Development of the ionisation signal analysis and the thermometry system for the LUX-ZEPLIN direct Dark Matter detection experiment

Andrew Stevens
St. Edmund Hall, Oxford

Thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Oxford

Michaelmas Term, 2021

Abstract

There is substantiated astrophysical evidence for a substance known as Dark Matter within the universe, at both galactic and cosmological scales. Despite this evidence, the true nature of Dark Matter remains one of the most profound questions in modern physics. One viable candidate is the Weakly Interacting Massive Particle (WIMP), which could nicely explain the observed measurements. The LUX-ZEPLIN (LZ) detector is searching for WIMP interactions with xenon nuclei, through both scintillation and ionisation signals. It comprises a time-projection chamber filled with 7 tonnes of liquid xenon, situated at the Sanford Underground Research Facility in South Dakota.

The experiment operates at low temperatures (about 175 Kelvin). During the cool-down of the experiment from room temperature and to ensure stability during running, the temperature needs to be monitored closely to ensure the detector is not damaged and that high-quality data can be collected. The unique implementation of the thermometry system within LZ necessitates bespoke solutions to overcome challenges related to the radiopurity and cleanliness. This thesis examines how these problems are resolved through the careful material procurement and development of a successful calibration procedure.

The ionisation signal plays an important role in the reconstruction of event energies and in interaction type discrimination. However, detector-specific phenomena cause an inherent position dependence. The origins of this position dependence are examined, along with the development of a method to correct signals using dispersed calibration sources. A consideration of the detector contraction is made along with a calibration of the level sensors. Together the position dependence in the ionisation signal, along with these level sensors enables the detector to be precisely levelled.

This work concludes by examining techniques to improve the sensitivity of the experiment below nominal capabilities. Specifically, this includes an examination of realistic pulse shape characteristics to improve the energy threshold of the experiment. Crucially, this provides a series of techniques which can be utilised within LZ analyses.

Statement of Originality

I confirm that the content of this thesis is my own work. Contributions from others are referenced in each specific section, however more general contributons are discussed below.

In Chapters 1-3, a review of literature is performed to motivate the later sections.

In Chapter 4, the thermometry system for the detector is included. The thermometry system was designed by previous researchers at Oxford, however the calibration design and implementation is my own work. The electronics for the thermometers have been designed by my supervisor, Prof. Hans Kraus, however the characterisation is my own work. I was heavily involved in the construction of these sensors and commissioning of the system on-site.

In Chapter 5, several elements are separately addressed relating to the position dependence of the ionisation signal. The position sensors and hexapod were designed and constructed by Dr. Kathryn Boast, a previous Oxford graduate student. The level sensors were designed and constructed by Dr. Feng-Ting Liao, another previous Oxford graduate student. The analyses conducted with these sensors are my own work. The detector levelling was conducted with several collaborators who performed some elements of the detector levelling, both remotely and on-site. The specific calibrations of the level sensors and the S2 response is my own work. The corrections to the S2 signal are my own work.

In Chapter 6, the basis for an ionisation only search is conducted. The original concept for this work was included in a collaboration paper, written by Dr. Brais López-Paredes (Imperial College) on behalf of the LZ collaboration. Results in this paper were cross-checked with the simulations I have conducted during its inception. The work I have produced goes beyond the work published. More realistic effects are included, which critically prepares the experiment for the upcoming analysis to low energy signals. These sections are my own work.

Acknowledgements

The past four years, studying in Oxford, have provided a fantastic opportunity for me to develop various aspects of my life. I've met some amazing people during the process who helped me through the highs and lows of the experience and it seems right to thank them here.

Firstly, I would like to thank my supervisor, Professor Hans Kraus for enabling me to study at Oxford. I am so grateful to him for all the support and advice he has given in different aspects of my work, through discussions and comments he has always been willing to provide. It was also such an enjoyable experience to be able to explore the Black Hills of South Dakota with him.

I would also like to thank my parents, Charles and Linda and my sister Julia for their continued support over the past few years, including the time when I was working from home in Cumbria. I especially appreciate their help as I coped with the logistics of various house moves. Additionally, I would like to thank my grandparents, who were also so supportive of me. Sadly they were unable to complete this exciting journey, but I am sure they would have been delighted to see the hard work come together.

The LZ group in Oxford have also provided invaluable help whilst I have been completing this work. I would like to make a special mention to Amy C., who has been regularly available to discuss elements of work over the past few years; the many hours of time, committed by her has helped me immensely, and I am especially grateful for this. Thanks must also go to the other members of the Oxford Group: Kim P., Feng-Ting L., Theresa F., Cees C., Eilish G., Aiham A., Niamh F. and Dan H.. Thank you for enhancing the Oxford experience, for your company and your inspiring questions when presenting work. Countless hours were spent in the cleanroom with many of you, cleaning and constructing

sensors. So, thank you for helping the time pass by more quickly and for getting through the difficulties together! A mention should go to those in the workshop who were able to quickly help with some hardware components of my work.

I would also like to thank members of the wider LZ community who have similarly enabled me to engage in discussions and for helping me to feel so welcome in the collaboration. I enjoyed many weeks on site in South Dakota and was fortunate enough to get to know a significant number of you. There are too many people to mention everyone by name (and would probably be foolish since I would miss someone out), but I am sure you know who you are. Nevertheless, there are a few that deserve particular mention. I am grateful to Andrew N. for his advice with both technical questions and for being there at different times, especially when things were tough. Many thanks also go to Will T. for his help with the ALPACA development, but also for assisting with other problems and for being a friendly face to speak with. I would also like to thank Brais L.P. who supported the work I was doing when he ran the LZ S2-only group for a few years and often enabled me to see problems from a different angle. I also thank Ishan K. who has also been helpful with many things. I am grateful to all four of you, it has been a pleasure getting to know you all.

I would like to thank various friends to whom I am grateful for their continuing support, both in Oxford and back home. Thank you to Josh T., Connor H., Liam S., Lilli H., Martin F., Rassin L., John C. and many more. I wish you all success with your endeavours. The Teddy Hall MCR always felt homely and enabled me to meet so many amazing people studying impressive things. I will remember my time there fondly. Thank you also goes to Liz B., Alex P., Migle S., and Aaron O'N. for the times we spent together as housemates on James Street in Oxford; I enjoyed getting to know you all.

I could not have completed this DPhil without you all, so thank you again!

The author with the inner most part of the detector - the TPC field cage and PMT arrays. Credit to Nick Hubbard Sanford Underground Research Facility, July 2019.

Contents

List of Figures

List of Tables

Introduction

There is currently evidence on different astrophysical scales for the presence of a substance, known as Dark Matter, which makes up about a quarter of the contents of the universe. This evidence suggests that Dark Matter constitutes a larger fraction of the energy density than the particles making up the Standard Model of Particle Physics. A leading candidate is the Weakly Interacting Massive Particle (WIMP), which is a general class of particles that interact at the weak scale and have masses in the GeV to TeV range.

The LUX-ZEPLIN (LZ) experiment is searching for WIMP interactions with xenon nuclei. It is situated 4850 feet below the surface at the Sanford Underground Research Facility in South Dakota. At the heart of the detector is a dual-phase container filled with 7 tonnes of xenon, operated at about 176 K. Dark Matter scattering from xenon nuclei in the liquid would lead to the production of both primary scintillation light and ionisation electrons. These electrons are drifted and extracted into the gas region under electric fields and produce a secondary scintillation signal. This arrangement is called a Time Projection Chamber (TPC). The primary (S1) and secondary (S2) scintillation signals are detected by photomultiplier tubes (PMTs) which provide information about the interaction location and the incident particle type. Understanding the S2 signal is a fundamental theme running throughout this thesis, culminating in the use of pulse shaping characteristics to improve the sensitivity to Dark Matter below the nominal energy threshold.

In Chapter 1 the evidence for Dark Matter is outlined and several potential candidates are discussed.

In Chapter 2 Dark Matter detection techniques are introduced. A specific focus is placed upon direct detection, the interaction of Dark Matter with Standard Model particles, utilised within LZ. Some other experiments are also considered.

In Chapter 3 interactions within the xenon and how these propagate through to signals within LZ are discussed. The components of the detector are examined individually, as well as the background sources to Dark Matter searches. Consideration is given to the different ways of calibrating the detector as well as the experimental strategy.

In Chapter 4 the thermometry system is introduced as a significant hardware contribution. This bespoke instrumentation informs operators of the temperature during cool-down and stable running. This is critical to conduct a Dark Matter search. The thermometers have associated precision and accuracy requirements to perform such a role, which are shown to be met. The readout electronics are shown to perform well and are utilised within the devised thermometer calibration scheme.

In Chapter 5 the position dependence of the S2 signal is considered. The cool-down of the detector is shown to cause a contraction that is consistent with one that is uniform. Additionally, the level sensors are calibrated and show good sensitivity (about 20 µm) to changes in the liquid level and so can be used to identify regions where time dependence in the S2 signal may be present. These level sensors are also used in the levelling of the detector, alongside a mapping of the S2 pulse width as a good proxy for the liquid to anode distance. This Chapter concludes by considering the correction of the S2 signal from xenon impurities and light collection effects, which is essential to apply to conduct a search for Dark Matter.

In Chapter 6 the S2 signal is considered further. At lower interaction energies some events have an S2 but no associated S1 signal. These are so-called S2-only events. The loss of information the paired S1 signal normally brings, leads to an increase in the background, a problem for such a Dark Matter search. This Chapter looks at regaining some sensitivity by accounting for the extra background. The pulse shape contains information about the depth of the interaction through the diffusion of the electrons when drifting. This Chapter

shows this can be used to improve the sensitivity of LZ below the nominal searches.

Chapter 1

Experimental evidence and candidates for Dark Matter

This Chapter discusses the experimental evidence for the existence of Dark Matter in the universe. Multiple sources of information suggest that such a substance is about 5.4 times more abundant than the luminous matter that makes up visible structures [1]. This evidence spans multiple cosmological scales, from considering the universe as a single system, through galactic clusters and down to individual galaxies. Various candidates exist for Dark Matter, having various theoretical motivations and are addressed here.

1.1 Astrophysical evidence

1.1.1 Early evidence

The first evidence for Dark Matter dates from the 1930s when Lundmark and Zwicky suggested the presence of non-luminous matter to explain an unexpected measurement of galactic dynamics in the Coma Cluster [2]. A discrepancy was observed in the velocity distribution and the inferred mass of the galaxy from one expected using the virial theorem. This states mathematically that,

$$\langle V \rangle = -2\langle K \rangle, \tag{1.1}$$

where $\langle V \rangle$ and $\langle K \rangle$ are the time averaged gravitational potential and kinetic energy respectively of galaxies within the cluster. It was shown that the mass required needs to be significantly larger than what is observed to explain the velocity distribution.

Almost 40 years later, works by Rubin and Ford observed similar disparities in mass distributions by measuring galactic rotation curves in the Andromeda nebula [4]. Simple

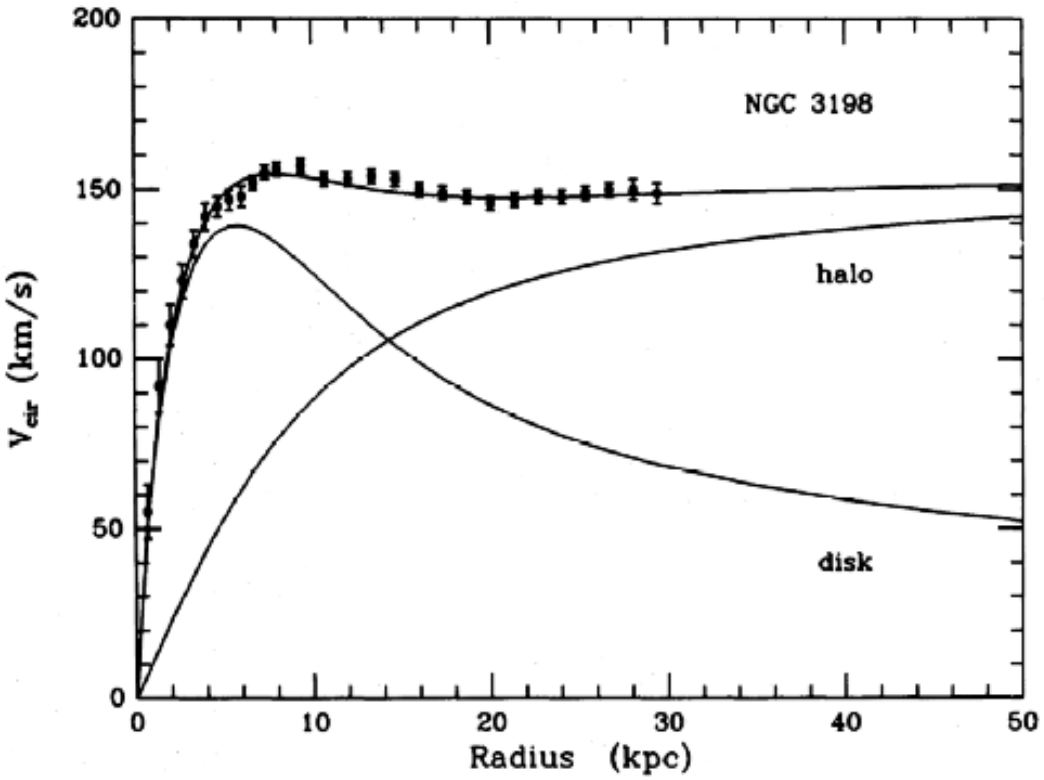

Figure 1.1: Galactic Rotation Curve for the NGC 3198 galaxy inferred from the experimental data points. Contributions to the circular velocity from the Dark Matter halo and the visible Galactic disk are also shown. © AAS. Reproduced with permission. Figure from [3].

considerations using Newtonian mechanics and gravity leads to the velocity profile being given as

$$v(r) = \sqrt{\frac{GM(r)}{r}}. \tag{1.2}$$

Here, G is the gravitational constant, r is the radius from the galactic centre and $M(r)$ is the mass enclosed within that radius. Based on the luminous matter seen in the galaxy which is concentrated towards the centre, the velocity distribution of stars should fall as $1/\sqrt{r}$ outside of the central region containing the majority of the mass. However, the observed distributions do not agree with this formulation. As Figure 1.1 illustrates for the NGC-3195 galaxy, it is observed that velocities of stars became approximately constant at larger radii, and independent of the radial distance to the centre of the galaxy. This suggests the presence of some non-luminous material or 'Dark Matter' which has a density distribution that is given by $\rho(r) \propto 1/r^2$ at larger radii. This is considered in more detail in Section 2.1.3.

Other theories could explain the rotation curve behaviour, without the need for Dark Matter. It is possible to modify General Relativity slightly and recover the velocity distributions, without the need for new particles. However, for the theory to be universal, it must explain all the observed phenomena [5]. Some evidence is more difficult to reconcile with modifications to gravitational theory, as are discussed later within this Chapter. An early modification is Modified Newtonian Dymanics (MOND), as was first proposed by Milgrom [6]. Within this formulation, Newton's Second Law of Motion ($F = ma$) has to be modified as

$$F = ma\,\mu\left(\frac{a}{a_0}\right). \tag{1.3}$$

The function $\mu\left(\frac{a}{a_0}\right)$ asymptotes to 1 for $a_0 << a$ and to a/a_0 for $a << a_0$, where a_0 is a constant. This ensures that the standard Newtonian physics is recoverable in the former case. In the latter case, the modified law becomes

$$F = \frac{ma^2}{a_0}. \tag{1.4}$$

Running through the same mathematics that leads to Equation 1.2, gives

$$v^4 = GMa_0. \tag{1.5}$$

Crucially the velocity is constant and independent of galactic radius. Therefore, in agreement with the rotation curve observations. This theory was further developed in the relativistic framework and is known as TeVes and is one of many modified gravity theories [7]. The modifications to General Relativity often add new fields to the existing model to reproduce the rotation curve behaviour [8].

1.1.2 Gravitational Lensing

One further piece of evidence for Dark Matter comes in the form of Gravitational Lensing. This is a phenomenon predicted by General Relativity and observed by astronomers, where light will bend around massive bodies [9]. It is known as gravitational lensing because it behaves in a similar method as an optical lens, where the focal distance is dependent on the mass content [10]. Distant objects can be lensed by closer ones and can therefore place bounds on that object's content. The observation of the lensing can be compared with the expected mass of the closer object from electromagnetic radiation. Similarly with previous measurements, disparities occur between ones involving gravity and those based in the electromagnetic spectrum of stars [11]. Hence, implying the presence of matter that is non-luminous.

An example of such a phenomenon is seen in the Bullet Cluster. This is a system whereby two galactic clusters collided in the 'recent' past. As Figure 1.2 shows, there is a disparity in the spatial arrangement of gas from the X-ray measurements to the mass distribution from gravitational lensing. Theoretically this can be explained by Dark Matter. As it is significantly less interacting, it can pass virtually unimpeded, whereas the gas interacted more readily and thus does not pass through as easily, leading to the separation of the two populations [12]. The stars making up the galaxies are considerably dispersed, so their distribution is similar to that of the Dark Matter. This evidence is hard to resolve with a simple modified gravity theory, such as the one discussed above. However, there have been attempts to try to explain this without some of the Dark Matter [13]. Despite this, the addition of particle Dark Matter does prove a rather elegant solution, without the need for tuning gravitational models.

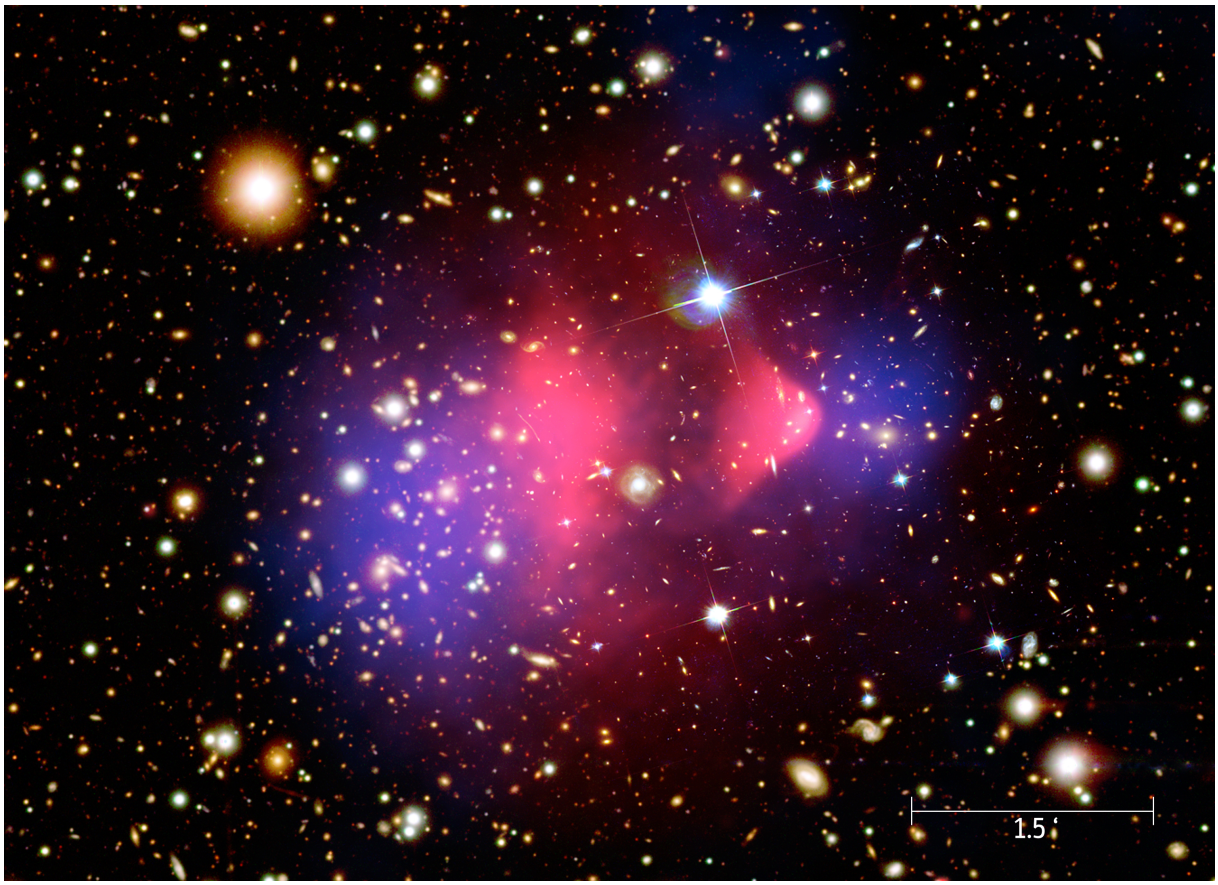

Figure 1.2: Composite image of the Bullet Cluster (1E 0657-56). The base image is from Hubble, x-ray emission based on data from the Chandra Telescope is shown in pink and the blue overlay shows the gravitational lensing inferred mass distribution. Images are based on X-ray measurements as well as lensing. Image from [14]. Credit goes to X-ray: NASA/CXC/CfA/M.Markevitch et al.; Optical: NASA/STScI; Magellan/U.Arizona/D.Clowe et al.; Lensing Map: NASA/STScI; ESO WFI; Magellan/U.Arizona/D.Clowe et al.

1.1.3 Large Scale Structure

The presence of the large scale structures observed in the universe is additional evidence for Dark Matter. The presence of clustering implies that the Dark Matter component must be cold. This is because warm Dark Matter would have been able to free-stream between the over densities and hence smooth them out. In this situation the onset of structure formation would have been delayed. This is inconsistent with current observations. Additionally, the presence of the Dark Matter facilitates the clumping of over-densities, aiding in structure formation [15]. The abundance has to be correct to explain the observed structure. Additionally, the models predict bottom-up structure formation which is also in support of the ΛCDM model [16].

1.1.4 Cosmic Microwave Background

The Cosmic Microwave Background (CMB) provides one of the most compelling pieces of evidence for Dark Matter. First observed in 1965 by Penzias and Wilson [17], this 'fingerprint' exists from early in the universe evolution. At a particular point the universe cooled enough so that photons became significantly less energetic and unable to ionise hydrogen atoms, and consequently the mean free path of such photons became comparable to the radius of the universe. This is a process known as decoupling [18]. As the universe has expanded the wavelength of these photons has become larger and hence formed the relic observed today. The microwave radiation forms a black body spectrum with characteristic temperature of 2.73 K [19]. Importantly, the CMB encodes information about the state of the universe at that time, when oscillations in the density became embedded within the photon distribution. Under gravity, particles coalesced leading to over densities. Outward pressure was then created as the material heated up preventing the full collapse. Together these processes are known as Baryon Acoustic Oscillations [1][20]. The Dark Matter content contributes to the collapse because it has insufficient thermal pressure from a low interaction strength. Hence, the amplitude of these oscillations is dependent on the content of the universe. Occurring over different length scales, these regions of

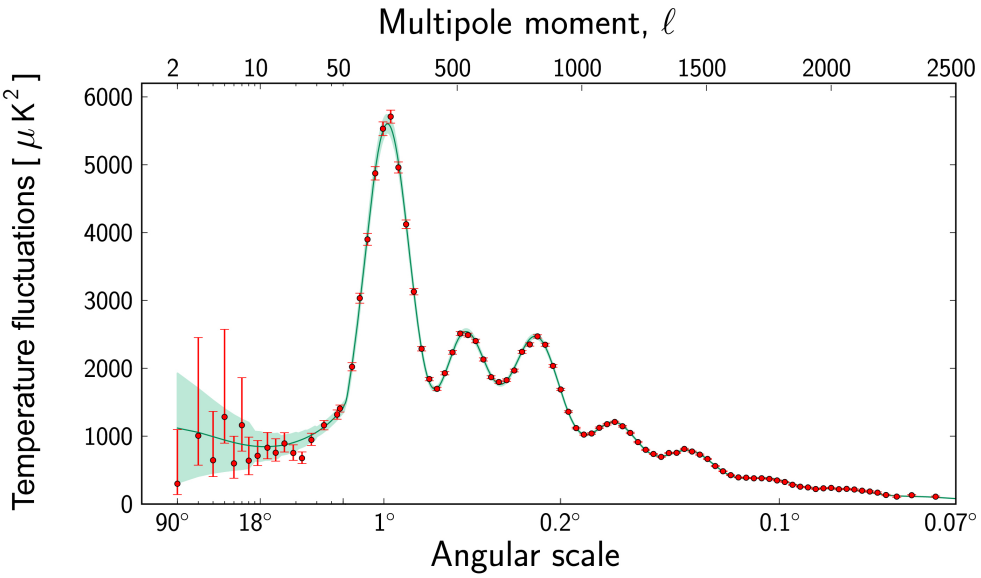

Figure 1.3: The angular power spectrum extracted from PLANCK data. This shows the average temperature fluctuations across different angular scales. Fit is from standard model of Cosmology, and constrains the quantities of Dark Energy, Dark Matter and Baryonic matter. Reproduced with permission from [22]. © ESA and the Planck Collaboration.

over and under densities led to fluctuations in temperature. These fluctuations can be examined across the sky for different angular scales to form an angular power spectrum. The PLANCK satellite [1][21] performed a measurement of such a spectrum in 2018, as shown in Figure 1.3. The spectrum shows remarkable agreement with the standard model of cosmology. The relative amplitudes of the peaks contain information about the abundances of the universe content. The position of the peaks provides information about the flatness of the universe [1] [21]. Observationally the universe is measured to be flat, which constrains the total density.

