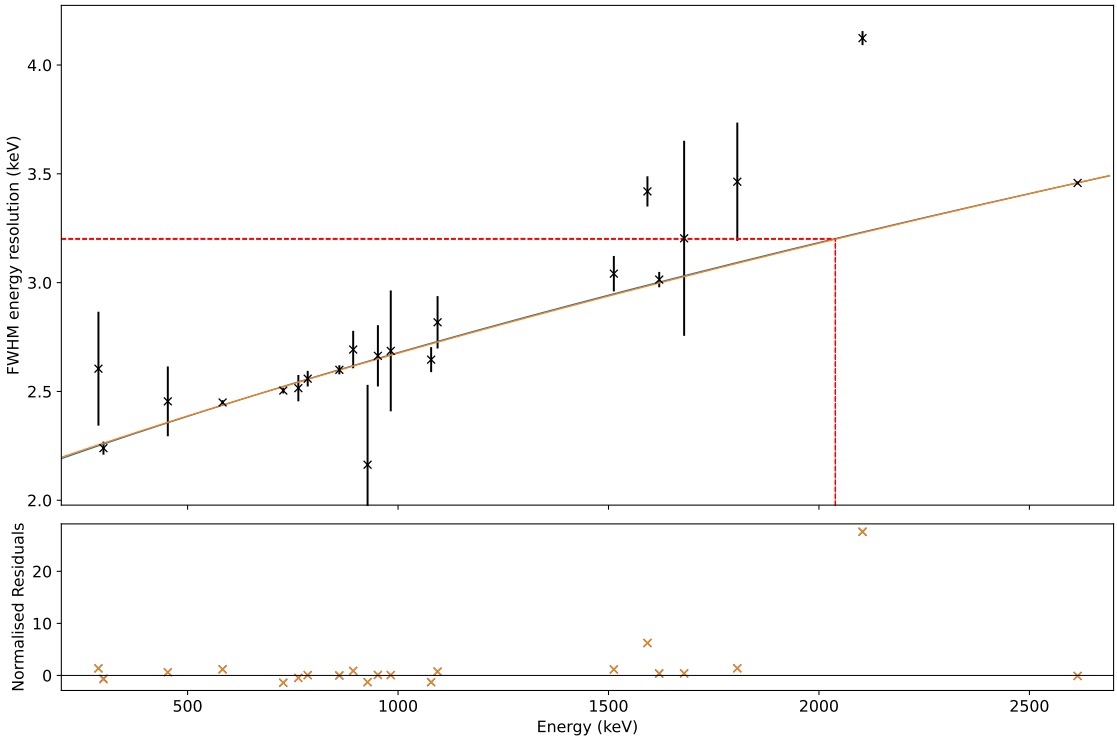

Figure A.2: FWHM fit for 1 detector. In blue is the linear fit and in orange the quadratic. In this case they are very close together implying the quadratic term is not necessary.

A.3 PSD

A.3.1 Current Filtering

The current is determined using the windowed waveform around the leading edge. After pole zero correction the waveform is differentiated, it is then upsampled from 16 ns samples to 1ns simply by assigning the same value to all the 16 samples in the previous sample. Finally, we apply 3 moving windows of 96 ns to smooth this pulse. The maximum of this pulse is then taken.

This method is used for all Mirion detectors: ICPC and BEGe. For the PPCs the current is determined using a slightly different parameter which was found to work better for these detectors. The current is determined using a moving slope filter. This sweeps through the waveform performing linear fits in a 200 ns window. The current parameter is then taken as the maximum of this.

A.3.2 Calibration

In the Compton bands the A/E distribution is characterised by a Gaussian single site band with a tail of multi-site events to lower A/E values. We therefore fit this distribution with a Gaussian signal and an exgaussian tail background A.2.3. This is done in a staged manner fitting first the Gaussian by itself in the signal

region, then the tail in the background region and then both together. The fit is only performed from -15 standard deviations of the signal peak to +5 standard deviations of the peak as in this region the background is approximately linear. All these fits are done as unbinnned extended maximum likelihood fits using the iminuit package [108].

To monitor the time stability of the A/E the low energy region from 1 to 1.3 MeV is fit with the centroid and standard deviation of the Gaussian extracted. Runs where is centroid deviates by less than 0.4 standard deviations can be grouped together for the final calibration. In the future we would like to switch to a run by run correction as in the energy calibration by this requires further improvements to the stability of the fitter.

Once this is done the runs are grouped together into a single spectrum. For the ICPCs the next correction is the drift time correction. This is done using DEP events as these have a clear 2 blob structure due to the topology of these events. The procedure is shown in Figure A.3. First the drift time spectrum is fit to extract these two areas and then each is fit in A/E. The drift time correction is then applied to the A/E using:

$$A/E_{CTC} = A/E(1 + \alpha dt_{eff}) \tag{A.3.1}$$

.