The amount of cold Dark Matter in the universe is constrained to be approximately 25% of the universe content. The remainder is Dark Energy ($\sim 70\%$) and ordinary matter ($\sim 5\%$) [1]. This includes a small fraction of neutrinos, which behave like hot Dark Matter. This model is known as the ΛCDM model of Cosmology.

1.1.5 Big Bang Nucleosynthesis

During the evolution of the universe the abundances of helium, deuterium and lithium are set. The predicted and observed measurements of the abundances of helium and deuterium are consistent with the ΛCDM model of cosmology [5][18]. The lithium abundance is however not consistent with this picture. Efforts are underway to understand if this is a systematic offset in the measurement [5]. Together the measured values still imply that about 5% of the universe is visible. Therefore, this leaves a considerable fraction of the universe that can be Dark Matter. The model would have to be incorrect by an enormous amount to produce a density consistent with measurements that the universe is flat [23].

1.1.6 Milky Way Dwarf Galaxies

The Milky Way has several satellite galaxies that provide evidence for cold Dark Matter. Firstly, the Dark Matter fraction observed in these galaxies appears to vary, and in some cases dominate [24]. This is a difficult observation to reconcile with modified gravity theories. Secondly, the number of such galaxies is dependent on the mass of the Dark Matter particles. Very light Dark Matter would be relativistic during structure formation and hence the number of over-densities forming these structures would be reduced [25]. Based on simulations, the expected number is into the hundreds [26], as compared with the measured ~ 60 [27]. Although this number is in disagreement, it is possible that some satellite galaxies are not yet observed due to their low luminosity. Understanding all this evidence can place experimental constraints on the nature of Dark Matter.

1.2 Dark Matter Candidates

Based on the evidence outlined above there is good support for the ΛCDM model of Cosmology. The model requires that the Dark Matter must meet a number of criteria. Firstly, the Dark Matter component must (constrained by experimental limits) not emit or absorb radiation and so is theorised to be electromagnetically neutral [5]. Secondly, evidence also supports that the Dark Matter does not readily interact with itself [28]. Thirdly, structure formation ensures that the Dark Matter must be non-relativistic when galaxies began to form [16]. Although the Dark Matter has to be predominantly cold, this does not necessarily place a constraint on the mass of the individual particles. Viable candidates that were in thermal equilibrium with the universe have their abundance 'frozen-out' when the rate of interaction of the particles drops below the expansion rate of the universe [16]. At this point the particle number density becomes fixed and hence would be expected to survive as a relic today. Figure 1.4 illustrates how this process works. Alternatively, there remains the possibility that the Dark Matter can be produced non-thermally through the decay of other species which are out of equilibrium. Several candidates are motivated that can explain such a particle and some are addressed here.

1.2.1 Weakly Interacting Massive Particles (WIMPs)

Weakly Interacting Massive Particles (WIMPs) are a set of models that meet an apparent coincidence known as the 'WIMP miracle'. Such candidates would be a cosmic relic from when the particle froze out in the early universe. If the assumption is made that the particles are all the same type, the interaction cross section is similar in magnitude to the weak interaction, and the mass is of order GeV to TeV then the particle could explain the total Dark Matter density observed today [30][31]. This apparent coincidence that these scales align is hence considered to be a 'miracle' [23][32]. Such a candidate would also potentially provide a solution to the hierarchy problem in the Standard Model of Particle Physics. This arises in the theory of the Standard Model whereby self couplings of the Higgs Boson contribute to the particle mass [33]. With only Standard Model particles,

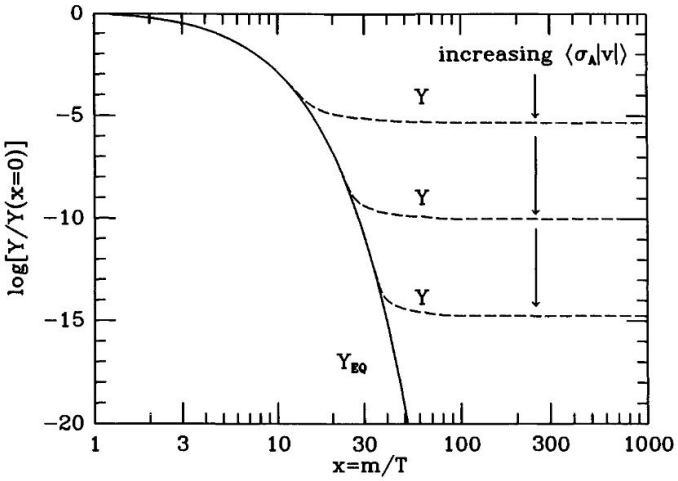

Figure 1.4: Sketch to explain how a thermal relic freezes out in the early universe. The cosmic abundance, Y evolves as a function of the dimensionless parameter x which is a function of the particle mass m and universe temperature T. The temperature is implicitly a function of time. The solid line is when the particle is in equilibrium with the universe and the dashed lines are when it freezes out. Different situations are possible depending on the annihilation cross section σ_A and speed v. Reproduced with permission from [29].

contributions would lead to the mass being much larger than observations show, and is non-calculable if the momentum of the virtual particles is allowed to go to the Planck scale. Supersymmetric particles, a viable extension of the Standard Model of Particle Physics, allow for the terms to the quantum corrections on the Higgs mass to cancel and hence resolve the problem [33]. The lightest supersymmetric particle (LSP) is stable in many theories. Should a particle exist, collider experiments show that it cannot interact strongly with the Standard Model particles over the range of masses for which they have sensitivity. Therefore, this proves to be an ideal WIMP candidate [33].

1.2.2 Axions

Another strongly motivated candidate is the axion. The presence of such a particle was introduced to explain the unnaturalness of the null observation of Charge-Parity (CP) violation in the strong sector of the Standard Model of Particle Physics. This would be allowed under a term in the standard model Lagrangian with a coupling which is experimentally limited to be close to zero [34]. Peccei and Quinn suggested that a U(1) symmetry could be added to the Standard Model to cancel out the CP violation term

[35]. Upon adding this into the theory, a spontaneous symmetry breaking of this field can have quantisations, producing particles called axions. These particles are predicted to have a sub eV mass and so initially appear to be poor candidates for Dark Matter. Their small mass would make them hot Dark Matter if they were produced as a thermal relic (like WIMPs). This is in contravention of observations of structure formation timescales as stated in Section 1.1.3. However, the extremely weak couping means that the particles were never in thermal equilibrium post inflation, so are cold Dark Matter candidates [35]. The Primakov effect provides a possible mechanism by which these particles can be detected through the decay of axions to two photons [36][37].

1.2.3 MACHOs

From observations, micro-lensing objects are observed in galaxies. These are objects where the amplification of a signal from gravitational lensing is detectable, despite the individual finer structure being non-resolvable optically [35]. An event such as this only occurs for objects which are comparable to a fraction of a solar mass; so called Massive Compact Halo Objects (MACHOs). It is expected that these objects are things such as brown dwarfs and neutron stars. These objects cannot explain the whole of the Dark Matter in the universe by current theoretical estimates however could constitute a fraction [38][39].

1.2.4 Primordial Black Holes

Primordial black holes are an additional source of matter content in the universe which is non-luminous. In 2015, the gravitational wave observatories (LIGO and Virgo collaborations [40]) saw the merger of black holes of 36 and 29 solar masses. The production of such massive black holes is difficult to explain astrophysically. Hence, the current hypothesis is that they originate from around the beginning of the universe [41]. More studies are required to determine their abundance to see if they can make up a sizeable proportion of the Dark Matter.

1.2.5 Neutrinos

Experimentally, the mass of neutrinos can be estimated from neutrino oscillation measurements. The small inferred masses mean these particles are relativistic and so cannot be good cold Dark Matter candidates [15]. However, some models predict sterile neutrinos. These are particles that can only be produced from oscillations of the three generations of neutrino and otherwise do not couple to other Standard Model particles. As long as the candidate is stable and has sufficient mass then it may be possible to explain some Dark Matter [42][43]. Experimentally, some of these models are ruled out by the lack of X-ray signatures in telescopes and by oscillation experiments such as MicroBooNE [44][45].

1.2.6 Dark Photons

The dark photon is another possible candidate for Dark Matter [5]. It can be a candidate should the mass be less than twice the mass of the electron to ensure it is stable. Such a particle needs to couple weakly to the electromagnetic charge in the Standard Model to explain the weak coupling of Dark Matter to ordinary matter [46].

1.3 Summary

Based on the evidence laid out here, there is very strong evidence for the presence of Dark Matter. Several candidates have been motivated, including the WIMP for which an emphasis is placed in this thesis. This is because LUX-ZEPLIN has sensitivity to these particles as is discussed in Chapter 3.

Chapter 2

Dark Matter Detection techniques and experimental searches

This Chapter discusses the possible ways Dark Matter can be detected. A particular focus is placed upon direct detection techniques which is the foundation for work in subsequent chapters. The Chapter concludes by examining the general principles by which such particles can possibly be detected and some of the experiments performing such a search.

2.1 Detection techniques

2.1.1 Complimentary methods of detection

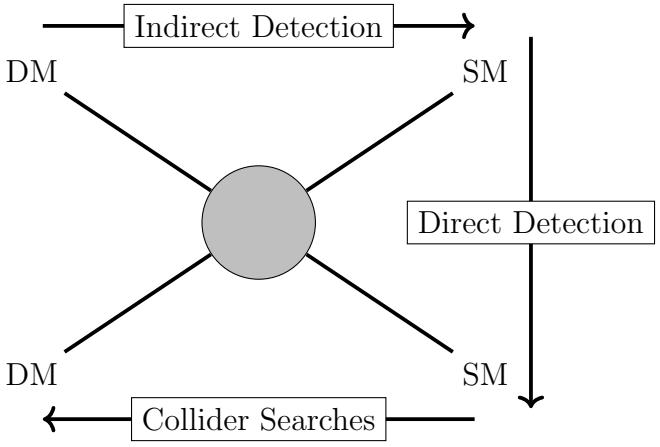

Figure 2.1: Three possible 'detection' mechanisms. Indirect detection involves observing the products of annihilation of Dark Matter particles from regions of high density. Collider Searches are from collisions of Standard Model particles to produce Dark Matter candidates at facilities such as the LHC. Direct Detection is the scattering of Dark Matter from Standard Model particles. Searches take places in experiments such as the range liquid xenon detectors considered in this thesis. The grey dot covering the vertex is to illustrate that the physics requires experimental data to determine the processes happening there.

The detection of Dark Matter would provide conclusive proof of the existance of such a substance. Throughout the remainder of this thesis, the assumption is made that an interaction is possible between Dark Matter particles and Standard Model particles beyond those that are gravitational. Figure 2.1 illustrates three mechanisms by which this may be possible. The diagram is a simplification of the physics, which is currently unknown, and hence is represented with a grey dot. Each method shows some complementarity as is discussed below.

Self annihilation of Dark Matter particles is one such possible process by which Dark Matter signals could be inferred. Here, Standard Model particles would be produced in regions where the Dark Matter density is higher, such as the centre of galaxies or stars [47]. The corresponding signal would be Standard Model particle pairs with the particle energies proportional to the mass of the Dark Matter particles that formed them.

The production of Dark Matter from Standard Model particles is an approach taken at colliders such as ATLAS and CMS at the Large Hadron Collider (Switzerland/France) [48][49]. For example the ATLAS detector examines the collisions of protons at centre of mass energies of up to 13 TeV and can examine the resulting arrangement of particles. A signal for the potential Dark Matter candidate would be missing transverse momentum when a particle which leaves the detector carries some momentum away. Feasibly, as the interaction strength would have to be low, it could meet the requirements of a Dark Matter particle. This method is not truly complementary to the other two methods, because it is possible that the collider can produce particles which are good Dark Matter candidates that have low cosmic abundance. This can occur if long-lived particles leave and decay outside of the detector, because they still have a short half-life when compared with cosmological time. Alternatively, if the annihilation cross section is high then the abundance can be low (Figure 1.4). Hence, a study would need to be performed with other methods, should a possible candidate be found.

The final mechanism is through scattering of Dark Matter particles. This is the approach taken in the direct detection experiments and requires knowledge of both galactic

properties and also the physics of the target material. This technology is used in the LUX-ZEPLIN experiment, on which this thesis is focused, so a more detailed examination is considered below.

2.1.2 Direct Dark Matter Detection

The direct Dark Matter detection methods rely on the ability of candidates, typically WIMPs, to elastically scatter from nuclei. The process can be described by non-relativistic kinematics due to the recoils being typically less than 100 keV [30]. The differential recoil rate (typically quoted in units of events/keV/kg/day) is given by [50]

$$\frac{\mathrm{d}R}{\mathrm{d}E_R} = \frac{\rho_0}{m_N m_\chi} \int_{v_{min}}^{\infty} v f(v) \frac{\mathrm{d}\sigma_N}{\mathrm{d}E_R} \, dv \,, \tag{2.1}$$

where ρ_0 is the Dark Matter density local to the detector and $f(v)$ parametrises the WIMP velocity distribution encompassing the astrophysical constraints. The mass of the WIMP is m_χ and the target material paricle mass is m_N. The term $\frac{\mathrm{d}\sigma_N}{\mathrm{d}E_R}$ is the differential cross-section for the scattering process, which is integrated as a convolution with the velocity distribution from the minimum velocity v_{min} to infinity. Where v_{min} is the minimum WIMP velocity that can produce a recoil of energy E_R. Specifically the recoil energy E_R is given as

$$E_R = \frac{\mu_N^2 v^2 (1 - \cos\theta^*)}{m_N} \tag{2.2}$$

where μ_N is the reduced mass of the WIMP-nucleon system, θ^* is the scattering angle relative to the incoming particle in the centre of mass frame. These contributions encompass the scattering kinematics and the nuclear/particle physics. Due to the dimensionality of the rate given in Equation 2.1 it is evident that a larger detector and a greater active run-time is going to increase the overall number of events.

2.1.3 Galactic properties and astrophysics

As contained in Equation 2.1, the astrophysics of the Dark Matter needs to be considered when calculating the rate of interactions. Phenomenologically, this is because the average

velocity of Dark Matter affects the rate at which a recoil energy E_R occurs. The Standard Halo Model (SHM) is often used as a common reference point to encompass many of the astrophysical effects within the Milky Way. By using a common standard for all experiments, it enables the results to be compared, even if the physical model could do with improvements. The SHM assumes that the Milky Way is a typical galaxy, and hence the Dark Matter environment around Earth can be approximated [51]. The values used within this section have been used throughout and enable comparisons to be made between sensitivity studies.

The density distribution must be able to produce the behaviour expected from rotation curves. Therefore, at larger radii the density must be approximately proportional to $1/r^2$ to be consistent with the flat velocity distribution as observed in Figure 1.1 [52]. Obviously, as the mass of the galaxy has to be finite, this behaviour cannot extend to small radii. Typically, modelled density distributions will flatten or fall with decreasing radius close to the centre of the galaxy. Modelling this behaviour is difficult, and several models lead to what is known as the 'cuspy halo problem' in low mass galaxies [53]. This problem is identified by comparing the inferred density distribution from astrophysical observations (a flattened density profile) to the higher density region predicted by the models.

The local Dark Matter density is assumed to be $\rho_0 = 0.3$ GeV/cm^3. As evidence has evolved, the estimation of this parameter has varied, but uncertainties still remain up to a factor of two [51]. Despite these uncertainties, which lead directly into uncertainties within the experimental limits, the same value needs to be assumed to ensure results can be compared.

The velocity $\mathbf{v}$ is given by a Maxwellian distribution of the form [50]

$$f(\mathbf{v}) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{|\mathbf{v}|^2}{2\sigma^2}\right) , \tag{2.3}$$

where the variance in the velocity is given by $\sigma^2 = \beta v_0^2$. Here v_0 is the velocity around the centre of the galaxy at the solar position and β is a constant which depends on the power law relationships used to model the density [54]. This relationship for the variance

in velocity comes from treating the Milky Way as having rotational symetry about the centre of the galaxy forming a sphere in hydrostatic equilibrium. More generally this is not the case and simulations are needed to model the relationships more accurately. Studies here assume a spherical model with $\beta = 1/2$ and $v_0 = 220$ km/s [51]. The Maxwellian velocity distrubution is cut-off at $v_{\text{esc}} = 544$ km/s;, which is the escape velocity of the Milky Way galaxy.

The velocity distribution introduced above does not encompass the peculiar motion of the Sun nor the orbital motion of the Earth about the centre of mass of the Earth-Sun system. Hence, a Galilean transformation needs to be made to obtain the velocity required in Equation 2.1. The transformation velocity has magnitude $v_e = 230$ km/s [51]. The orbital motion of the Earth leads to a small modulation of the velocity distribution, and hence a modulation of the rate of WIMP scattering over the period of one year. This is due to the motion of the earth being in one direction with respect to the halo for 6 months of the year and in the opposite direction for the other 6 months. In terms of an experiment, this would manifest in the rate of Dark Matter interactions varying sinusoidally over a one-year period. In order to be sensitive to this effect the rate of interactions would have to be high enough that the behaviour could be observed. Any experiment measuring this modulation would have a higher likelihood that any signal was truly Dark Matter.

2.1.4 Interaction of Dark Matter with target nuclei – the cross section

As shown within Equation 2.1, there is the differential cross-section term $\frac{\mathrm{d}\sigma_N}{\mathrm{d}E_R}$. The cross-section σ_N is the interaction strength between the WIMP and the nucleus. It depends on kinematic phase space factors and the underlying quantum mechanics. In combination, the nuclear and particle physics of the interaction is modelled as the sum of spin-dependent and spin-independent contributions.

Firstly, the spin-independent cross-section, which arises from scalar or vector inter-

actions involving Dark Matter is considered [50]. Due to complexities of the density distribution inside the nucleus, it is often factorised through the use of the nuclear form factor F^2, defined to be unity at zero momentum transfer. Hence,

$$\sigma(q,r) = \sigma_N F^2(q,r), \tag{2.4}$$

where σ_N is specific to the interaction physics and is through a currently unknown mechanism. The form factor comes from a Fourier transform of the nuclear density which in general is a function of the momentum transferred, q [55]. Using the Helm representation which softens the nuclear density distribution at the edge of the nucleus, the factor is given by [56][57]

$$F^2(q) = \left(\frac{3j_1(qR_1)}{qR_1}\right)^2 \exp\left(-q^2s^2\right). \tag{2.5}$$

Where, $j_1(x)$ is a Bessel function of the first kind, the constant $s = 1$ fm governs the softening and $R_1 = 1.2A^{1/3}$, where A is the atomic mass of the nucleus. The interaction differential cross-section can then be written as [50]

$$\frac{\mathrm{d}\sigma}{\mathrm{d}E_r} = \frac{M_N \sigma_N F^2(q)}{2\mu_N^2 v^2}, \tag{2.6}$$

where M_N is the mass of the nucleus and μ_N is the reduced mass of the WIMP-nucleus system and v is the relative velocity of the Dark Matter particle with respect to the nucleus. If the interaction of Dark Matter with a proton and neutron are similar, then

$$\sigma_N \approx \left(\frac{\mu_N}{\mu_p}\right)^2 A^2 \sigma_p^{SI}. \tag{2.7}$$

Where the WIMP-proton spin-independent cross-section (σ_p^{SI}) is scaled with the reduced masses of the proton-WIMP (μ_p) and nucleus-WIMP (μ_N) systems and the atomic mass A [58]. The factor of A^2 originates when considering the coherence of the scattering with all the nucleons. This assumption is broken for much higher energies [56]. Hence, this shows for the detection of such a particle, under the assumptions made, larger nuclei with high A give a higher event rate. The spin-independent event rates for a cross-section of 10^{-46} cm^2 are shown in Figure 2.2. This is a typical cross section which is just beyond current experimental limits in the 10 to 100 GeV/c^2 mass range. This includes the form factors for different isotopes, and hence the shapes of the spectra are subtly different.

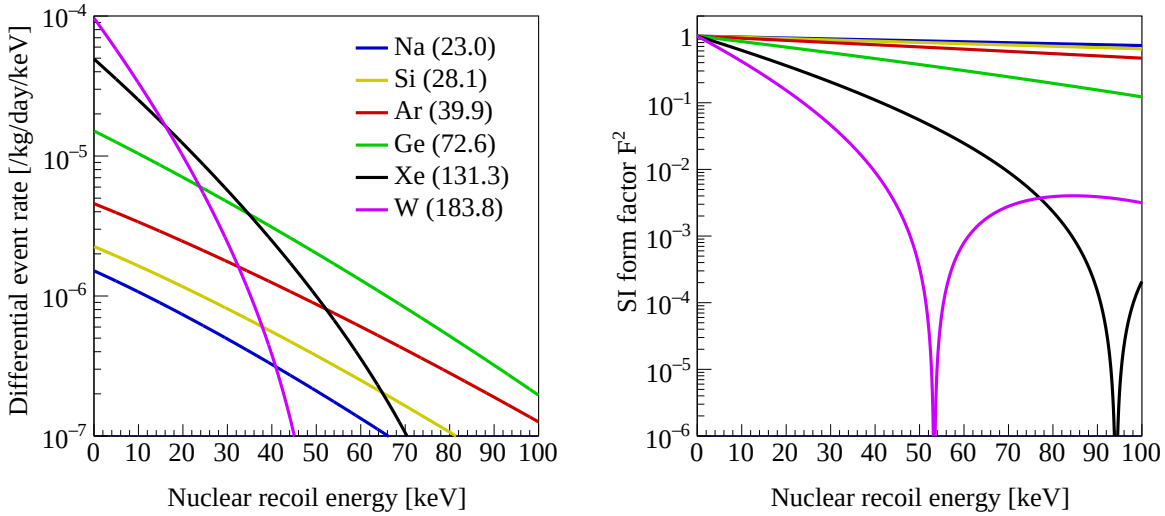

Figure 2.2: Left: Differential event rate as a function of the nuclear recoil energy for different target materials. The cross-section is assumed to be 10^{-46} cm^2. Right: The spin-independent form facor (F^2) for different target materials. Figure from [59].

In addition, spin-dependent scattering which depends on the total nuclear spin, J has, in contrast with Equation 2.6, a differential cross section given by [50]

$$\frac{\mathrm{d}\sigma}{\mathrm{d}E_r} = \frac{16M_N}{\pi v^2}\Lambda^2 G_F^2 J(J+1)\frac{S(E_R)}{S(0)}. \tag{2.8}$$

The Fermi constant is G_F and the function S is the equivalent of the form factor for nuclear spin. This describes the distribution of spin throughout the nucleus and is often parametrised. Importantly, J is only significant when the number of nucleons is odd, because contributions cancel for pairs of nucleons. Throughout this thesis focus will be placed on the spin-independent contributions.

2.1.5 General experimental techniques

The general principle of direct detection is common to a large array of experiments, though it is possible to use different target materials in order to attempt measurements of WIMP-nucleus scattering (Figure 2.2). The main requirements of such a search are similar. The signal rate is usually on the order of 1 to 10 events/tonne/year/keV as is shown in Figure 2.2 for a 10^{-46} cm^2 WIMP-nucleon cross-section. In all cases the low rate can easily be masked by many background events which can mimic a signal. These backgrounds are considered in Section 3.4 for Xenon based technologies, but in summary usually arise from

cosmic rays and residual radioactivity from materials and contamination. These processes can produce neutrons which undergo nuclear recoils in detectors and gamma rays which can interact with the electrons. As experiments become more sensitive, then the coherent scattering of solar and atmospheric neutrinos on nuclei will become a more dominant nuclear recoil background [60]. A common theme is that experiments have to use various strategies to reduce these backgrounds. This can include the selection of high purity detector materials that have lower radiogenic contamination and placing experiments deep underground which reduces the cosmic muon flux by several orders of magnitude [61]. Without implementing many of these measures the obtained sensitivities of many experiments would not have been achievable. Detectors are built that aim to have a low energy threshold as possible because the event rate is higher at the lowest recoil energies (Figure 2.2). Additionally, as the event rate is inversely proportional to the quantity of target material and exposure time, maximising these quantities ensures the sensitivity to lower WIMP scattering cross-sections can be improved. Experiments aim to measure the $\sim 10-100$ keV recoils from WIMP-nucleon scattering. They do this by observing quanta from the interactions with a variety of detector technologies. Phonons (heat), photons, and electrons (ionisation) can all be produced. Typically, nuclear and electronic recoils distribute their energy differently between these three quanta depending on whether the interaction took place with the nucleus or with the atomic electrons. WIMPs are expected to interact with the nucleus, so many experiments will detect multiple quanta which aids in the nuclear-electronic recoil discrimination efforts. Figure 2.3 shows how the current phase space is excluded for spin-independent scattering. The curve showing the coherent scattering of neutrinos is included. This is at the point where an experiment with that sensitivity to WIMP-nucleon cross sections would observe coherent neutrino scattering events which mimic Dark Matter. This is considered in more detail within Section 3.4.

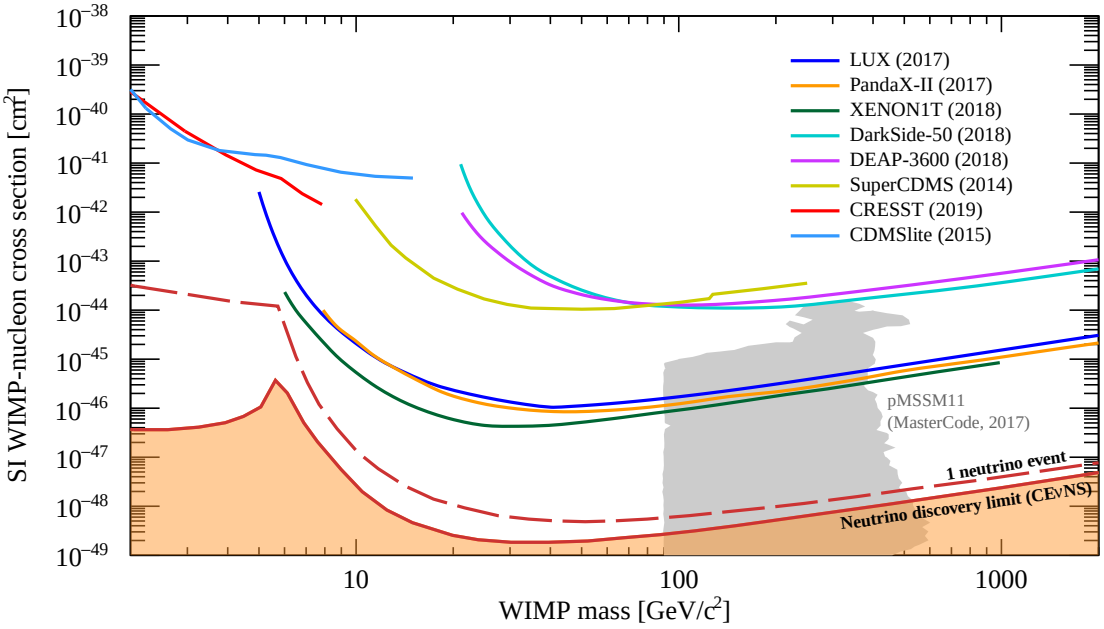

Figure 2.3: Experimental upper limits (at 90% confidence level) for spin-independent WIMP-nucleon scattering for various experiments as discussed within the text. Included is a region of theoretically motivated phase space. The red dashed line and the shaded region show regions where the coherent scattering of neutrinos begins to be a more significant background, forming the neutrino floor. LUX [62], PandaX [63], XENON-1T [64], DarkSide [65], DEAP-3600 [66], SuperCDMS [67], CRESST [68], CDMSlite [67].

2.2 Direct Dark Matter Detection Experiments

2.2.1 Liquid noble gas detectors

Experiments using Xenon as a detector medium have been running for the past few decades. Detectors can be single-phase (liquid only) and observe scintillation photons from interactions. Typically, this is done with photomultiplier tubes which have single photon detection capabilities. Alternatively, detectors can be dual-phase, and observe both the ionisation and scintillation light. As previously introduced, this dual-phase concept aids in nuclear-electronic recoil discrimination due to the different partitioning in these two quanta. Xenon based detectors also benefit from the detector medium itself being self-shielding. Many of the external background source radiation will interact in the outer few cm of xenon so the background event rate close to the centre of the detector is lower than at the edge. In detectors with a good position reconstruction, a region around the edge of the detector is sacrificed which has a much higher rate of backgrounds. This process is known as fiducialisation.

The DAMA collaboration operated a small chamber in the mid 1990s that showed the feasibility of using liquid xenon as a target medium [69]. The ZEPLIN range of experiments then operated in Boulby mine in the United Kingdom from the late 1990s to early 2010s with active masses of a few kg. ZEPLIN-I was a small single phase detector and used pulse shape discrimination techniques (Section 3.1.1) to discriminate nuclear and electronic recoils from scintillation light alone [70]. ZEPLIN-II was the first dual-phase xenon detector [71] and showed the technology could be used to perform a search for WIMPs. ZEPLIN-III was larger still and set competitive WIMP-nucleon cross-section limits at 3.9×10^{-44} cm^2 for a 50 GeV/c^2 mass WIMP [72][73]. Additionally, the LUX experiment, which operated at the Sanford Underground Research Facility, South Dakota, United States from 2012 to 2016 had an active mass of around 250 kg. It set world-leading limits in the 10 to 1000 GeV/c^2 mass range as shown in Figure 2.3[62]. The XMASS detector has recently set the best limits with a single-phase detector of 2.2×10^{-44} cm^2 for a 60 GeV/c^2 mass WIMP [74]. The experiment operates from the Kamioka observatory

and makes use of a circular geometry to maximise coverage with photomultiplier sensors. The XENON experiment series have operated largely in parallel within the Gran-Sasso lab in Italy and have similarly increased their target masses over the past decade or so. PandaX operates from the China Jinping Underground Laboratory (CJPL) and is similar to the LUX and XENON experiments [75]. The world leading limits of these experiments are also included in Figure 2.3. These detectors all use dual phase technology and show that detectors of tonne scale are a feasible proposition for Dark Matter searches. This technology used in the LUX-ZEPLIN experiment, for which this thesis focuses on, is considered in more detail in Chapter 3.

Liquid Argon experiments are also pursuing a similar region of phase space as the Xenon experiments. The DarkSide collaboration have proposals for a new experiment which is 20 tonnes in size, a considerable scale up from the 50 kg which ran previously [76][77]. Argon has the advantage over xenon in that it is financially cheaper to obtain, and has better pulse shape characteristics, aiding in the rejection of electronic recoil backgrounds (Section 3.1.1) [78]. The technique is less effective at the lowest energies when the numbers of quanta are small. The self-shielding of argon is substantially less effective than with xenon. Hence, during fiducialisation more of the detector has to be excluded from the analysis and so the experiments have to be larger. The sensitivity is also lower due to the smaller value of A for the nuclei. From Equation 2.7 the suppression due to the A^2 factor is large and hence the rate is smaller for argon-WIMP scattering. Hence, further requiring the detector to be larger. The argon is also not transparent to the scintillation light produced and so wavelength shifting materials have to be used [78]. These points are considered in considerably more detail in Section 3.1.3 when the advantages of Xenon are considered. DEAP-3600 [79] has operated with an active mass of 3.6 tonnes and is situated at SNOLAB in Sudbury, Ontario, Canada. It is a single phase detector using pulse shape discrimination, in contrast to DarkSide which is a dual-phase detector analogous to the Xenon experiments considered previously.

2.2.2 Solid state detectors

Developments before the liquid-noble gas detectors were mainly with solid state detectors operating at cryogenic temperatures. Materials such as sodium iodide, silicon, calcium tungstate, and germanium have been used. The technology uses the detection of lattice vibrations in the crystals which manifest as phonons. Sensitive thermometers which operate using superconductor transitions can detect these small vibrations. This eV sensitivity gives the experiments the ability to probe to lower mass WIMP-nucleon interactions than liquid-noble gas experiments. Unfortunately this leaves them susceptible to thermal noise from temperature fluctuations which limits the use of this technology. Some detectors also detect secondary quanta (photons or ionisation charge) which aids in the nuclear and electronic recoil discrimination. CDMS operates from the Sudan Complex in Minnesota, United States, where germanium and silicon crystals are used to search for Dark Matter using phonons [80]. They also collect the ionisation charge at electrodes to aid in discrimination efforts. By placing these sensors on two sides of the crystals they allow for surface event rejection, and thus reducing some of the background [81]. Their nominal search produced limits of 1.5×10^{-44} cm^2 for a 46 GeV/c^2 mass WIMP [67]. Competitive limits in the sub 10 GeV mass range were also set using a run called CDMSlite through a dedicated run with the experiment in different electrostatic conditions to convert the ionisation charge to more phonons reducing the energy threshold considerably [82]. In addition, EDELWEISS (germanium crystals) [83] operates from the Modane facility, France. There phonons were also detected as well as ionisation signals. They developed a system of interleaved electrodes which helped reject surface events that was developed further in CDMS [84]. CRESST (calcium tungstate crystals) [85] operates from the Gran Sasso lab, Italy. Scintillation light is collected through a silicon calorimeter and detected with more thermometers. When constructed, the experiments set world leading limits, however these detectors struggle with scalability. Hence, most have decided to focus on the few GeV/c^2 mass range where the low energy thresholds lead to better limit setting capabilities than some xenon experiments. This is illustrated in Figure 2.3 where the 2019 CRESST result

is included.

The DAMA experiment has claimed evidence for a Dark Matter signal, including measuring an annual modulation effect [86]. This detector is a sodium iodide crystal, that observes scintillation light through photomultiplier tubes. This result is in major disagreement with many other experiments including the recently published results of ANAIS-112 [87], which utilised the same technology. Various other detector mediums also rule out such a result.

2.3 Future plans

Using the knowledge and experience from these experiments the next generation of experiments will aim to probe the parameter space further. For example the XENONnT [88], DARWIN [89] and LUX-ZEPLIN [61] experiments all have predicted sensitivities that aim to do this. Particular focus is placed on the LUX-ZEPLIN detector, on which this thesis is based.

Chapter 3

The LUX-ZEPLIN Detector

This Chapter discusses The LUX-ZEPLIN (LZ) experiment, a two phase liquid-gas xenon time projection chamber searching for Dark Matter. A consideration of particle interactions in liquid xenon and how the signals are produced and observed is made. This is subsequently utilised when considering interactions in xenon in Chapters 5 and 6. In addition, several advantages of xenon as a detector medium are discussed, for example the large nuclear size gives an increased event rate at lower energies. The different components of the LZ detector are examined, which can be used together to improve the signal and background discrimination. The typical backgrounds contributing to a Dark Matter search are introduced in addition to calibrations of the detector response. These backgrounds also apply to the lower energy searches as considered in Chapter 6. The Chapter concludes with a consideration of the simulation and analysis frameworks that lead to the projected sensitivity estimates.

3.1 Liquid xenon detectors

The ability to search for Dark Matter with xenon has been used in several experiments, as discussed in the previous Chapter. LUX-ZEPLIN is similar to these experiments and builds on the foundations established over the past few decades.

3.1.1 Interactions in xenon

Xenon exists chemically as a monatomic element. When particles interact, they can do so with the nucleus or with the atomic electrons. Gamma rays and electrons typically interact with the atomic electrons through electromagnetic interactions. Particles such as neutrons and WIMPS are more likely to interact with the nucleus [90]. This interaction can occur through a coherent scattering process, as discussed in Section 2.1.4. Particle interactions produce quanta: photons from scintillation processes, electrons from ionisation process

and phonons from heat losses are all generated. The fractions of the three quanta produced are highly dependent on interaction type and aids in particle identification. Nuclear recoils lose significant proportions ($\sim 85-90\%$ [91]) of the recoil energy to translational motion, and only a fraction is transferred to the atomic electrons. Liquid xenon detectors do not detect this heat and therefore this leads to an offset in the energy scales for the two interaction types. More typically these detectors will detect a combination of an ionisation and scintillation signal, as is the case within the LZ detector. As detailed below, the ratio of the signals from electrons and photons is dependent on the particle interaction type, which aids in particle identification. This then improves the sensitivity to WIMP interactions.

Specifically, the xenon microphysics shows how the production of these quanta occurs. Ultimately these quanta lead to the detectable scintillation and ionisation signals. These processes occur over short length scales close to the vicinity of the original scatter (typically ~ 0.1 µm [92]) so many interactions can be considered point-like. Firstly the interaction (shown here with an electron) can form an exciton (Xe^*) which then leads to the production of primary scintillation light of around around 175nm. This is commonly referred to as vacuum ultra-violet (VUV) [93]. The process proceeds as [30]

$$e^- + Xe \longrightarrow Xe^* + e^- \tag{3.1}$$

$$Xe^* + Xe \longrightarrow Xe_2^{*\nu} \tag{3.2}$$

$$Xe_2^{*\nu} + Xe \longrightarrow Xe_2^* + Xe \tag{3.3}$$

$$Xe_2^* \longrightarrow 2\,Xe + \gamma. \tag{3.4}$$

Here this production method involves the production of $Xe_2^{*\nu}$, an excimer with both electronic and vibrational excitations. Alternatively production of this light can also occur if the original xenon atom is ionised in the process

$$e^- + Xe \longrightarrow Xe^+ + 2\,e^- \tag{3.5}$$

$$Xe^+ + 2\,Xe \longrightarrow Xe_2^+ + Xe \tag{3.6}$$

$$e^- + Xe_2^+ \longrightarrow Xe^{**} + Xe \tag{3.7}$$

$$Xe^{**} + Xe \longrightarrow Xe^{*} + Xe \ (+\text{heat}) \tag{3.8}$$

$$Xe^{*} + 2\,Xe \longrightarrow Xe_2^{*} + Xe \ (+\text{heat}) \tag{3.9}$$

$$Xe_2^{*} \longrightarrow 2\,Xe + \gamma. \tag{3.10}$$

Here the interactions produce the xenon in a higher electronically excited state and is denoted as Xe_2^{**}. In addition, charged xenon molecules (Xe_2^+) are also produced in an intermediate step. The stages depicted in Equations 3.6 to 3.9 are the recombination processes. Here, the ionisation electrons from Equation 3.9 recombine to add to the scintillation yield. Consequently, the recombination process ensures the number of electrons and photons are anti-correlated. The recombination process under a larger electric field is reduced as electrons are removed before they have the opportunity to recombine [94].

At energies above about 1 MeV the probability that further processes can occur is larger and hence the numbers of quanta produced are affected. For example a process called the Penning effect can occur [30]

$$Xe^{*} + Xe \longrightarrow Xe^+ + Xe + e^-. \tag{3.11}$$

Additionally, the quenching of two excitons can also occur

$$Xe^{*} + Xe^{*} \longrightarrow Xe^+ + Xe + e^-. \tag{3.12}$$

In both of these processes the quantity of excitons is reduced. These excitons can then not de-excite to produce light, but release electrons instead, hence affecting the ratio of scintillation and ionisation signals. As these processes only contribute at energies much

larger than the typical recoils considered by WIMPs, these effects can be seen as higher order corrections to the recombination effects considered earlier.

The final number of electrons n_e and number of photons n_γ in an interaction are then related to the number of excitions N_{ex} and ions N_{ion} prior to the recombination process taking place as

$$n_e = N_{ion}(1 - r) \tag{3.13}$$

$$n_\gamma = N_{ex} + rN_{ion},$$

where r is the recombination probability. Therefore, the relationship

$$n_e + n_\gamma = N_{ex} + N_{ion} \tag{3.14}$$

trivially holds from these assumptions. Recombination does not affect the overall sum of quanta but does however change the ratio of charge to light produced in an interaction.

It has been shown that the initial ratio of excitons to ions is about 1 for nuclear recoils and, is about 0.06 for electronic recoils and is largely independent of energy and electric field [91][92][95]. Therefore, the ratio of ionisation to scintillation signal can provide information about the interaction type. The recombination probability is known to be very similar for both electronic and nuclear recoils and so does not affect these interaction differences [92]. This discrimination power is illustrated within the context of a liquid xenon detector in Section 3.1.2. Additional discrimination capabilities come from studying the timing characteristics of the scintillation light. Equations 3.4 and 3.10 show the de-exitation of an excimer that can be initially produced in a singlet or triplet state. The singlet state decays with a characteristic decay time of 3.27 ns whereas the triplet state decays with a characteristic decay time of 23.97 ns [96]. Nuclear recoils and electronic recoils produce different proportions of the singlet and triplet states, so the resulting scintillation light has a slightly different time distribution. The S1 pulse shape can sometimes be distinguishable in each case, and can aid in the discrimination of recoil type. It is utilised more widely in liquid argon detectors due to the larger difference in decay times for each state [30]. To measure these time constants there is a

need to disentangle the pulse shaping characteristics of the readout electronics and the recombination time scale which can both affect the pulse shape.

The energy of an electronic recoil is known to be proportional to the number of quanta, hence

$$E_R = W(n_e + n_\gamma), \tag{3.15}$$

where W is the work function of the xenon [97]. This is the average energy that is needed to produce an ion or exciton. The value is determined experimentally to be about 13.7 eV [92][98]. However, for nuclear recoils the energy needs to be expressed as

$$E_R = \frac{W}{\mathcal{L}}(n_e + n_\gamma), \tag{3.16}$$

where $\mathcal{L}$ is the 'Lindhard factor' and is in general energy dependent, and reflects the large contribution of kinetic energy dissipated during the recoil process [91]. Based on these relationships the energy scale for the detector can be used to determine the energy deposited. Although the sum of the quanta is largely independent of the recombination, the individual yields can be measured for nuclear and electronic recoils as a function of energy. Physically this yield is just the number of quanta for a given energy deposit. Experimentally, for low energy electronic recoils the charge yields (Q_y) are considerably larger than light yields (L_y) [99].

3.1.2 Signals in a Time Projection Chamber

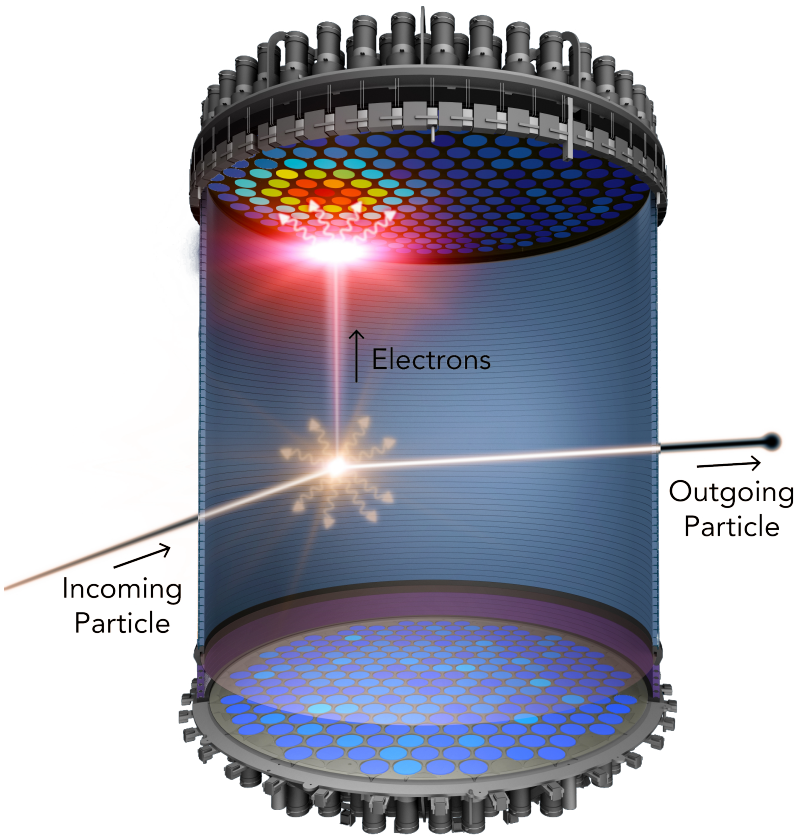

Figure 3.1: An interaction in a TPC instrumented with photomultiplier sensors. This Figure is an LZ collaboration adapted image of one originally shown in [61].

Interactions with liquid xenon can occur within a time projection chamber (TPC) such as one illustrated in Figure 3.1. In this case the scintillation light released during the interaction can be detected by photo-sensors. This is commonly referred to as an S1 signal. An electric field allows the released electrons from Equation 3.5 to drift through the xenon. Typically, a much larger electric field is used to extract the electrons from liquid into gas. This large electric field is required to overcome the potential barrier which is present at the liquid surface [30]. This is commonly achieved my the use of two grids which span the liquid surface. Electrons which are extracted then can excite the gas molecules in a process similar to the one discussed above and hence produce secondary scintillation light which is similarly detected by photo-sensors. This is commonly referred to as an S2 signal. The drifting process is orders of magnitude longer then the scintillation. The time

delay between the S1 and S2 signals can provide a good measure as to the depth of the interaction. An example event topology is illustrated in Figure 3.2 from LUX-ZEPLIN simulations showing such an S1 and S2 signal. As is shown in Figures 3.1 and 3.2 the hit pattern on the photo-sensors can be used to provide an indication as to the position in the plane perpendicular to the drifting electrons.

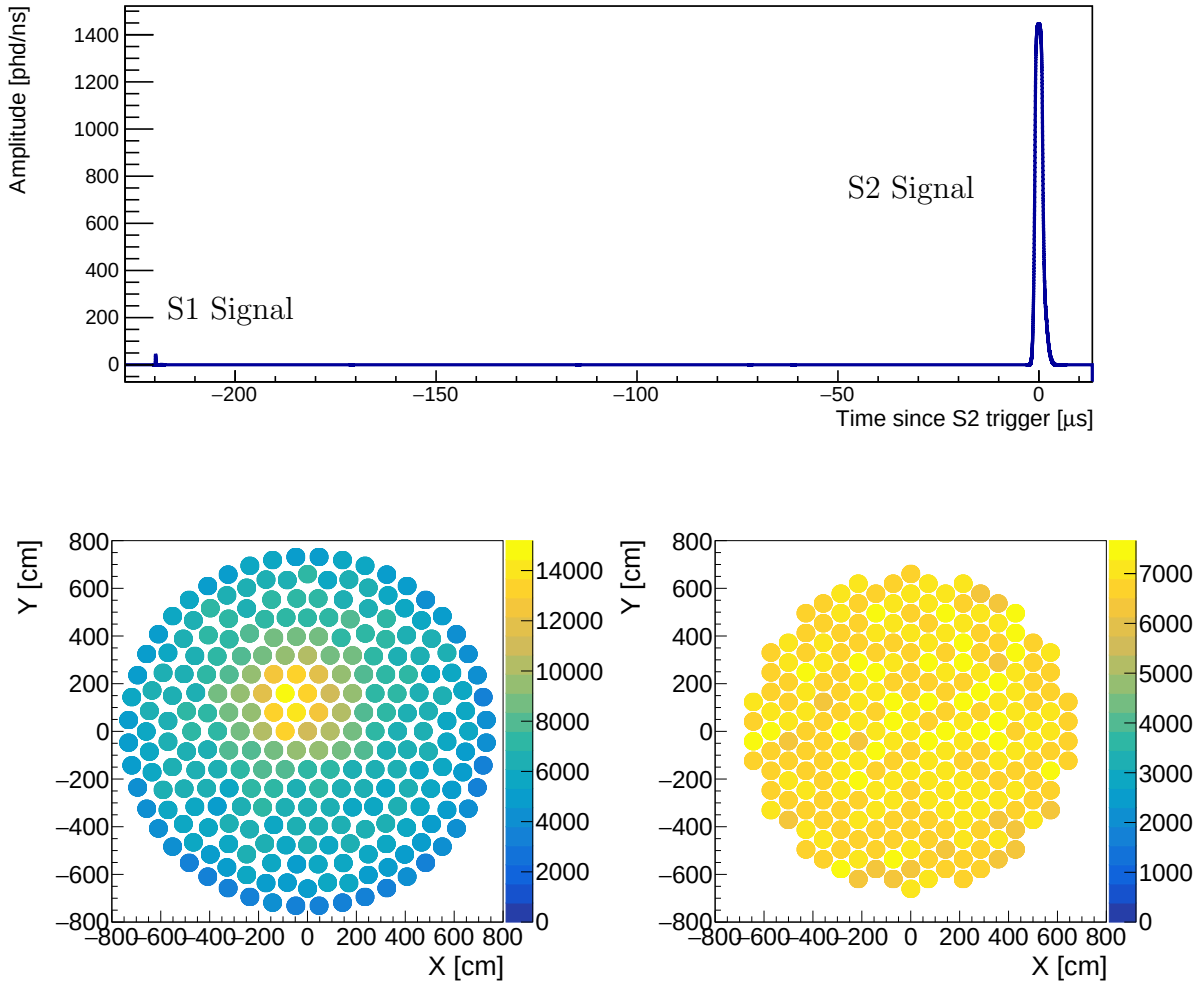

Figure 3.2: (Top) An example simulated high energy event with S1 and S2 signal. Not all physical phenomena affecting the signals is simulated for clarity. The top (bottom left) and bottom (bottom right) photomultiplier arrays are shown for the duration of the S2 signal showing the localisation of the signal about the electron extraction position.