The final correction is the energy correction. The Compton region from 900 keV to 2.3 MeV is sliced into 20 keV bands with the gamma lines excluded. The A/E is then fit in each band and the centroid and standard deviation extracted. These are both then fit versus energy, as shown in Figure A.4, and a classifier defined using the equations below. The validity of this approach can be seen in Figure A.5 where after correction the Compton bands are more uniform.

$$\mu(A/E) = a + bE \tag{A.3.2}$$

$$\sigma(A/E) = \sqrt{a + b.E^c} \tag{A.3.3}$$

$$A/E_{classifier} = \frac{A/E/\mu(A/E) - 1}{\sigma(A/E)} \tag{A.3.4}$$

Finally, the cut value is calculated using DEP events. A sweep through of possible cut values from -8 to 0 is performed and for each the survival fraction is calculated. This is extracted by fitting both the events surviving and failing each cut value in energy to get the number of events in the DEP peak. We can then simply interpolate using this distribution to 90 % survival fraction to get the cut value. The survival fractions for the other peaks are also calculated to track performance and this is shown for different cut values in Figure A.6.

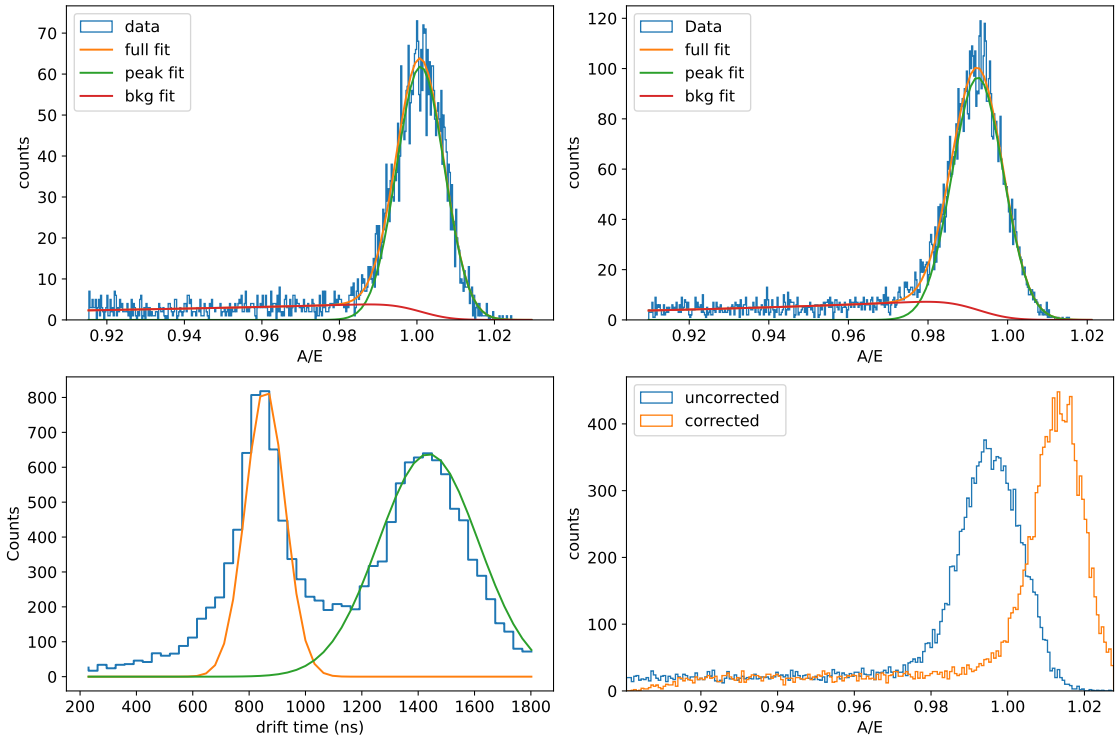

Figure A.3: Example of drift time correction. The detector is first fit in the drift time space using 2 peak top fits as in the bottom left panel. Each peak top is then fit in A/E as shown in the top panels and the correction calculated. The effect of the correction is shown in the bottom right panel.