A consideration of both signals is also used to provide discrimination between electronic and nuclear recoils. Due to the difference in the initial exciton and electron ratios, discussed in Section 3.1.1, the S2 signal is smaller for nuclear recoils. By using calibration sources, such as that of tritium, shows that the events form banded populations (Figure 3.3). Considering only events with an S2 signal smaller than one standard deviation below the median of the nuclear recoil band, ensures better than 99.5% discrimination.

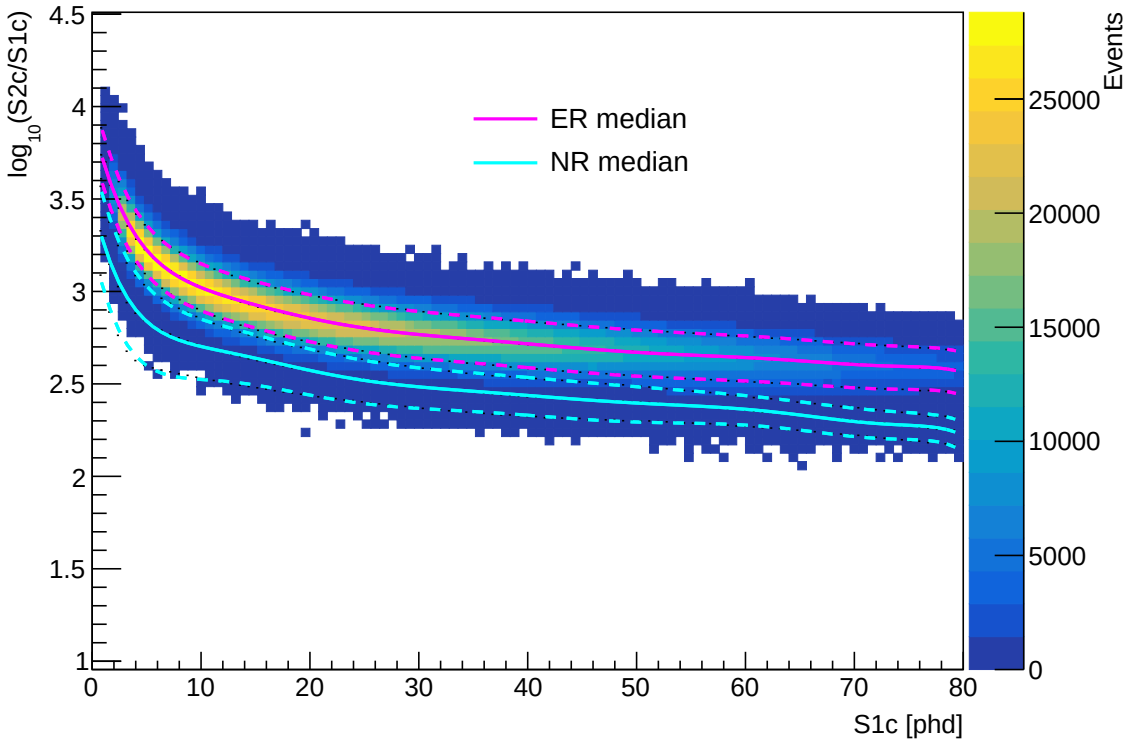

Figure 3.3: Simulated 10^7 tritium events (electronic recoil) for a nominal LUX-ZEPLIN electric field configuration using the NEST simulation package [99]. Median and 95% contours for nuclear and electronic recoils (NR and ER respectively) are also shown, which can be determined experimentally from calibration. The logarithm of the ratio between the spatially corrected S1 and S2 signal sizes (S1c and S2c respectively) acts as a discriminator between interaction types as a function of S1c. Figure provided from [100].

The S1 and S2 signals incorporate position dependent efficiencies in the detector. These can come from differences in the light collection efficiency, geometry, xenon purity and non-uniformities in the electric fields. Additionally, there are efficiencies related to the photo-sensors that need to be taken into account for both signals. The position dependence is discussed in more detail in relation to the S2 signal in Section 5.4. For both signals the position dependence is removed using mono-energetic sources to form corrected signals. S1 and S2 signals are therefore corrected to $S1_c$ and $S2_c$. The S1 and S2 signals are measured as the number of photons detected (phd). This is based on a calibration using the time integrated charge deposited on the sensor from a single photoelectron. The number of photons within each signal is therefore proportional to the initial number of photons and electrons released at the interaction site. The proportionality factor g_1 (g_2) quantifies the detection efficiency for the S1c (S2c) signal and is required to reformulate Equation 3.16 as

$$E_R = \frac{W}{\mathcal{L}} \left(\frac{S1_c}{g_1} + \frac{S2_c}{g_2} \right) \tag{3.17}$$

where $S1_c$ and $S2_c$ are the time integrated signals for each signal corrected for position dependence. As discussed, for electronic recoils, the Lindhard factor $\mathcal{L}$ is unity but needs to be included for nuclear recoil events where it is $\sim$ 0.15 at 5 keV.

3.1.3 Advantages of liquid xenon as a detector medium

Building an experiment using a liquid noble gas TPC as outlined in Section 3.1.2 has several attractive features. Firstly the detector medium itself is inert and so will not readily react with the detector materials. This means that the detector can be stable over a long period of operation. Additionally, the scintillation light produced (175 nm) is able to propagate through the medium [30]. Xenon does not have any long-lived isotopes, which means that much of the radioactive contamination will decay away once stored underground. The light yields in xenon are also sufficiently high (about 70 photons/keV for typical recoil energies) so that interactions can be detected [30]. The coherence of the nuclear recoils with the xenon also ensures that the event rate is sufficiently larger than other target materials. This is because the scattering scales with the square of the atomic

mass. This effect is seen specifically within Figure 2.2. The presence of ^{136}Xe with an odd number of nucleons also makes it possible to perform a spin dependent search for Dark Matter. Xenon detectors are also scalable with larger masses and more photo-sensors. Of course, in the process of doing this, the challenges of such a detector change, as do the relative importance of different backgrounds and additionally requirements on the purification and circulation systems.

One major advantage of xenon based technologies is the self-shielding ability. Essentially this is due to the attenuation of mainly gamma rays in the outer regions of the xenon detector. This means that the backgrounds from outside the experiment reduce exponentially towards the centre [61]. Larger experiments have larger shielding and hence they have much lower rates within the centre of the experiment than near to the walls. The self shielding ability is greater in xenon than other detector target mediums due to the high density. Hence the region suitable for performing a Dark Matter search can be significantly larger than with other detector materials.

3.2 LUX-ZEPLIN

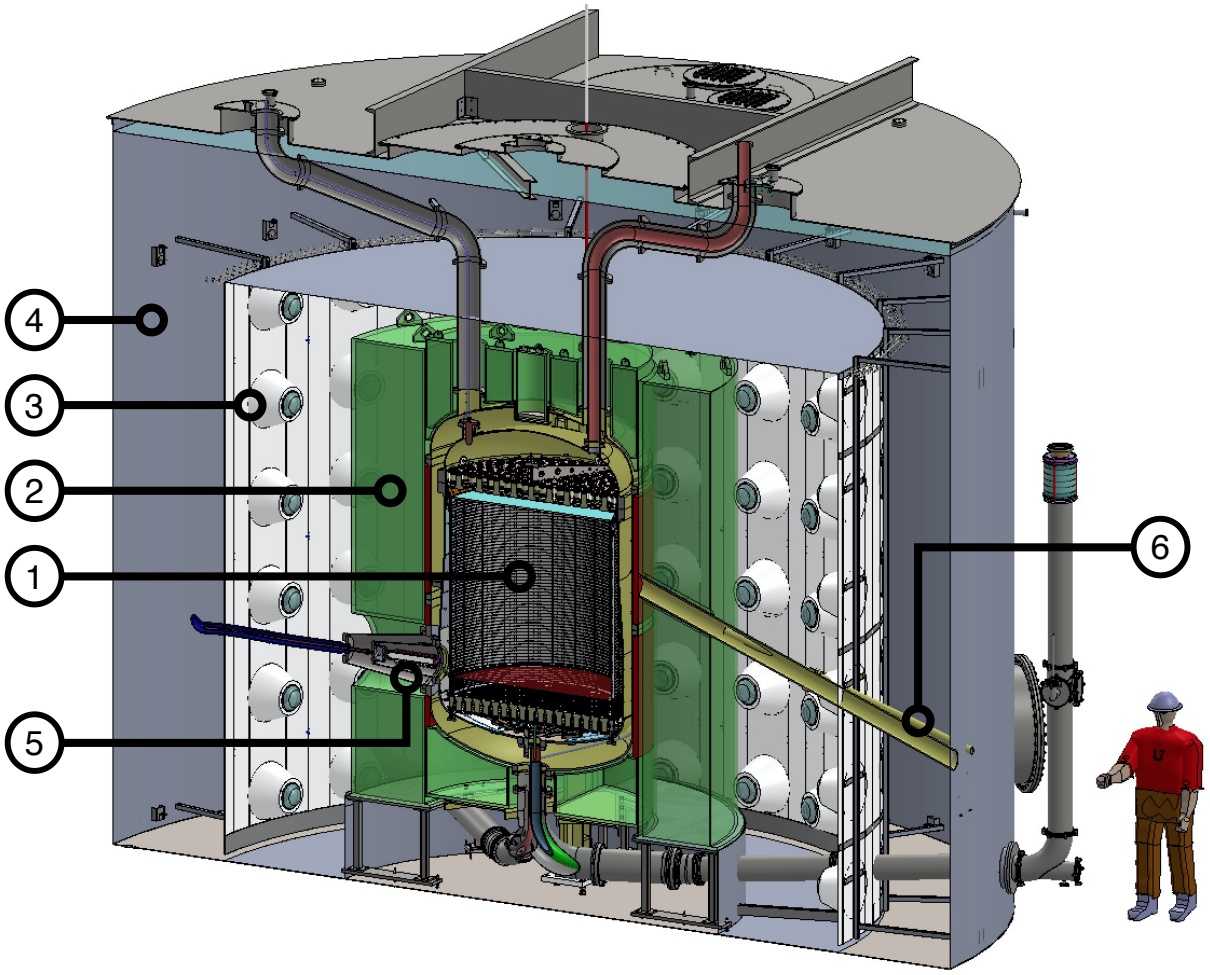

Figure 3.4: Overview of the LUX-ZEPLIN detector systems. 1-The inner Xenon detector/TPC, 2-Outer detector liquid scintillator tanks, 3-Outer detector Photomultiplier tubes and Tyvek sheeting, 4-Water tank, 5-The cathode high voltage connection, 6-One of the neutron calibration conduits. Figure from [101].

The LUX-ZEPLIN experiment is located 4850 feet underground at the Sanford Underground Research Facility, Lead in South Dakota. The detector utlises the cavern that was first hollowed out for the Homestake experiment [102], then extensively redeveloped for LUX [62].

3.2.1 Time Projection Chamber - TPC

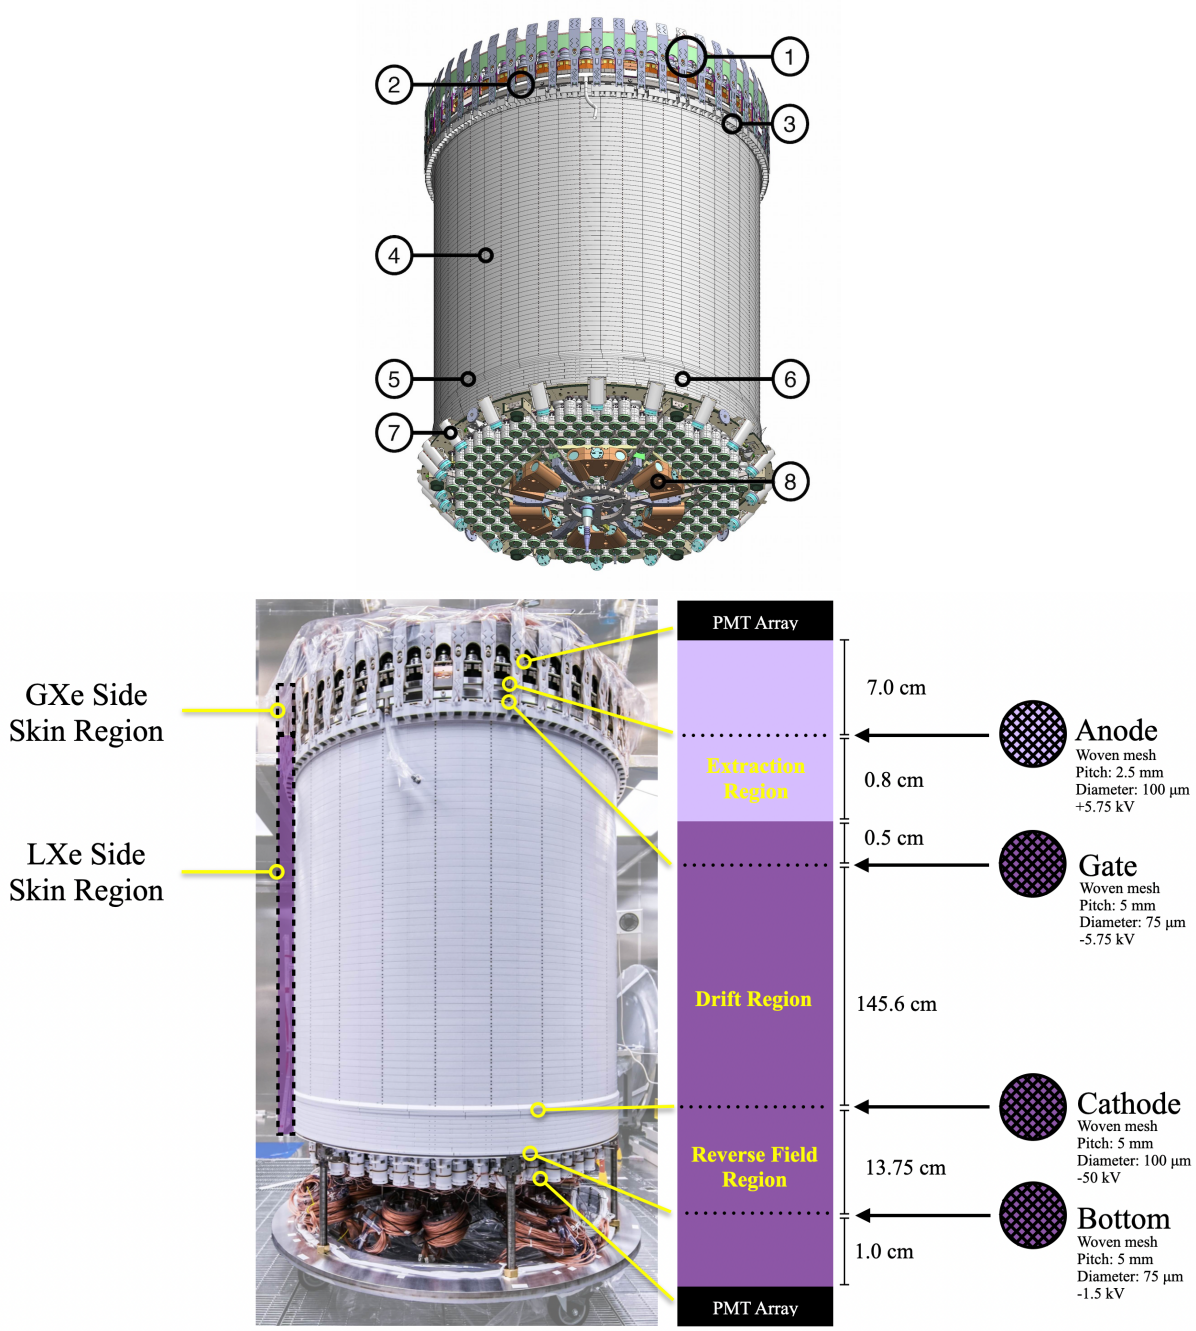

Figure 3.5: (Top) A rendering of the outside of the TPC with a photograph of the completed detector. 1- Top photomultiplier array, 2-Weir region, extraction region and liquid level, 3 - Photomultipliers for the xenon skin, 4-The main TPC field cage, constructed of Teflon and titanium filed shaping rings, 5-Cathode, 6-The reverse field region, 7 and 8 Further instrumentation for the xenon skin. (Bottom) A photograph of the finished assembly alongside a schematic of the dimensions of the different regions of the xenon detector and TPC. Here, the lighter shade signifies this region is gaseous xenon. Photograph taken by Matthew Kapust, Sanford Underground Research Facility. Figures from [101] and [103] respectively.

The Time Projection chamber (TPC) is filled with about 7 tonnes of liquid xenon (LXe). The TPC is constructed from PTFE, which is highly reflective to the xenon VUV light (better than 99.5%) [104]. The outside of the TPC is illustrated in Figure 3.5. Four finely woven grids [103] define the electrostatic environment within the chamber as illustrated in Figure 3.5. These grids are custom-built for the experiment using a loom. Encased within the TPC walls are metal rings which form a field cage. This is used improve the uniformity of the electric field within the detector. Between the gate and cathode is the drift region or xenon bulk. Here, interactions with liquid xenon can produce ionisation electrons and scintillation light (Section 3.1.1). The scintillation light forms the S1 signal. Under the nominal electric field of 310 V/cm the electrons drift to the top of the detector by the gate. The anode and gate grids define the electroluminescence or extraction region. Between these grids is the liquid surface. Here electrons from the interaction that have drifted are accelerated under a much larger electric field ($6 - 8$ kV/cm), and are extracted into the gas and undergo electroluminescence contributing to the S2 signal. The mechanics of this process are important to the thesis, which leads to variations in the pulse shape and signal size and are discussed throughout. This is particularly the case for the low energy analysis techniques discussed in Chapter 6. The bottom grid is used to shield the bottom photomultiplier tubes from the 50 kV potential that is placed on the cathode. The electric field in the region between the bottom grid and cathode (the so called 'Reverse field region' or 'RFR') means that electrons drift downward and hence S2 signals cannot be produced from interactions within this region. The detection of the light from the S1 signal is almost instantaneous to the interaction, but the S2 signal is delayed by a few to a few hundred µs due to the time required to drift electrons. As the drift velocity is approximately constant throughout the TPC volume, then this delay provides a good proxy to the depth of an interaction and forms the basis for a TPC operation.

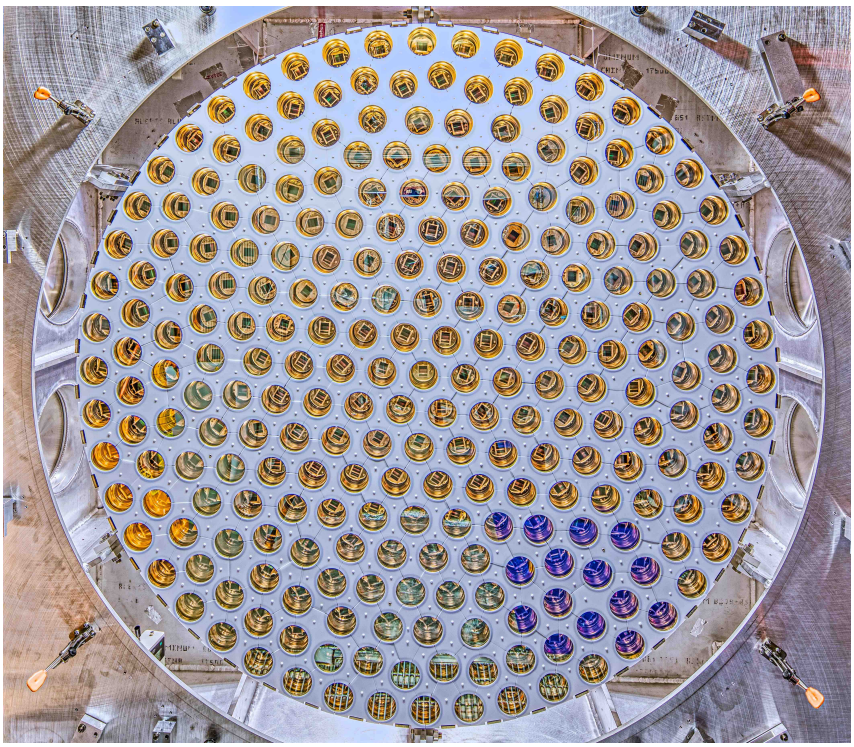

Figure 3.6: The photomultiplier sensors forming the top array within the assembly enclosure. The sensors form a hexagonal structure within the centre, but move to a circular pattern on the outside. This has been optimised to aid in light collection and position reconstruction of events. Photo by Matthew Kapust, Sanford Underground Research Facility. Figure from [101].

3.2.2 Photomultiplier arrays

Instrumenting the detector are 494 Photomultiplier Tubes (PMTs) as illustrated in Figures 3.6 and 3.7. PMTs are able to detect individual photons using the photoelectric effect. A photon colliding with the photocathode within the PMT is able to liberate an electron. This single electron is amplified with a system of dynodes. Ultimately the signal is increased enough that a measurable signal is produced. Due to the dispersion of the electrons the signal does not form an instantaneous response, but rather a continuum, often referred to as a pulse. LZ uses R11410-22 (Hamamatsu) tubes that were specifically produced for the experiment. These sensors were selected for their relatively high radiopurity and quantum efficiency. At a value of 20-30% the quantum efficiency defines the probability that an incident photon is detected [93]. Of the 494 PMTs, 253 tubes are located above the TPC in the gas region. Their arrangement was optimised to ensure that the position resolution in the plane parallel to the grids is as good as possible. The

position reconstruction is based on the S2 signal hit pattern on the top array photomultiplier sensors using the Mercury algorithm [105]. About 2/3 of the S2 electroluminescence light is detected on this array and is highly localised to the point where the electrons are extracted from the liquid. From calibrations (Section 3.3), the Mercury algorithm is tuned to the light response as a function of the position and the resulting light response functions are then used throughout operation. For any given event the Mercury algorithm uses these calibrated functions to reconstruct the position of the event. In combination with the drift-time measurement, this therefore provides full three-dimensional reconstruction of the event position. The LZ coordinate system has the vertical z axis, with the x and y axes completing the Cartesian basis. The PMTs on the bottom array are used to enhance the light detection of small S1 signals. The bottom array often observes more S1 light due to the reduced optical transmisibility at the liquid surface [61]. In Chapter 6 the effects of not being able to do this three-dimensional reconstruction are examined in considerable detail.

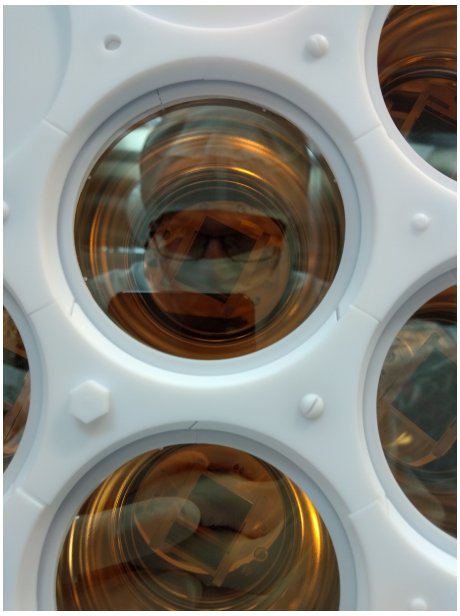

Figure 3.7: A photograph of several photomultiplier tubes on the bottom array. The photocathode is visible within the photograph.

The signals from the PMTs are amplified and digitised using the LZ Data Acquisition system (DAQ). This system contains amplifiers for the signal, digitisers and the ability to write to disk. From there the data is able to be transported for analysis in numerous

locations worldwide.

The LZ signal from each PMT is fed into two analogue amplifiers in parallel. One of these amplifies the signal by a factor of 4 and is referred to as the low gain (LG) or high energy (HE) channel. The second amplifies the signal by a factor of 40 and is referred to as the high gain (HG) or low energy (LE). The two channels are used in combination to provide information about an event. Then these signals are fed into the digitisers. They have 14 bit precision operating over 2 V range, at 100 MHz frequency [61]. These digitisers will saturate for the highest energy events in HG. Hence, the need to have the LG capabilities. LG is not used universally because the amplification of the signal leaves the possibility that post signal noise from the electronics is larger in relation to the signal trace. The system has a trigger to ensure only data forming pulses is digitised (Pulse Only Digitisation) which reduces the disk writing burden [61]. This is also used in conjunction with an online trigger to save to disk only events that have a measurable response.

3.2.3 Inner and outer cryostat vessels

The TPC sits within the titanium inner cryostat vessel (ICV), filled with about 10 tonnes of xenon [106]. This vessel is situated within a larger outer cryostat vessel (OCV) to form a vacuum layer to maintain thermal integrity of the ICV.

3.2.4 Veto detectors

To aid in the background discrimination from signal events the detector has inbuilt veto capabilities. Particles such as neutrons and gamma rays can often interact multiple times and so an observation of a multiple scatter within the TPC can help reduce backgrounds. However even greater power comes from considering optically separated regions of the detector. The detector is constructed with two such regions that can veto events. Specifically a region called the xenon skin surrounds the TPC volume (Figure 3.5). Contributing to the self shielding within the detector, the region is instrumented with further PMTs which can detect scintillation light. The outer detector contains gadolinium doped liquid

scintillator and sits outside the OCV. Gadolinium has a large neutron absorption cross section [61][101] and is hence used to capture neutrons which may have also interacted in the TPC. As shown in Figure 3.4, the tanks surround the cryostat vessel. The outer detector is then further surrounded by a water shield which also attenuates particles. From the instrumentation of the outer detector and water tank with PMTs the scintillation light from the gadolinium interactions can similarly be detected. In the case of the xenon skin and the outer detector, an observation of a signal in the veto detectors, in coincidence with one within the TPC, enables the events to be removed from the analysis. WIMPs would not interact multiple times within the detector due to the very small cross section being probed. Passively the veto detectors also have the advantage of introducing more material between the cavern walls and the TPC that attenuates more particles, and thus reduces backgrounds within the TPC.

3.2.5 Circulation system

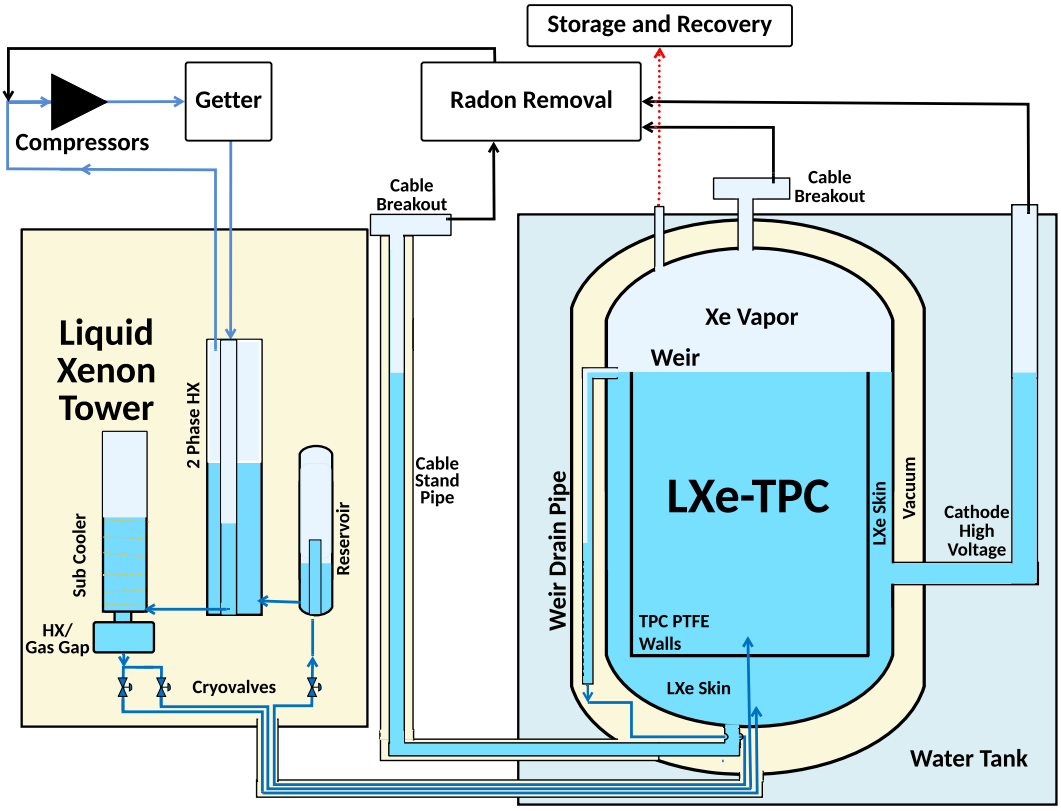

Figure 3.8: Schematic diagram showing the basis of the circulation system. Figure from [101].

The liquid within the experiment circulates using a network of plumbing as simplified within Figure 3.8. Liquid overspilling the weir within the TPC is collected in a reservoir within a structure known as the 'Xenon Tower'. This is used to stabilise the liquid level within the experiment. From there the liquid flows into a heat exchanger to provide some cooling to the physically separated xenon flowing into the experiment. When the xenon has boiled off then it is compressed and fed into a device called the getter. Here impurities in the xenon are removed. These impurities include particles which have out-gassed from the detector components and that are present naturally. For example, oxygen ions can be formed from the ionisation of out-gassed oxygen. The getter operates by flowing the gas through zirconium at several hundred degrees Celsius. The presence of many of these elements hinders the collection of charge within the TPC and so is essential to remove as much as possible. The zirconium reacts with non-noble gas elements readily and hence prevents them from being recirculated in the system [107]. The gas returns to the heat exchanger and is cooled from the outgoing gas until it liquifies. Upon entering the sub-cooler the liquid is further cooled until it is below that of the detector. This is done using systems containing liquid nitrogen and electrical heaters. At this point the xenon then flows into the xenon skin or the TPC to complete the cycle. The system is designed to operate so that the detector contents are circulated every few days. The circulation system also contains a charcoal based radon removal system which helps remove radioactive radon from the experiment [108]. This is required because the getter is insensitive to noble gas elements due to their chemical properties being similar. The system relies upon the different velocities of the radon and the xenon through the charcoal column. The ability to circulate the xenon enables the impurities to be removed.

3.2.6 Instrumentation sensors

The LZ detector is instrumented with various sensors which monitor the conditions within the detector. This can help provide feedback to influence the circulation system. With the detector being sealed, these sensors are required to operate without intervention for the duration of the experiment. Thermometers instrument the detector and are discussed

in considerable detail within Chapter 4. Furthermore, level sensors form part of the circulation system and TPC to ensure the detector can be levelled successfully and monitor long term stability. Six sensors instrument the extraction region in order to perform these tasks, capable of readout to tens of µm. These sensors are considered in Chapter 5 where the calibrations are performed and the levelling is considered. A series of six position sensors instrument the gas region above the top PMT array and aid in the knowledge of how the detector cools. These are also discussed in Chapter 5. Loop Antennae are fixed to the PMT arrays and can be use to veto TPC events associated with electrostatic breakdown. Acoustic sensors are attached to the outside of the ICV vessel and are capable of detecting bubbles and vibrations. In addition to the pressure sensors, any remaining knowledge must come from the xenon response of the detector.

3.3 Calibrations

In order to understand the xenon response, efficiencies and detector effects calibrations must be performed. Calibration sources are typically selected to each do a specific job, but together can produce a complete picture. Calibrations are used to characterise the energy reconstruction, understand time offsets in the detector or to calibrate the electronic and nuclear recoil bands (two populations in the S1-S2 parameter space that can be used to perform particle discrimination). Many of the sources are deployed using the calibration source systems which can lower them using motors. Due to the self shielding of the liquid xenon, the full extent of the detector is particularly hard to calibrate in this way. As a result, calibrants are mixed into the circulation system and allowed to circulate into the system. Once homogeneously mixed, these calibrants are used to measure the position dependence within the liquid. This is discussed in more detail within Chapter 5 where this procedure is considered.

3.3.1 Sources external to the ICV

There are 4 deployment points for external sources. One such point is on top of the outer cryostat vessel (OCV). This is specifically reserved for photoneutron sources such as YBe. The principle operation consists of a gamma source surrounded by a material that will undergo a photoneutron reaction. This source requires considerable shielding surrounding this material to improve the neutron to gamma photon ratio. It is intended to perform some detector nuclear recoil response at the top of the TPC and includes responses at lower energy [61]. The other three deployment points are located between the inner and outer cryostat walls. In three tubes, a capsule containing the source is capable of being lowered to below the cathode level. Most deployment plans however will include examining the response at many points between these levels. One source that is used is that of ^{57}Co, which emits predominately $\sim$ 122 keV gamma rays and is useful for xenon skin calibrations [109]. The ^{57}Co sources were also used when levelling the detector to increase the event rate sufficiently within the extraction region. This is discussed in more detail in Chapter 5.

In addition to these source points a deuterium-deuterium generator is used to produce neutrons. This generator produces a monoenergetic beam of neutrons from a fusion reaction. By placing the generator at the end of evacuated conduit, then attenuation between the source and the detector is minimised. One such conduit is shown in Figure 3.4. The source was successfully deployed to calibrate the LUX experiment nuclear recoil response from 0.7 keV to 75 keV [110].

3.3.2 Dispersed sources

To understand the features of electronic recoils and to identify where in the parameter space these events lie, beta sources are injected into the detector and allowed to disperse. The response can be measured for a broad range of energies due to the unequal partitioning of energy during beta decay between electron and neutrino. In the case of LZ ^{220}Rn is chosen for this purpose. Most of the ^{220}Rn will have decayed within the circulation system

before reaching the TPC to ^{212}Pb, so this beta spectrum is the one that is measured. This source, does not have long lived decay products (all $<$ 12 hrs). Hence, residual activity within the detector will drop to background levels after a few days and thus not compromise searches for Dark Matter. As an alternative tritiated methane (also a beta emitter) could be used, as was successfully done in the LUX experiment [61]. This source is disfavoured due to the longer half-life of 12.3 years which would require active removal in order to prevent the addition of background [61]. The effectiveness of the removal process needs to be evaluated based on the achieved circulation rates before this calibration is attempted.

Mono-energetic sources are also injected to act as standard candles. For example, one internal calibrant used within LZ is ^{131m}Xe. This metastable isotope decays with 164 keV gamma emission [101]. It is produced in this exited state from the decay of ^{131}I. The half-life of 11.9 days ensures that the source can become homogeneously dispersed throughout the TPC and xenon skin, and so can be used for intra-day stability monitoring. It is also used to characterise the attenuation from xenon impurities, as discussed in detail in Chapter 5. This source has energy deposits larger than the typical WIMP search range in LZ (6 keV$_{\mathrm{NR}}$ to 30 keV$_{\mathrm{NR}}$), and causes some saturation of the digitiser from the largest pulses. The lower energy decays of the calibration with ^{83m}Kr are not as susceptible to such effects. This isotope is therefore used in addition to ^{131m}Xe. The most typical decay is two successive de-excitations which lead to a total energy deposit of about 41 keV [111]. The second decay has a half-life of 154 ns and so some pulse shaping characteristics may be subtly different from standard electronic recoil backgrounds. The first decay has a half-life of 1.8 hours and so any injected source will decay away quickly. Numerous events can be produced in a short time period with this source, and so this is good for calibration of the position reconstruction and examining spatially dependent effects (as discussed in Chapter 5.4).

3.4 Backgrounds

During the search for Dark Matter the detector observes background events. Some of these events are capable of passing the typical data analysis cuts considered in Section 3.5.2 despite many mitigation techniques. Hence, these events are indistinguishable from WIMP scattering on xenon nuclei. The contributions to the background are divided into several main categories, shown in Figure 3.9. The next few Sections discuss these backgrounds in more detail.

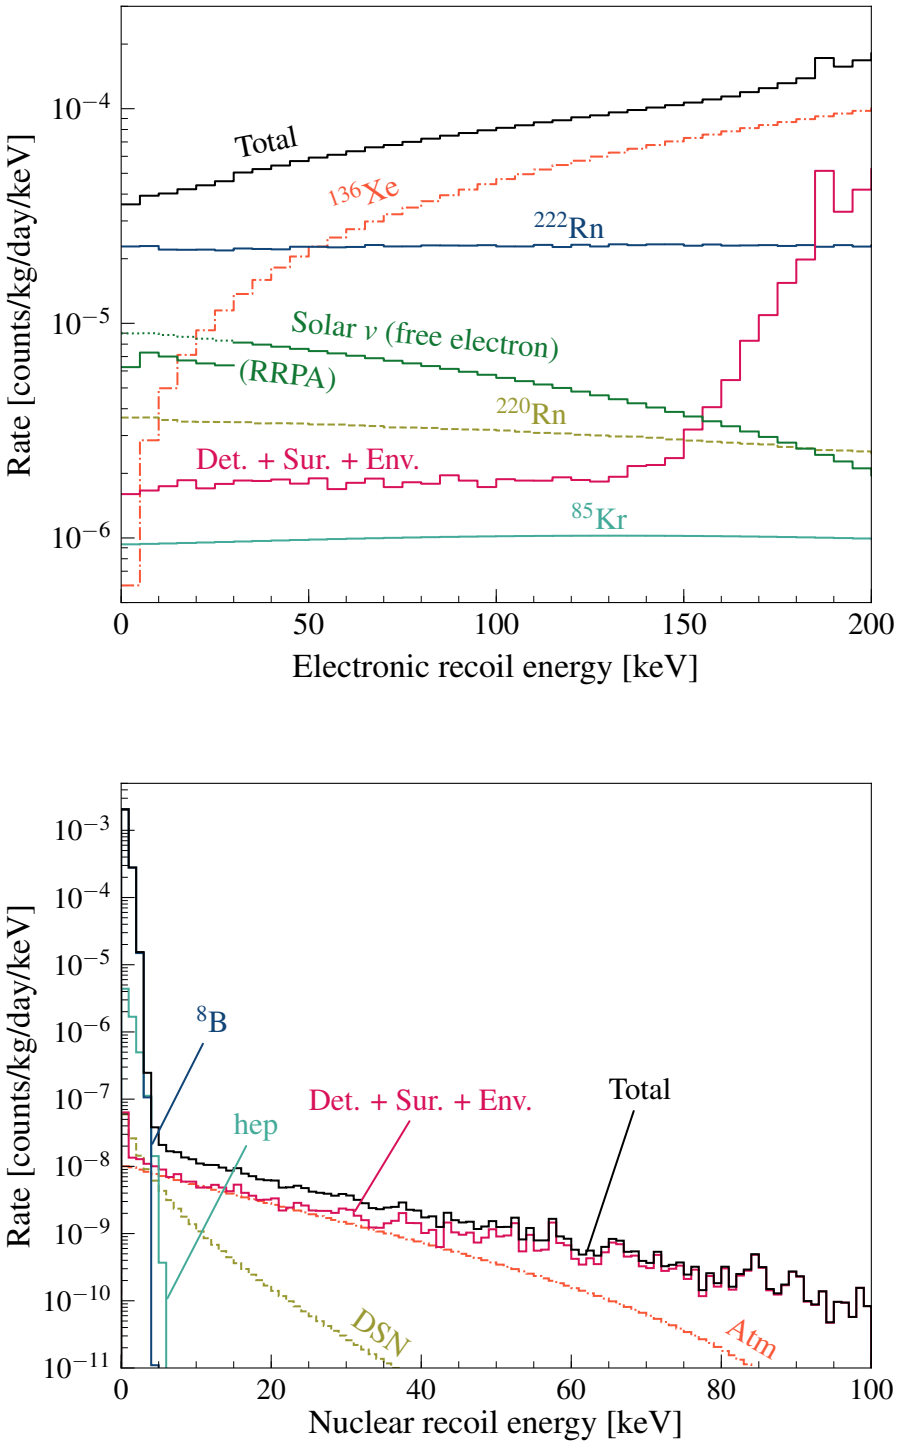

Figure 3.9: Electronic (Top) and Nuclear (Bottom) recoil backgrounds for the standard WIMP search as a function of recoil energy. No veto cuts or the detector efficiencies are applied to these plots. Figure from [112].

3.4.1 Astrophysical

One astrophysical component to the background comes from cosmic rays. By placing the experiment underground, the background from cosmic rays, which predominantly manifest as muons following showering processes in the atmosphere, is reduced by several orders of magnitude. The factor of reduction of such muons is about 3×10^6 relative to the surface [61][113]. Even particles that make it to the experiment are relatively easy to identify due to the high ionising capabilities of muons. It is possible that the muons could induce showers within the rock and these particles could interact within the TPC. Such a background though is considerably sub-dominant [61]. Additionally, cosmogenic activation of xenon or detector components will occur, but will reduce after they are taken underground.

Astrophysical neutrinos also contribute considerably to the background in LZ. Solar neutrinos from the pp chain can scatter from atomic electrons contributing to an electronic recoil background (Solar ν)[114]. Additionally, neutrinos formed in cosmic ray showers in the atmosphere (atm) and diffuse neutrinos from supernovae (DSN) can also incite nuclear recoil interactions. ^{8}B neutrinos and helium-proton (hep) neutrinos from solar processes contribute to the nuclear recoil backgrounds below 5 keV_{NR}. The nuclear recoil background is through coherent elastic neutrino-nucleus scattering (CEνNS) which is similar to a hypothesised WIMP interaction. Scattering from ^{8}B neutrinos is particularly problematic as it can mimic a 6 GeV/c^2 WIMP [115]. The presence of such a background is a physics signal in its own right, and would be the first observation of solar neutrino coherent scattering [61]. Additionally, such a signal is good validation of the analysis techniques close to the detector threshold. The rate of scattering is sufficiently low below about 10 keV that the expected rate is very low. Future larger detectors will have to contend with this extra background which is commonly referred to as the neutrino floor [60].

3.4.2 Trace radioactivity in detector materials

A component of the background arises from decays of radioisotopes within the detector materials (Det). Extensive methods were used to obtain ultra-pure titanium [61] and to obtain low radioactivity values on the photomultiplier tubes. Additionally, other components were screened to detect low levels of ^{40}K, ^{137}Cs, ^{60}Co as well as isotopes of uranium and thorium. Despite these extensive efforts, very small quantities of such isotopes remain. These isotopes can undergo various decays and eject particles into the xenon volumes either directly or via subsequent (α,n) processes. Of particular concern are species that emit neutrons that can mimic WIMP interactions. Increased contamination within photomultiplier, Teflon or titanium components means that more of the xenon volume has to be excluded from the analysis due to a high event rate, thus impacting the sensitivity performance. Of additional concern is that some daughter particles can be ejected or can emanate into the main xenon detector and become homogeneously dispersed. Once this happens decay modes that have a 'naked' beta emission are particularly hard to remove in data analysis. Such decays are present within both the ^{220}Rn and ^{222}Rn decay chains when a fraction of lead isotopes decay to bismuth [61]. As such particles are continuously produced then they cannot easily be removed.

The xenon itself contains further trace elements which are radioactive. Specifically this includes ^{39}Ar, ^{136}Xe, and ^{85}Kr. Due to these isotopes also being noble gasses they are not removed by the getter. The elements are largely removed from the xenon during a pre-running purification using charcoal chromatography [61]. This leaves the concentration of ^{85}Kr at 0.015 parts per billion g/g [61]. The beta decay of these elements leads to electronic recoil events. ^{136}Xe is also present within the detector and has a double beta decay mode which contributes to the electronic recoil background.

Additional sources of backgrounds include a contribution from the cavern walls (Env). Radioisotopes in the rock contribute to the electronic background from the emission of gamma rays [116] and for the nuclear recoil background from neutrons. The outer detector and water tank reduces much of this background by 6 orders of magnitude [112].

3.4.3 Surface contamination and dust

Additionally, the environment in which detector components are stored, constructed and assembled can impact the detector backgrounds through surface contamination (Sur). This contribution comes from two areas. Firstly the exposure to air containing a high proportion of ^{222}Rn, which the progeny of the decay can plate out onto surfaces and remain there. Once the detector is running these progeny can become mobile within the detector medium or can decay on the walls leading to additional backgrounds. To mitigate this the detector is constructed in a reduced radon atmosphere and parts are purged in nitrogen gas. Components are bagged within nylon, where possible, which has been shown to reduce the radon exposure of components [117]. Secondly dust will add to the backgrounds due to the presence of radioisotopes. LZ has a requirement to have less than 500 ng/cm^2 within the inner parts of the experiment. The detector was constructed within a clean room of better than class 1000. These points will be considered in more detail within Chapter 4 when the considerations are made for the LZ thermometry system.

3.4.4 Additional sources

There are additionally some event topologies of interactions that can lead to additional backgrounds. One example is from the presence of scatters in the reverse field region. The reverse field region does not allow for the collection of ionisation electrons. Therefore, a multiple scatter event of a gamma ray with one interaction within this region can have an uncharacteristic S1/S2 ratio and hence the event could look more like a nuclear recoil event. Studies are ongoing to reduce this type of background. Backgrounds can also occur from the reconstruction of fake events. This is the process where S1 and S2 signals are paired together that are uncorrelated. A lone S1 signal can be produced by the pileup of PMT dark counts or Cherenkov events within the PMT quartz windows [112]. It is also possible in some cases to produce an S2 signal without an associated S1. The lone S1 and S2 pair can therefore together form a fake event. The ‘S2-only’ events are discussed as part of a separate analysis in Chapter 6.

3.5 Experimental strategy

3.5.1 Overview

In order to detect Dark Matter, a number of candidate events have to be observed. The number measured will be compared to the number expected. Prior to detector turn-on, it is predicted that about 6.49 events will be observed within the parameter space where WIMP interactions are expected from the backgrounds discussed above [112]. This expectation comes from the extensive assay measurements of detector components [118], tape lifts to measure dust exposure and from simulation. By careful material procurement and clean-room operation, the initial predicted number of events is reduced to this low level. In addition, astrophysical backgrounds are approximated using available data based on previous experiments and modelling. The number is likely to be refined when studies of backgrounds are performed. Observation of rarer Dark Matter scattering, producing single numbers of events is only possible if the number of irreducible background scatters is kept to this low level. In order to do this requires extensive analysis and calibration. Many event topologies can be separated from possible signal and background events using data analysis cuts.

3.5.2 Data analysis cuts

The analysis to remove background events is a multi-step process. Firstly an event in the xenon has to be reconstructed. The LZ analysis package (LZap) performs this function consisting of the identification of pulses, their classification and parametrisation. There are inputs from calibration at this stage to define the energy reconstruction and remove position dependence of signal size. Additionally, analysis is performed to standardise the response to single photons for each photomultiplier tube.