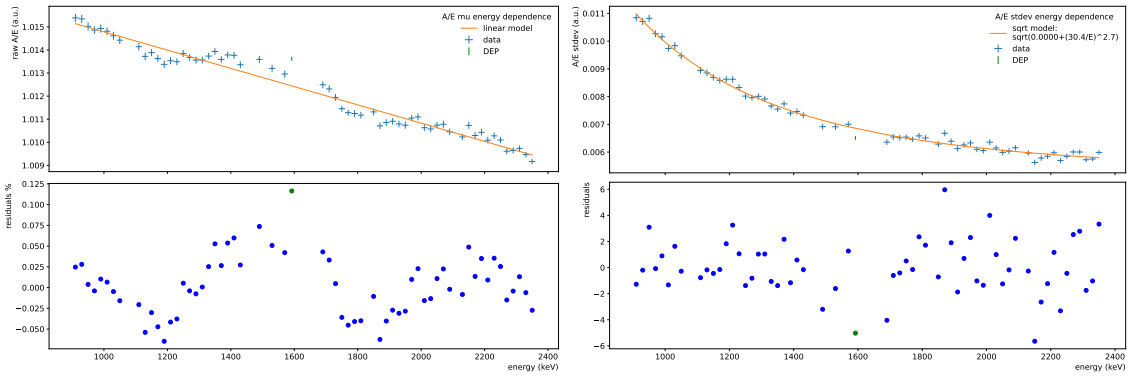

(a) Example of fit to A/E means with energy. (b) Example of fit to A/E sigmas with energy.

Figure A.4: Fits for the A/E energy correction.

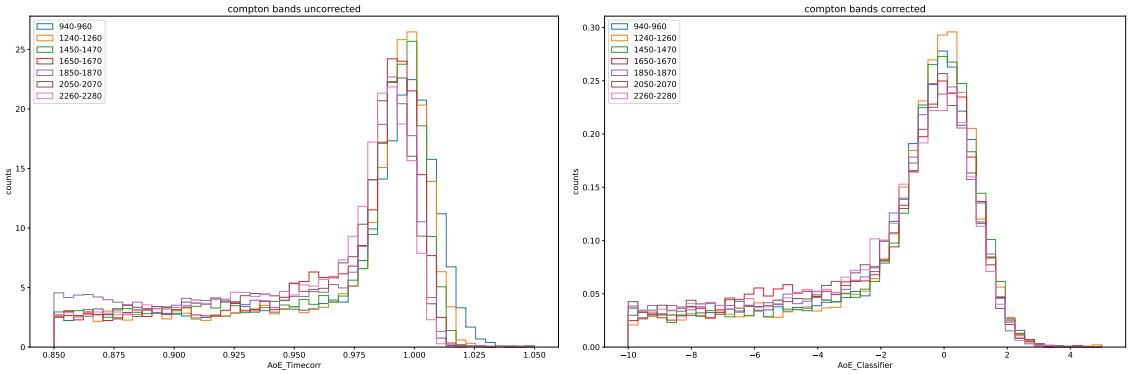

Figure A.5: Compton bands before and after A/E energy correction.

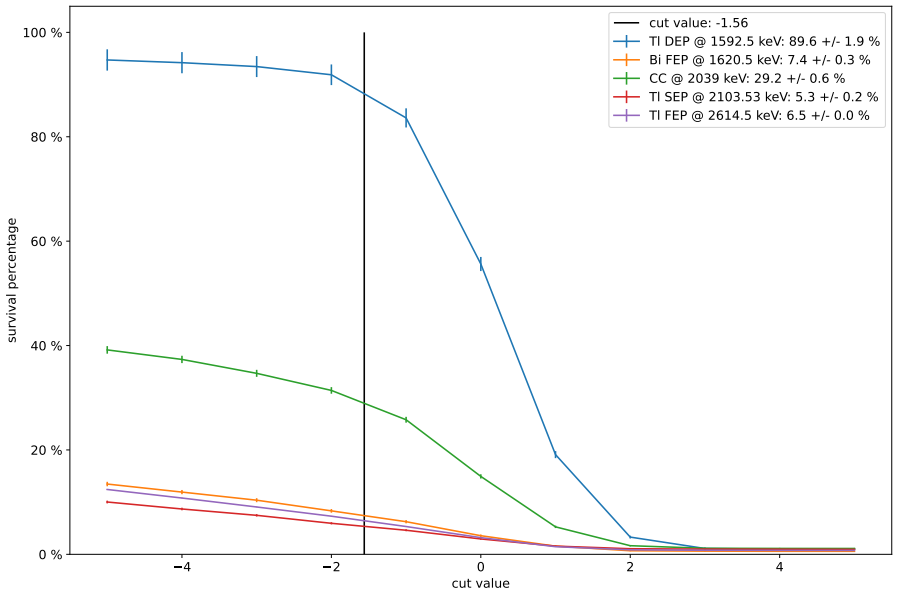

Figure A.6: Survival fractions for varying values of the A/E cut.