Once the pulses are identified then an appropriate event classification can be allocated. This looks for the number of S1 and S2 signals and their relative sizes. Multiple prominent S2 pulses are indicative of multiple iterations, therefore incompatible with a WIMP

scattering. These events can be hence removed from the analysis. This is also extended to include events that have energy deposits within the veto detectors coincident in time with the TPC.

Secondly, the self shielding of the xenon means that the background rises close to the walls of the TPC. It is therefore common to consider only events that are within a smaller region of the detector than the geometrical walls. In order to do this the position of the events must be reconstructed. As explained in Section 3.1.2 the event depth is estimated using the time difference between the start of the S1 and S2 pulses. The position in the plane perpendicular to the drift is estimated based on the PMT hit pattern [105]. Events can then be removed that are within 4 cm of the TPC walls, 13.5 cm of the gate or 1.5 cm from the cathode. This then leaves a fiducial volume, from which events are considered.

Events can also be removed based on the signal size. Any S1 signal must have three-fold coincidence. This means that three PMTs must see a signal characteristic of at least a single photon within a pre-defined time window in order to pass the cut. This cut is to reduce the accidental pileup of dark counts from PMTs creating a fake S1 signals. This coincidence requirement is the driver on the lower energy threshold in the experiment. A minimum signal size from an S2 is decided to be at 420 phd which ensures adequate position reconstruction ability. This is discussed in more depth, where the choice of value is obtained in Section 6.2.5. The choice of these values will be dependent on the detector parameters measured during the analysis. Numbers here are included to provide an indication of the nominal expectations.

Additional data quality cuts could also be applied, such as to ensure the detector conditions are stable. For example, the level sensors could be used to remove events during times when the liquid level changes. Furthermore, the environment following a large S2 signal could have delayed extraction of electrons present. These electron-trains could require a temporal region following the S2 to be excluded from some analyses.

The remaining events following all data cuts are fed into the LZ Statistics framework as addressed later in this chapter.

3.5.3 Simulation

Simulations are used to validate the analysis methods, simulate backgrounds for a statistical analysis and examine detector efficiencies. LUX-ZEPLIN uses two simulation packages [119].

The first simulation package is a GEANT4 based simulation which models particle interactions within detector components in a simulated geometry. The simulation interfaces with the NEST [99] package which simulates the xenon physics as outlined in Section 3.1.1. The model is based on empirical data from various experiments. The quanta are propagated through until photons terminate on PMT cathodes. Many effects are included such as the varying reflectivity of detector components, the intricacies in the expected electric field configuration, and the expected position dependent effects within the experiment. The simulation also includes interactions within the outer detector and xenon skin. Once light has reached photocathodes, then the electronics response is also modelled using a paramaterised model based on the analogue and digital electronics. This simulates the response of the single photon through the PMT, the subsequent amplification and digitisation in the front end electronics. The output files from this simulation are in the same format that data takes, and hence the simulation output can be interfaced with the analysis framework LZap.

The second simulation framework parametrises much of the xenon response. Again using NEST, this simulation takes in an expected energy spectrum of recoils. This is either calculated analytically or with a full simulation as detailed above. Given a specified energy deposit, or source, the simulation produces the quanta and then propogates this through to obtain the detected signals. The conversion from the quanta to the S1 and S2 photons detected is done using different g_1 and g_2 values depending on how the detector is set up. These factors are efficiencies of detecting S1 and S2 light respectively, as used for energy reconstruction, introduced in Equation 3.17.

3.5.4 LZ Sensitivity to WIMP interactions

Using the simulations as detailed in the previous Section, the expected capabilities of the experiment can be estimated. In order to do this, background and WIMP scattering events are simulated using the parametrised model detailed in the previous Section. Following data analysis cuts, the probability of finding events within different regions of the S1-S2 signal size parameter space is determined. These distributions provide the foundations to the profile likelihood ratio (PLR) method, used in the statistical analysis. This is examined in more detail within Section 6.2.6 where a consideration of the same framework is applied to a low mass Dark Matter search. The method calculates the statistical likelihood for a given mass when the cross section is varied. Many toy experiments are generated to enable the median sensitivity to be calculated. These toy experiments can capture the systematic uncertainties in the rates of the different backgrounds. Based on this model the expected sensitivity is shown in Figure 3.10. This shows the expected sensitivity of the experiment at 40 GeV/c^2 is 1.6×10^{-46} cm^2. At low mass the sensitivity is much worse due to the energy threshold of the experiment when low numbers of quanta are produced. At larger masses the lower abundance means the rate will fall due to the lower number density. Here the assumption is made that the detector conditions are well understood and time varying phenomena can be accounted for. In producing this sensitivity, astrophysical parameters are assumed which are detailed in Section 2.1.3.

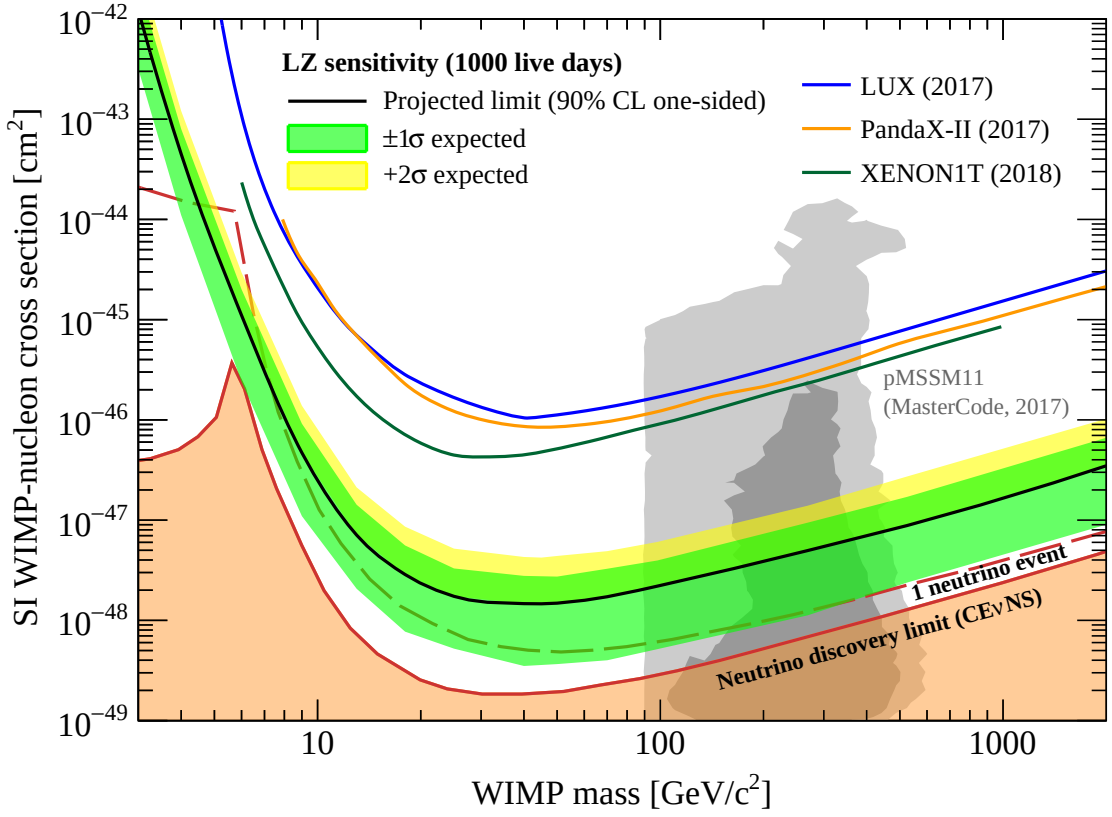

Figure 3.10: Projected spin-independent WIMP-nucleon cross section sensitivity of LUX-ZEPLIN as a function of WIMP mass. This shows that LUX-ZEPLIN should achieve several orders of magnitude improvement on previous experiments which are also shown. alongside theoretically motivated regions for Dark Matter. Sensitivity lines at which this technology will more readily detect coherently scattering neutrinos are included which could limit use of this technology. Figure from [112].

3.6 Summary

The LZ detector is specifically designed to be able to detect the nuclear recoil scattering of WIMPs. LZ uses both the scintillation and ionisation channels, which aids in the position reconstruction and interaction type discrimination. The detector is constructed of multiple sub-detectors which have to work in unison to be able to achieve the nominal sensitivity. Achieving this also requires a good understanding of backgrounds, an extensive detector calibration and robust event reconstruction.

Chapter 4

Thermometry for the LUX-ZEPLIN experiment

This chapter discusses the thermometry system for the LZ detector, including a consideration of how the construction and calibration procedures of the thermometers meet LZ requirements. This is necessary for LZ to be able to achieve the nominal goals for Dark Matter searches, both to ensure the integrity of the detector during cool-down and stability during data collection. The thermometry system also needs to be robust for long periods and not significantly increase the activity rates from intrinsic radioisotopes. It is shown that the front-end electronics are capable of readout within the above specifications and the thermometers themselves can be calibrated to this level.

The LZ detector is fitted with an extensive system of thermometry, consisting of 90 thermometers that have stringent radio-purity and cleanliness requirements, necessitating a bespoke solution. They are based on a PT100 resistor. The locations where the bespoke thermometry system is mandatory are illustrated in Figure 4.1. There are 18 thermometers located within the liquid xenon, below the bottom PMT array, used to measure the temperature of the xenon flowing into the detector. During normal operation, the ICV is filled with liquid xenon up to a weir, with the region above in xenon gas. Liquid flows over the weir and enters a series of purification and cooling steps before recirculating into the detector. The weir region is illustrated in Figure 4.2. A further 18 thermometers are located in the gas region to measure temperature profiles there. A series of 36 thermometers sit on the outside of the ICV within the vacuum space that are crucial for monitoring the temperature while filling, when it is necessary to have a detailed knowledge of the detector state and for temperature profiling once cold. They are separated vertically by approximately 30 cm for this purpose. When filled, these sensors operate close to the nominal liquid temperature of 175 K depending on the pressure and phase of the xenon

in the vicinity of the thermometer [61]. A further 18 thermometers are located on thermosyphon heads, which sit on the outside of the ICV, close to the main flange near the top (see Figure 4.1)) and form part of the system part that provides cooling power to the detector and establishes a controlled thermal environment for the xenon inside the ICV. This system uses liquid nitrogen ($\sim$ 77 K) to establish the $\sim$ 176 K inside the ICV and hence, thermometers here operate between these temperatures.

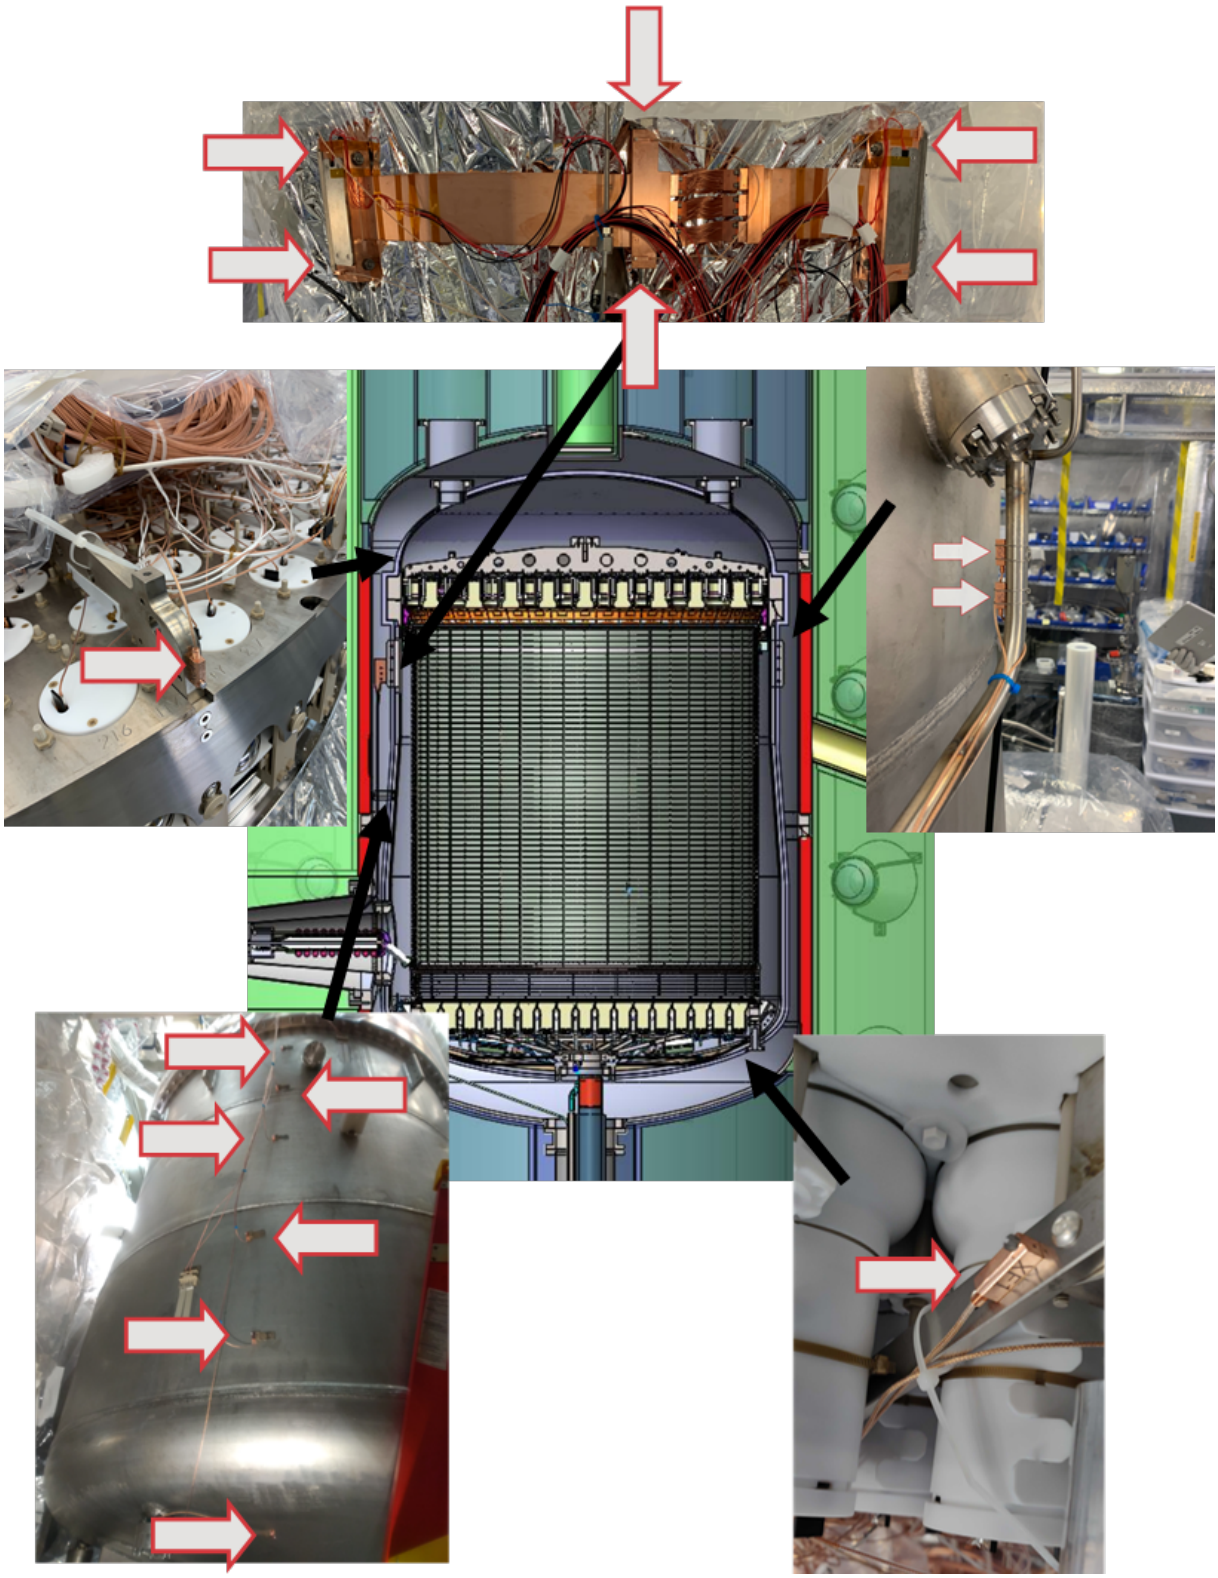

Figure 4.1: Images of different thermometer locations, identified with arrows: (top) on one of three thermosyphon evaporators, (centre left) on the edge of the top PMT array in gas, (centre right) on one of the weir drain lines, (bottom left) thermometer 'ladder' on the outside of the ICV within the vacuum space, (bottom right) back side of the bottom PMT array in liquid.

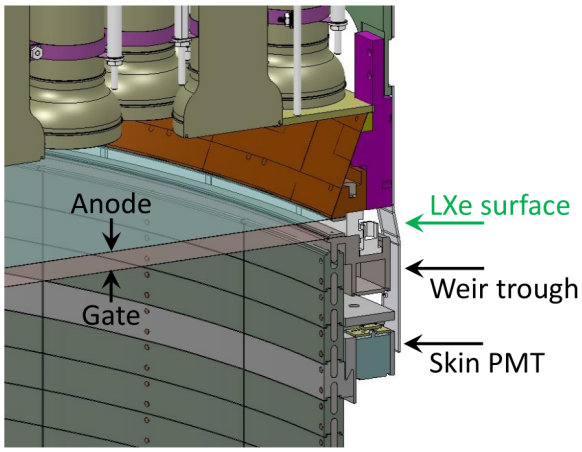

Figure 4.2: Diagram of the weir and electroluminescence region (between liquid level and Anode). Liquid spills over into the weirs and flows down tubes on the outside of the ICV in the weir drain lines (not shown), and enters the circulation and purification system. The geometry is complex in this region, with lots of features, the Anode and Gate grids and the weir are within 20 mm of each other. Figure from [61].

4.1 Thermometry requirements

The LZ thermometers have accuracy and precision requirements of 1 K and 0.2 K respectively. These are driven by experience from previous experiments as well as a consideration of the roles in which the thermometers have to undertake.

4.1.1 Stable running

The detector operates at a temperature of around 175 K and a gas pressure of about 1.8 bar. Motivations behind these values lie in the signal generation and properties of the xenon. The gas pressure principally affects the generation of the S2 signal. Figure 4.3 shows that at this temperature and pressure, an expected $\sim$ 800 photons per electron are produced at the nominal 11.5 kV anode-gate potential difference. Due to light paths not terminating on a photocathode and a PMT quantum efficiency of $30.9 \pm 2.5\%$ [93], most photons are never detected. Indeed, a pessimistic scenario of 5% for the photon detection efficiency yields only 40 detected photons. This is considered to be the limit in achieving close to 100% single electron detection efficiency and adequate position resolution on the events [61]. This minimum anode-gate voltage is hence necessary to meet this requirement.

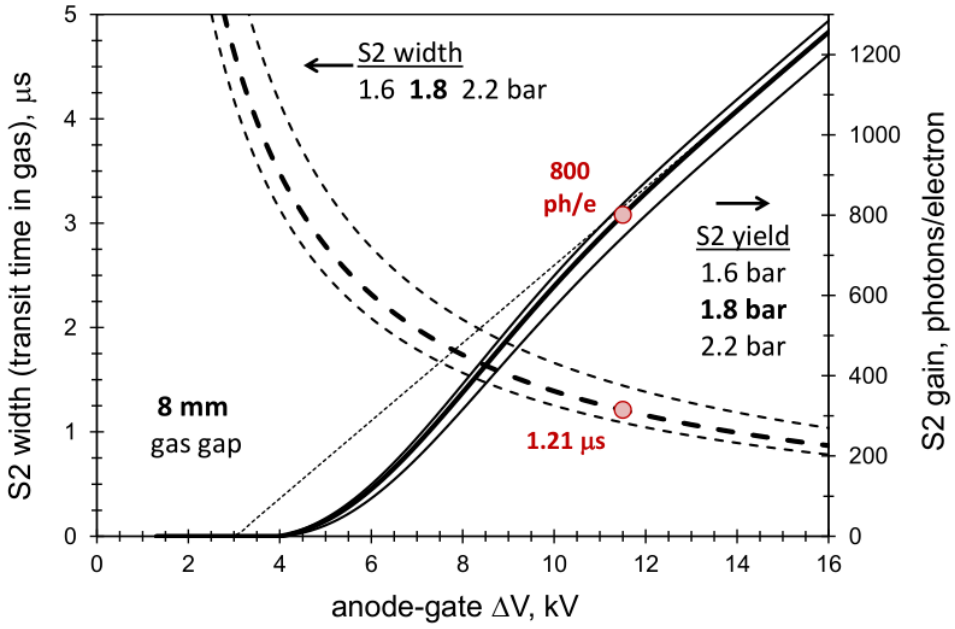

Figure 4.3: Variation of the S2 width (transit time in gas) and amplification (number of photons per electron) as a function of the gate-anode potential difference for an 8 mm gas gap. Lines show the variation for different pressures. Figure from [61].

Higher anode-gate voltages also have the desirable feature that the transit time of electrons in the gas is smaller, which is required for a good position resolution along the vertical axis of the detector. Achieving these high voltages is challenging, especially when increased field gradients lead to the enhanced risk of electrostatic breakdown, alongside the increased probability of spurious emission. This is particularly true when considering the complex geometry as illustrated in Figure 4.2.

The number of photons produced per unit distance within the gaseous part of the extraction region is given by [30][122]

$$\frac{\mathrm{dN}_{ph}}{\mathrm{d}x} = \alpha E - \beta P - \gamma, \tag{4.1}$$

where E is the electric field strength and P the pressure. α, β and γ are experimentally determined to be approximately ~ 0.137 /V, ~ 177 /bar/cm and ~ 45.7 /cm respectively. Some variations in these parameters have been observed as reference [122] summarises. The effect of pressure in both cases is largely driven by the mean free path of the electrons, which is larger when the pressure is lower. The pressure cannot be too high, as this will reduce the signal size. Conversely, a pressure which is too low would lead to more variation

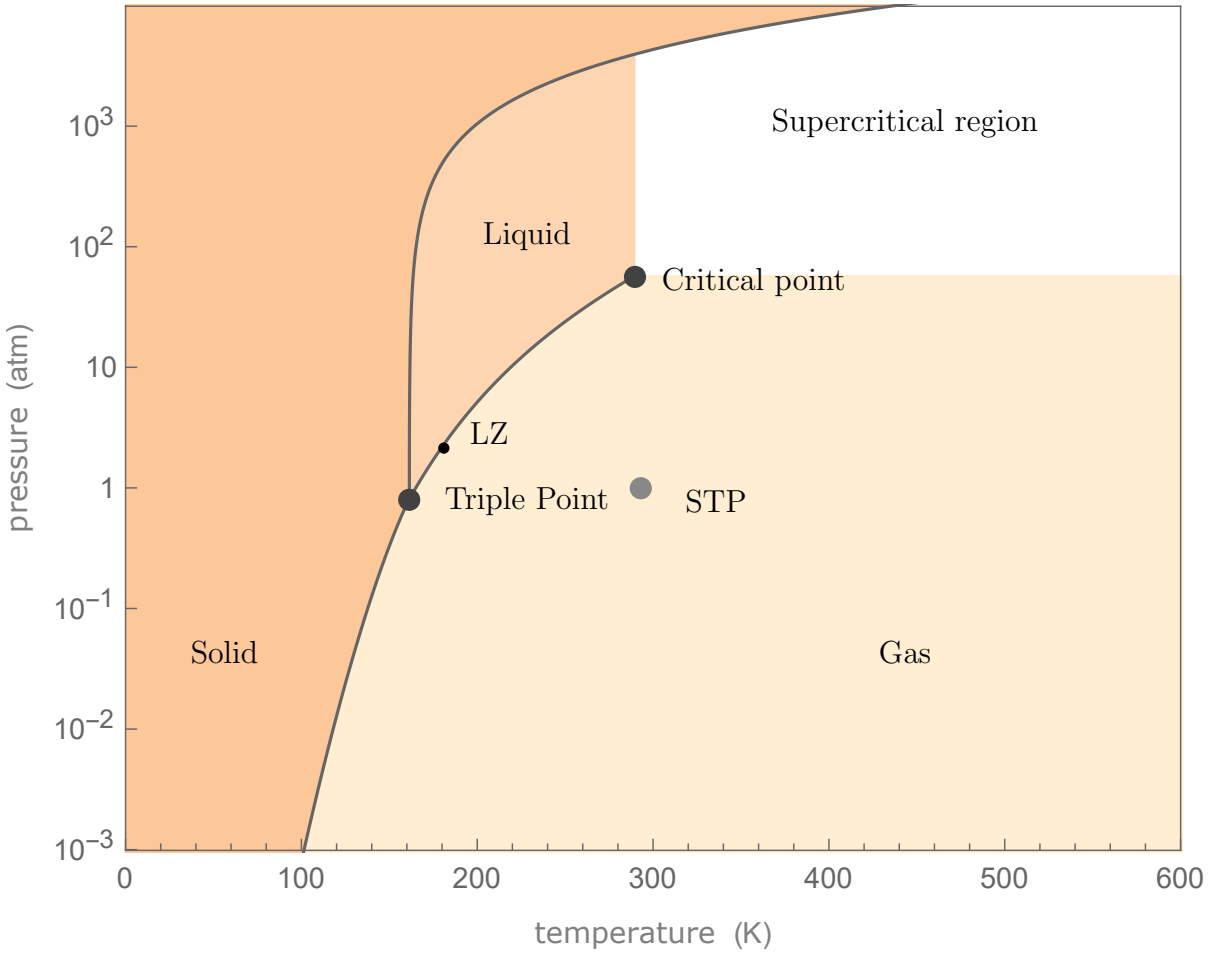

Figure 4.4: Phase diagram with the approximate nominal conditions for LZ included. Importantly, LZ has operating conditions in a position in the phase space which means that gaseous and liquid xenon can co-exist. The STP is the phase at standard temperature and pressure. Amended from a plot produced using the Mathematica package [120]. This Figure uses data from [121].

in light production due to Townsend amplification [30][123]. This is the process where charge multiplication occurs, particularly close to the anode wires, where the electric field is strongest. Due to the position on the xenon phase diagram (see Figure 4.4), lower pressures would place the conditions closer to the triple point and make the detector more marginal to freezing. Therefore, during cooldown, a gas pressure is chosen in the nominal range of 1.6 to 2.2 bar. With the detector state on the saturation curve, the design gas pressure of 1.8 bar, requires a temperature of 175 K at the liquid-gas interface. Lower in the detector, where colder xenon enters, hydrostatic pressure from the fluid overburden ensures the liquid state is maintained. For normal operation, the cooling is controlled to maintain an optimal flow of liquid xenon overspilling the weirs and circulating through the system. The allowed range of pressure, between 1.6 and 2.2 bar gives a temperature range of 6 K between about 174 K and 180 K [124]. Therefore, knowledge of the temperature is important to ~ 1 K accuracy.

4.1.2 Cool-down

Most thermometers have their main use during the cool-down procedure when significant work is needed to fine-tune the circulation and cooling systems. Experience with previous xenon detectors and the test setup shows that an accuracy of 1 K was sufficient to control the thermal profile of the xenon space during this time. When more stable, greater emphasis is placed on temperature changes, as this establishes even finer tuning of cooling parameters. This enables the dampening of oscillations in temperature and pressure to be controlled. Hence, sub-kelvin precision is required.

Some more nuanced uses of thermometry occur during the cool-down process when detector parts contract. For example, the height of the TPC reduces by several mm when cooling by more than 100 K from room temperature. Combined measurements of temperature and data from sensors, which sit on top of the TPC are used to measure this contraction of the top PMT array relative to the ICV. These sensors form a device known as the hexapod and are discussed in Chapter 5. Vacuum space sensors complete

this picture by providing information about the contraction of the titanium ICV. In combination, this provides information as to the true absolute change in the height which is important when calibrating true vertical position to drift time. These particular sensors also provide essential knowledge of thermal gradients which could create unnecessary stresses in the material. This is a particular concern on the numerous flanges that are part of the vessel. Thermometers below the bottom array can also provide an indication of the liquid level in the regions at the base of the detector during filling. This can provide a crude estimate where the geometry is complex and hence the mapping between xenon mass and height is non-trivial. Additionally, sensors that sit near the xenon inlet can examine the flow into the detector to measure the heat load from the transfer of xenon from the xenon recirculation system to the detector.

Once stable conditions are achieved and WIMP search data collects, the thermometers are used to monitor the condition within the detector and the thermosyphons continue to control the cooling through the flow of liquid nitrogen.

4.1.3 Accuracy, precision and radiopurity requirements

The LZ thermometers have formal associated requirements of 1 K accuracy and 0.2 K precision. As mentioned previously, the numerical values for these requirements are based on experience from previous successful experiments. Specifically, the accuracy requirement is to identify thermal gradients across the ICV during cooling and inform when condensing is occurring. To finer tune cooling parameters and maintain detector stability during running requires sub-kelvin precision to be able to identify temperature trends.

Radiopurity of the thermometers is also highly important and also has associated requirements. As introduced in Chapter 3 the WIMP search operates in an energy regime of $\sim$ keV recoils, meaning radiation from decays of material components produces a detectable background in the xenon. Hence, pre-emptively, materials were procured which have been confirmed to have low radioactive content. This was validated through an extensive scheme of germanium screening and radon emanation. Isotopes of uranium

(^{238}U,^{235}U) and thorium (^{232}Th) and their daughters, in particular, emit various species until they reach stable lead isotopes and are major contributors to the experiment background. Additionally, (α, n) reactions and decays which produce neutrons, contribute highly to the background since nuclear recoils of neutrons can mimic WIMPs, and are hard to separate in some interaction topologies. The gamma decays of isotopes of potassium, caesium and cobalt contribute to the electronic recoil background [125].

Requirements also exist to minimise the surface contamination of detector components. Radon and dust are the two primary concerns. Radon daughters in the air can plate out to surfaces and give a ^{212}Pb background. Components must have $< 10,000$ days of exposure to air of 10 Bq/m^3, however, this should be minimised wherever possible. Dust naturally contains radioactive contaminants, thus prevention of exposure to dust is important. This motivates the need to perform the construction in the clean room. The LZ requirement is for < 500 ng/cm^2, which can be validated using UV light inspections, examining components under a microscope and estimating with tape lifts.

Additionally, the sensors should be reliable and not have a high failure rate. This is especially true as these sensors are inaccessible for the several year operation of LZ so cannot be replaced.

4.2 Design and Construction

The principal design of the thermometers and accompanying electronics are driven by the requirements as laid out in Section 4.1.3.

4.2.1 Design and material procurement

The package of LZ thermometers is illustrated in Figure 4.5. The sensor itself is a PT100 [126], which has a nominal temperature coefficient of 0.385 Ω/K around 0 °C. Importantly, these sensors only guarantee this linearity down to about −30 °C, so there is a need to correct for any manufacturing differences in calibration below this point. The outer shells

are made from copper to enable good thermal conductivity and ensure direct response to temperature changes. The cable consists of 4 interwoven wires surrounded by the conductive sheath which also carries the electrical ground of the sensors. At the warm end of the cable, on the breakout box, the cable is split into the four wires, which allows for the interface with the Oxford electronics as described in the next section. To ensure the device measures the temperature at the thermometer head, heat ingress needs to be estimated. Typical copper conductivity is 400 W/m/K [127] which is contrasted with the lower values from the phosphor-bronze cable sheath to reduce heat ingress from the wiring. For this cable the thermal conductivity of ~ 10 W/m/K was estimated using the Wiedemann–Franz law based on measurements of the electrical resistance of the cable sheath [128]. The contribution of the heat is given by

$$Q = -kA\frac{\Delta T}{L}, \tag{4.2}$$

where k is the linear conduction coefficient, A is the cross-sectional area, and ΔT is the difference in temperature over distance L. Hence, the contribution from the cable is suppressed by both the conductivity, the cross-sectional area and the distance between the PT100 and room temperature being $\sim$ 10 m. Arguments also apply to the inner four wires, which have higher resistivity and smaller cross-sectional area. Therefore, the relative contribution from the cable is negligible. To maintain good thermal contact with the base material, partially threaded stainless steel screws are used to ensure the copper can be forced onto to the surface where the temperature is to be measured. A copper washer and screw are used to provide strain relief (see Figure 4.5), to prevent the delicate cables from being pulled out and snapping during assembly.

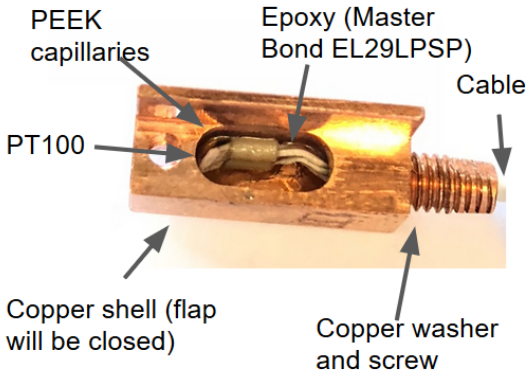

Figure 4.5: The LZ thermometer head. The PEEK capillaries and epoxy ensure that the insulating sleeve does not move, which could cause an electrical short. Each thermometer has $\sim$ 10 m of cable which connects them from the location on the detector to the breakout boxes. The stainless steel screw to affix the thermometer to the surface is not shown here.

Copper also has a high radiopurity, so makes a good choice of material. The cable is wrapped in Teflon and so had to be cleaned carefully to remove dust which affixes through electrostatic charging effects. Final cleaning on site was done under de-ionising fans to prevent this. These fans produce a constant stream of ionised air, which neutralises any charges on the plastic, thus preventing dust accumulation. A summary of the germanium screening for the final design is shown in Table 4.1. Approximately 0.5 kg of thermometers are installed within the detector. The main contributing factor in this table is the ceramic on the PT100. Due to the relatively small quantity inside the detector and the fact that they are far from the main fiducial volume, this is deemed satisfactory. A Monte Carlo simulation confirmed that the estimated rate in LZ for the standard 6 keV_{NR} to 30 keV_{NR} energy range is less than 1/100th of the total background for both nuclear and electronic recoils. The cable, considered separately, is less problematic for these backgrounds due to the increased distance from the active TPC region for most of the cable length. Additionally, based on measurements of similar cable the rate per kg is lower than the thermometer head.

Source	U early	U late	Th early	Th late	Co60	K40
Activity (mBq/kg)	43.46	17.24	6.82	8.87	3.06	482.67

Table 4.1: Activities for the completed thermometer enclosure excluding the cable. Contributions are dominated by the ceramic on the PT100.

4.2.2 Construction procedures and protocols

The construction of the sensors was performed in the clean room in the Denys Wilkinson Building, Oxford. Prior to work commencing, measurements showed that the radon levels were $\sim$ 22 Bq/m^3 [129] whilst dust assays verified that dust levels were around that expected of a class 1000 clean room [130]. Final cleaning was done at Imperial College in London which has lower variability in dust levels and use of the space was solely for LZ. At all stages, full local and LZ clean room procedures were followed.

All parts went through the rigorous LZ cleaning procedure which involved various steps. This included wiping parts and cleaning in solvents. This was interspersed by the LZ quality assurance procedure which involved electrical checks, UV inspections, and examination under a microscope. Parts were bagged when not in use to reduce radon exposure. Radon levels will reduce naturally in the bag once the components have been sealed and hence reduce this burden [117].

4.3 Electronics

Custom built electronics are used to readout the thermometers. In order to avoid cabling and additional connectors in the underground space, a dedicated front-end readout card has been designed. These readout cards plug directly into the feed-through flanges which separate the xenon/vacuum space from the cavern.

Design and precision

The ability to meet the accuracy and precision requirements for the thermometers drive the design of the electronics used to read out the resistances of the thermometers.

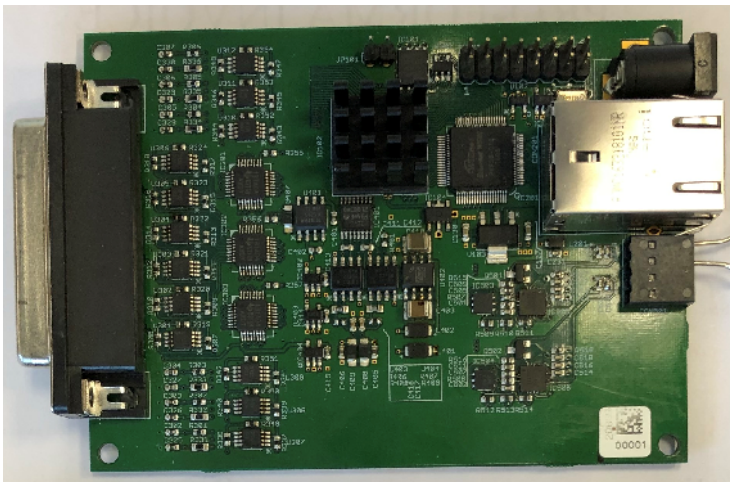

Figure 4.6: A prototype Oxford 6-channel readout card which performs 16-bit true four-wire measurement for the LZ thermometers. To the right is the RJ45 port which communicates with the slow control system and to the left is the DB25 which interfaces with the flange where the thermometers are connected.

The Oxford 6-channel frontend readout card, as shown in Figure 4.6, is capable of reading 6 channels of true four-wire resistor measurements to 16-bit precision. Four-wire readout has two wires carrying the current and two wires connected to the voltmeter. This removes the dependence on the internal resistance of the wires as well as those effects affecting individual cables, such as temperature dependence or the electromagnetic environment. This gives accuracy benefits when compared with some commercially available systems that use three-wire readout. An illustration of the internal electronics described above is shown in Figure 4.7. The voltage drops over two resistors (R_{PT100}

and the 49.9 Ω reference resistor on the frontend electronics) are individually measured in a differential configuration and amplified by a factor of 21. The ADC measuring the reference resistor (49.9 Ω) provides a measure of the current flowing through the PT100. In conjunction with the measurement of the voltage drop over the PT100,the resistance can be obtained using Ohm's law. The operating current is very small ($\sim$ 250 µA) and hence the Joule heating on the PT100 resistor is negligible.

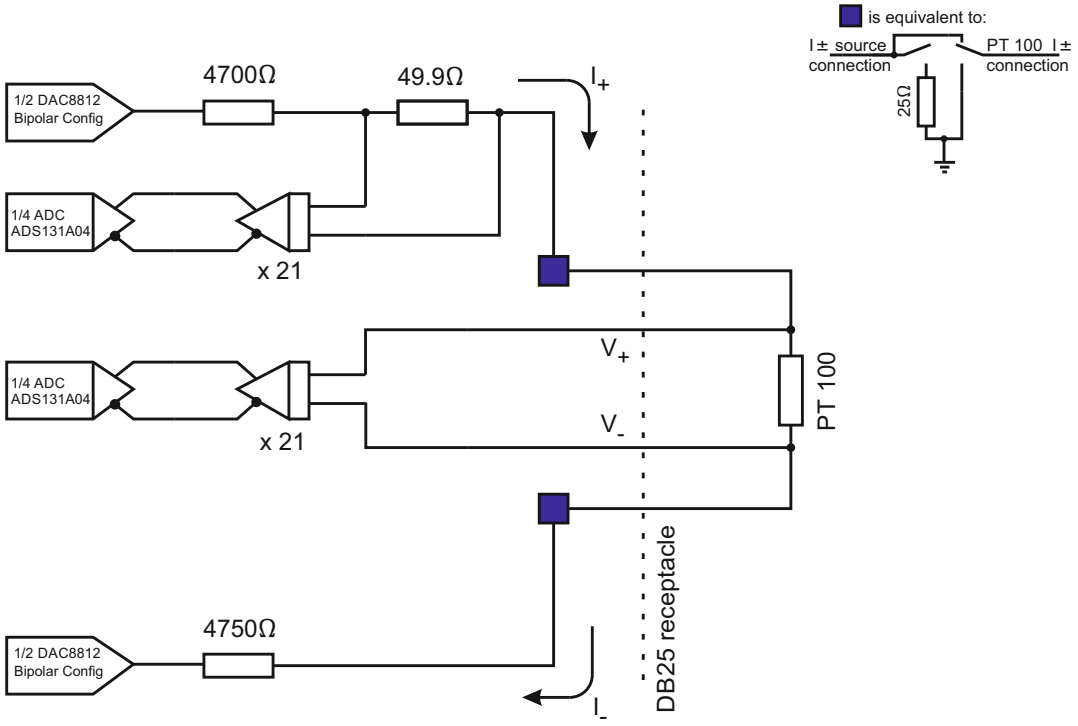

Figure 4.7: Schematic diagram of the front-end electronics card used for readout of the resistance values on the PT100 thermometers. Blue boxes represent the sub circuit, shown on top right which includes the ability to turn off individual channels through the MODBUS interface.

The excitation current is provided through two DACs which have a settling time of less than 0.5 µs. Both of the DACs must have a voltage that sums to zero to set up a square wave voltage source which transitions between ± 1250 mV, as is illustrated in Figure 4.8. A sine wave transition has been used to soften the the transients in the square wave, to remove effects due to higher-order Fourier modes. By ensuring that the excitations are ±1250 mV and they are in phase, removes systematics from the polarity

such as from thermopower. The frontend card also has an electronic switch to enable the individual channels to be electrically shorted to ground, as shown in Figure 4.7. This prevents thermometers acting like antennae of electrical interference in the event of the failure of one wire.

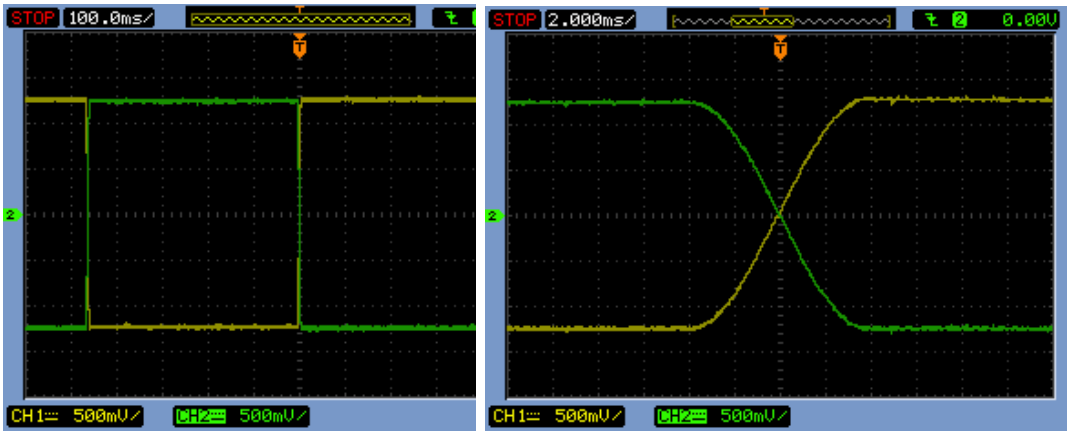

Figure 4.8: Oscilloscope trace for the two DAC outputs (Figure 4.7) showing a full cycle of the order 1 s to the left and the transition on the right which takes approximately 8 ms. The two voltages of ±1250 mV are centred around a voltage of zero.

To calibrate the frontend electronics, five test resistors (known to 0.1% accuracy from manufacturer specifications) were used. The values were checked using the frontend readout card post calibration and conform to the accuracy requirements needed.

Manufacturer's Value	10Ω	20Ω	50Ω	60Ω	100Ω	Average
Channel 0	1.25	1.69	2.25	2.11	2.99	2.06
Channel 1	1.46	1.61	2.09	2.17	3.06	2.08
Channel 2	0.97	1.59	1.85	1.97	2.94	1.86
Channel 3	1.33	1.52	1.90	2.11	3.13	2.00
Channel 4	1.15	1.49	2.18	2.34	2.69	1.97
Channel 5	1.17	1.70	1.96	2.29	3.22	2.07

Table 4.2: Standard deviation values in mΩ for the resistors that were each sampled 2000 times. The values are substantially less than the 0.04 Ω that would be ∼ 0.1 K variation (assuming the 0.385 Ω/K).

Using the same resistors enables the precision of the electronics to be measured when the resistance is not changing. This enables the separation of thermal effects of thermometers from those of the electronics. Data for 2,000 samples were examined, and as expected, gave resistance values around a mean value in a Gaussian distribution. Results for the standard deviation of each resistor per channel are shown in Table 4.2. Impor-

tantly though the standard deviations are considerably better than the required precision. Indeed, assuming Gaussian behaviour, a 5 sigma discrepancy on any one value would not even register as a temperature difference of 0.1 K at LXe temperature (assuming the 0.385 Ω/K conversion factor) where $R_{PT100} \approx 60$ Ω. There is some channel to channel variation but even in the worst case, the specification would still be met comfortably.

Temperature dependence of electronics

In designing the readout electronics, care had to be taken in ensuring that the frontend readout card measures the resistance and that any contributions from the thermal environment in which the board operates are sub-dominant. Semi-conductor devices exhibit inherent temperature dependence and so expert electronics design is required to minimize such effects. This can be done by procuring components that have a small temperature dependence or designing the circuits so effects due to temperature cancel out. To quantify this effect for the thermometer measurement, the frontend electronics operating temperature was varied within a temperature-controlled chamber, between −10 °C and +30 °C. Data was taken every 4 s in cycles where the temperature rose and fell between the two limits over several hours in a profile as shown in Figure 4.9. Test resistors, mimicking thermometers, were placed outside of the test chamber, via ribbon cable in a space in which the temperature did not change. Channel 0 was used to monitor the temperature of the board through a PT100 sensor, which was attached to the board with heat-conducting paste. The output from this channel and the temperature to which the chamber was set, is shown in Figure 4.9. The temperature profile of one of the channels is displayed in Figure 4.10. Again the value of 0.385 Ω/K has been used to convert to temperature. Values of the gradient varied between 0.21 mK/K and 1.8 mK/K. Even the larger variation with temperature is acceptable, considering an unrealistically large departure of 10 K cavern temperature would only produce a variation of $\sim$ 0.02 K. This is negligible, especially when compared with requirements.

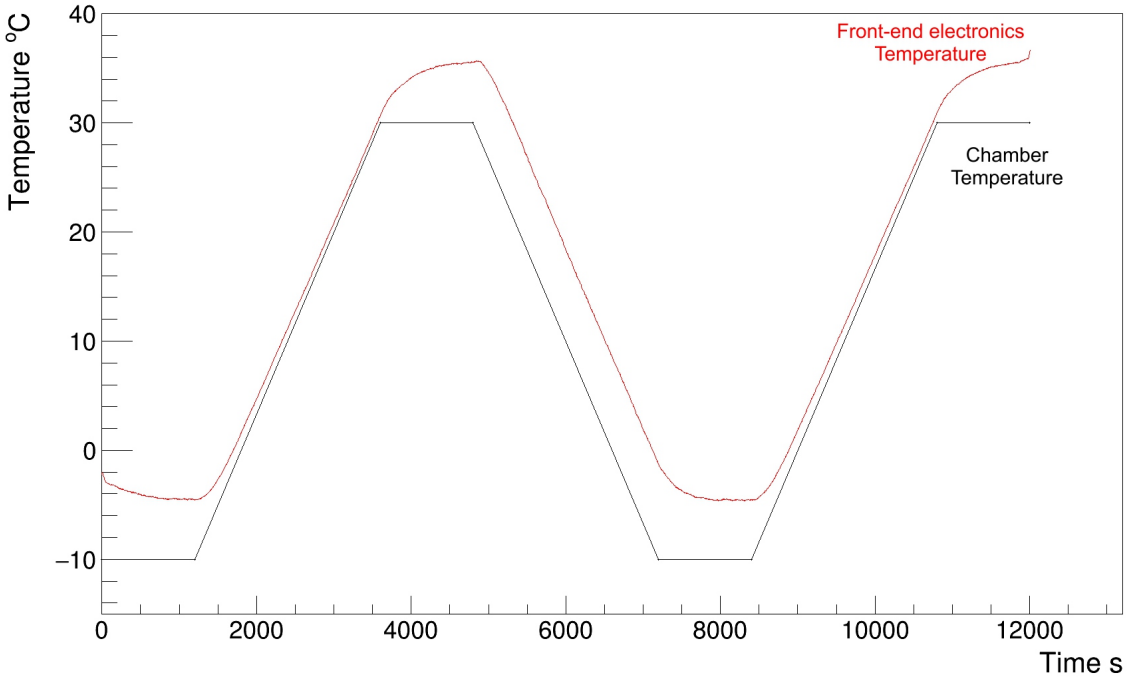

Figure 4.9: Temperature profile of the chamber (black) and the PT100 affixed to the board (red). The electronics are warmer than the ambient room temperature with a characteristic lag of a few minutes. For the board temperature a conversion factor of 0.385 Ω/K is used.

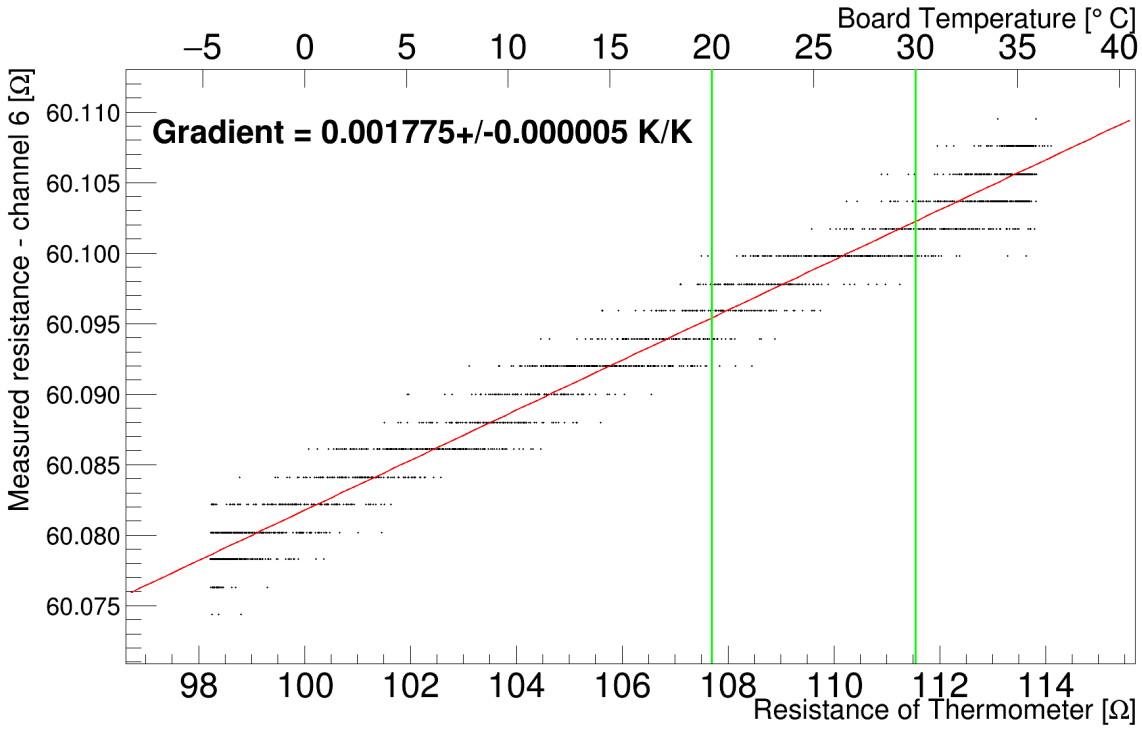

Figure 4.10: Values from the thermometer frontend readout board to examine the correlation between a PT100 thermometer mounted on that board to measure its temperature and a resistor on the outside of the chamber. A conversion factor of 0.385 Ω/K has been used to convert between resistance and temperature. Green vertical lines illustrate the extreme values that are expected from the ambient conditions. The minimum and maximum values here form an arbitrary range for the test and do not correspond to any physical limit for the electronics.

4.4 Calibration procedures

4.4.1 Setup

An algorithm using successive PID (Proportional-Integral-Differential) controllers was used to do the calibration in a semi-autonomous way between $\sim$ 150 K and $\sim$ 320 K. This is essentially a system whereby the power output is dependent on the temperature difference from the target (the proportional part), the total differences accumulated (integral part) and the rate of change of the temperature in a time step (differential part). The temperature was measured for all runs with a reference thermometer from Lake Shore with accuracy better than 25 mK across the temperature range [131].

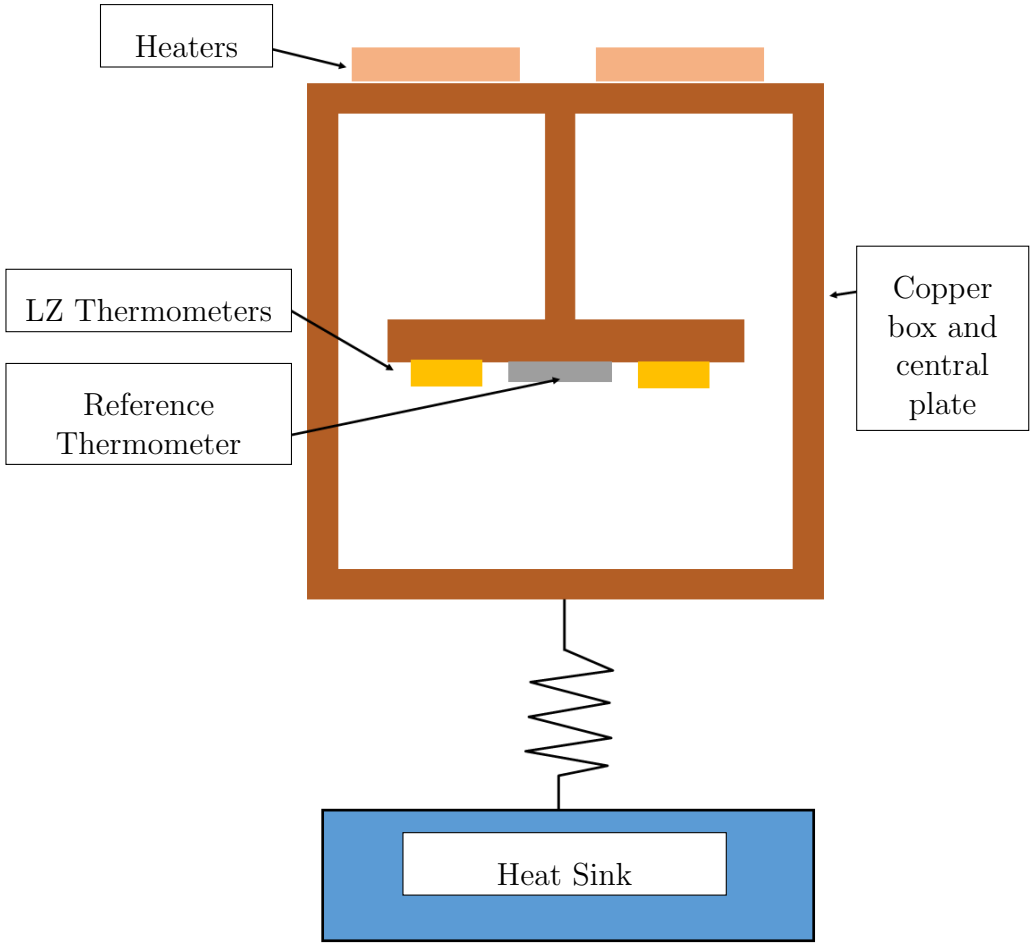

Figure 4.11: Diagram to illustrate the calibration setup. The heaters were controlled by the PID controller. The copper box was connected to the heat sink through copper braids by suspending over liquid nitrogen.

A copper cylinder was suspended above a liquid nitrogen heat sink. Copper braids

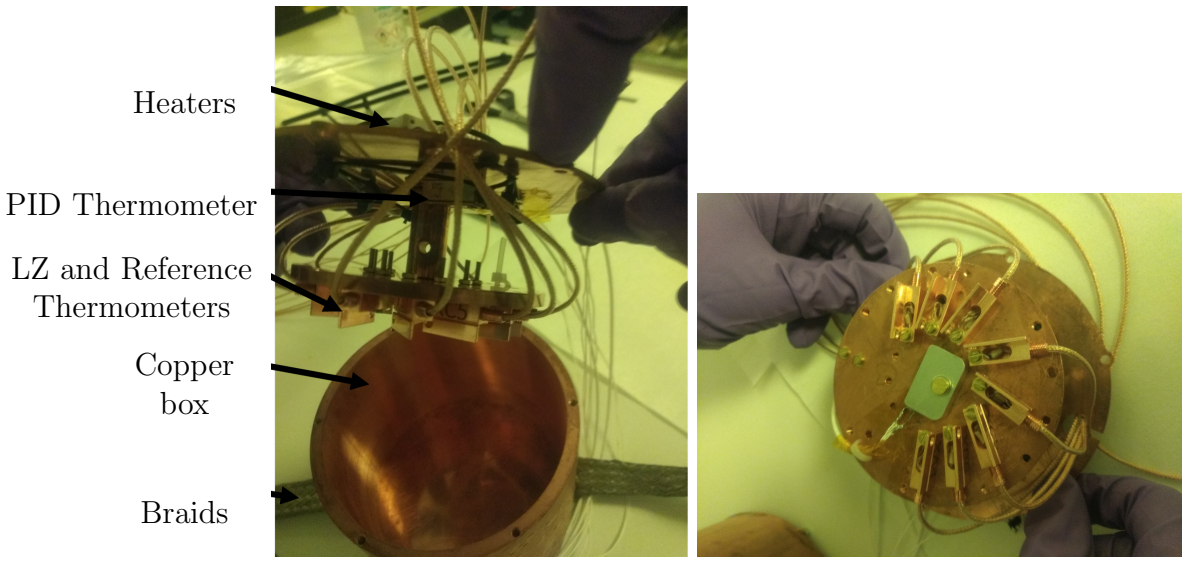

Figure 4.12: Images to illustrate the setup that was suspended over liquid nitrogen to calibrate the thermometers. The braids dangled into the liquid nitrogen to couple to the heat sink. This is illustrated in Figure 4.11.

formed a thermal link between the cylinder and the nitrogen to provide cooling. The temperature was controlled by using heaters that were mounted on the top of the cylinder. This is illustrated in Figures 4.11 and 4.12. Normally, the point at which heating and cooling of a structure occurs should be in close proximity. However, in the setup discussed here, this was not possible due to space restrictions. It has been verified that the contributions from radiative transfer between the top and bottom plates are negligible. Thermometers were then mounted on a middle plate which is thermally connected to the top plate via a copper rod. The difference in the temperature from the middle to the edge of the thermometer plate is small due to the high copper conductivity and the short distance of ~ 1 cm. This can transfer a few watts per kelvin through the copper to the thermometers and hence maintain equilibrium across the disk. Despite the larger temperature gradient in the cable the contribution is still sub-dominant in the measurement of the temperature. Contributions of further systematics, from the way the PT100 is coupled to the shell and the nitrogen atmosphere temperature surrounding the thermometer are difficult to quantify. These are likely to be sub-dominant because all PT100s were similarly mounted within the shells, and the copper box reduces the variation of the temperature of the nitrogen so should be similar for all thermometers.

The power output of the heaters was controlled to get temporal regions where the temperature was changing extremely slowly (< 1 K/5 min). This was done using an algorithm, which, through successive PID control steps, updated a power supply output every second. Hence, during these periods it was assumed that the plate where the thermometers were mounted was in thermal equilibrium. Mathematically the individual PID controller can be expressed as:

$$V_i = K_p(T_i - T_g) + K_i \sum_{i=0}^{i} (T_i - T_g)\Delta t + K_d \frac{(T_i - T_{i-1})}{\Delta t} + \text{bias}; \tag{4.3}$$

where K_p, K_i and K_d are coefficients, T_i is the temperature at step i, T_g is the 'target' or 'goal' temperature, and Δt is the time interval which for this setup is 1 s. The goal resistance (as a proxy for temperature) was incremented in 6 equal intervals of approximately 10 Ω. This enabled the fitting of a cubic polynomial with two remaining degrees of freedom. The parameters best describing the PID controller were determined, to be given by values described in Table 4.3. These values were obtained using a testing procedure in the lab before cleaning and re-assembling things inside the clean room.

	R<55	65<R<55	120<R<65
K_p	5	20	20
K_i	0.01	0.02	0.045
K_d	0.01	0.01	0

Table 4.3: Table of values for the PID controller, values of R are in Ohms.

A set of different PID parameter sets was required to allow the system to operate in such a wide variety of temperatures. Various exit conditions had to be added to the algorithm on top of the basic PID, when the device converged on a temperature that was stable but at some displacement from the original target. This was essential because the PID relaxation time was too large, and for this calibration, the absolute temperature value is less relevant. A significant overshoot, which was typically present in the first data point during cool-down runs, proved to be inevitable when allowing for enough cooling power to reduce the temperature to below that of liquid xenon. The power supply could

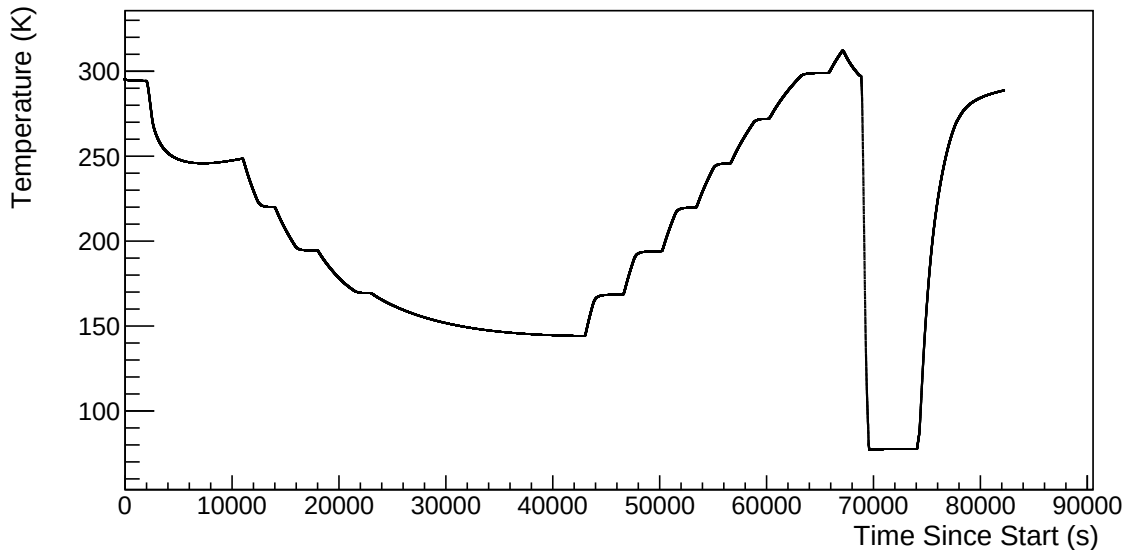

Figure 4.13: Temperature profile of the reference sensor during a calibration run illustrating the temperatures reached. The warm-up period followed the cool-down period automatically during the overnight run. The final part is when the copper vessel was submerged in liquid nitrogen.

P_0 (K)	P_1 (K/Ω)	P_2 (K/Ω^2)	P_3 (K/Ω^3)
210.059	2.49251	1.31873×10^{-3}	-5.32424×10^{-6}

Table 4.4: Average thermometer calibration coefficients for all thermometers given by Equation 4.4. Non-linear contributions from P_2 and P_3 are small.

only deliver 25 V, thereby limiting the linearity of the proportional term. An example of a run is shown in Figure 4.13.

Data was taken for six warm-up and six cool-down points and both data sets were fitted with a polynomial of the form

$$T = P_0 + P_1(R - 75\Omega) + P_2(R - 75\Omega)^2 + P_3(R - 75\Omega)^3, \tag{4.4}$$

for each case, where R is the resistance in Ω and T is the temperature in K. Equation 4.4 with the average set of parameters given in Table 4.4 contains negligible contributions from the quadratic and cubic terms, hence the function is effectively a straight line. To remove a correlation between the intercept and the gradient, the abscissa of the points is often shifted by the mean value. In this case due to the nominal operation being 60 Ω then this choice was biased to the lower end of abscissa values. The average from the fitted functions during the warm up and cool-down modes was taken, subject to the standard

quality assurance process, which requires the difference between the two curves to be small. This is discussed in more detail in the next section.

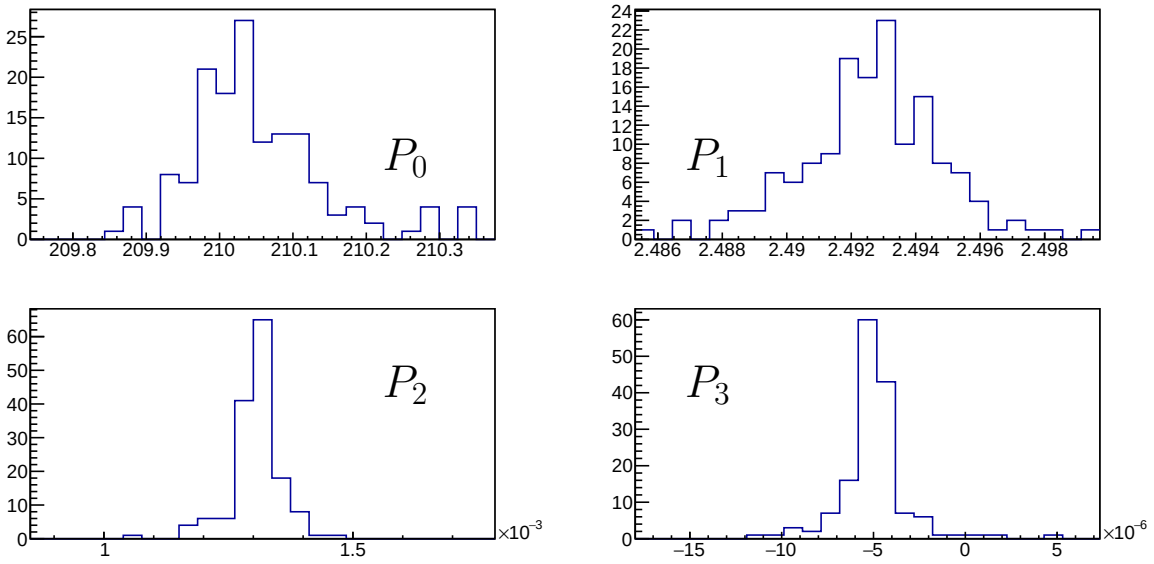

Figure 4.14: Values of the parameters for the calibration, based on those in Equation 4.4. These are distributed around a mean value which has been taken as the average set of parameters.

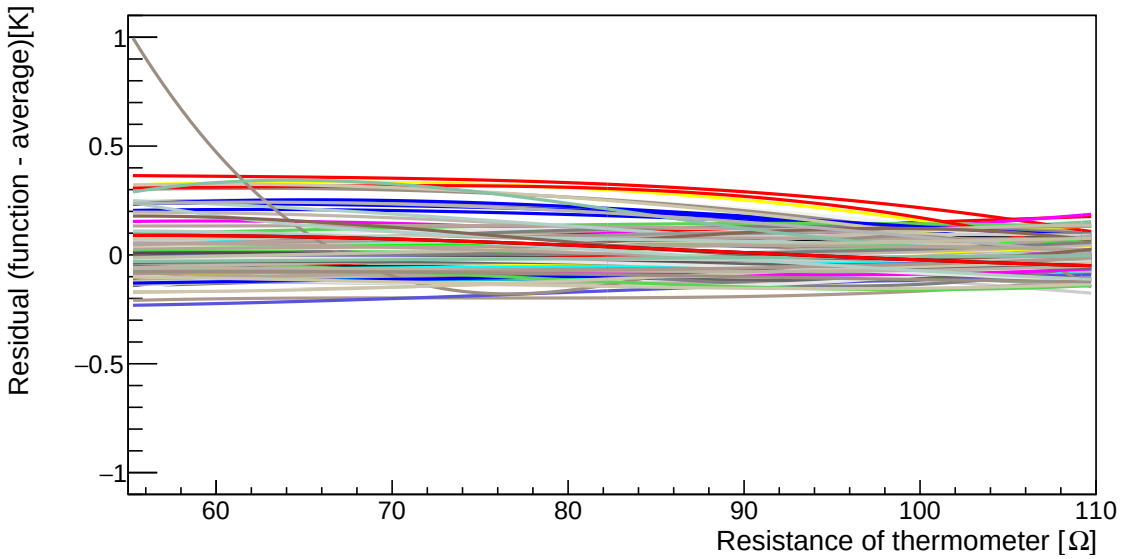

Figure 4.15: Calibration curves plotted for all thermometers when compared with the average curve. Thermometers meet requirements for accuracy, assuming the average is the best calibration, when any biases for given thermometers will cancel out. The sensor with the largest departure at lower temperatures does not show significant departures in comparison to other nearby sensors during operation.

The thermometers all have similar trends, validating that cross run systematics are

small. This can be seen in two ways: firstly, Figure 4.14 shows coefficients from all the calibration runs and secondly, Figure 4.15 shows a way of comparing the different calibrations to the average. In the former, the value of p_1 is ~ 0.4 which is comparable to the nominal constant of 0.385Ω/K, and the values are distributed in one peak. This also informs that the variation is small between PT100s in the batch. In the latter, each calibration curve is compared to the average calibration. This provides some indication as to the accuracy, given that all thermometers agree to within 1 K during the normal operating window with any set of parameters used.

4.4.2 Quality Assurance

Once the thermometers had been through the calibration procedure, a validation sequence was required. A series of tests that thermometers must pass to be used in LZ are detailed below.

Firstly, there must be no permanent hysteresis effects. This would highlight that the sensor had some sort of issue and may have broken. A failure rate of 3 in 130 thermometers was observed during calibration. To identify these faults the difference in the resistance from the reference one at room temperature at the beginning and end of the run was checked. If this differed by more than 0.07 Ω then the thermometer was rejected. This would be symptomatic of a change of about 0.2 K

Data was taken in both cool-down and warm-up configurations to confirm there are no hysteresis effects from either process. Calibration curves were fitted separately for each of the configurations and then results combined. In doing so, function values across the range were compared. Typically, differences of less than 0.05 K across the operating range, at any given resistance value, were observed. This showed that the hysteresis effects are small and also gives some indication that the uncertainties following cooling and warming are similar.

Estimating the uncertainties in the method was extended by repeating measurements with a set of four thermometers. Typically, observed differences in the calibration curves of

less than 0.01 K validate that the procedure is good. The conditions within the copper box are tightly controlled using the PID controller. Hence, dominant contributions between calibration runs are likely to come from external factors such as the level of the liquid nitrogen and the cable length. This test validates that these effects are sufficiently small and so can be neglected.

A further understanding of the thermometer accuracy is obtained by examining the differences in the measured resistance between sensors whilst running. Figure 4.16 shows the difference in ohms between each sensor and the reference sensor on channel 1 (red). The timing of each set of points is based on the data run in Figure 4.13. Only the regions where the temperature is stable are shown and for the basis of the 6 calibration points. Between these regions of stability, the values fluctuate significantly and so have been omitted for clarity. This is symptomatic of the sensors not being in thermal equilibrium. During these stable periods, the temperature is still changing, but at a continually differing rate such that the second derivative of temperature with respect to time is not zero. This suggested that equilibrium was being maintained throughout as the constant difference between the two sensors is maintained. The first set of points is taken at room temperature data when the PID system is not switched on. The greater range, which is on the order of $\sim$ 0.05 Ω or $\sim$ 0.13 K, between the first and subsequent set of points, suggests a small systematic offset during operation. This is conservative, as the comparison is not done at the same temperature, which can be seen in Figure 4.13. This is possibly due to differences in coupling with the plate with the PT100. Any systematic effects are likely smaller than this, hence if comparing the LZ thermometers (channels 2-11) any variation here suggests the 1 K accuracy requirement is met.

The thermometers were also examined at liquid nitrogen temperature. From the typical behaviour of PT100s, approximate linearity of response should hold to this point. If the calibration curves, when extrapolated, can reproduce the temperature of the nitrogen this acts as a test for the calibration. Specifically, it shows that the curve is valid beyond the bottom data point at $\sim$ 160 K and can provide some check that thermometers meet

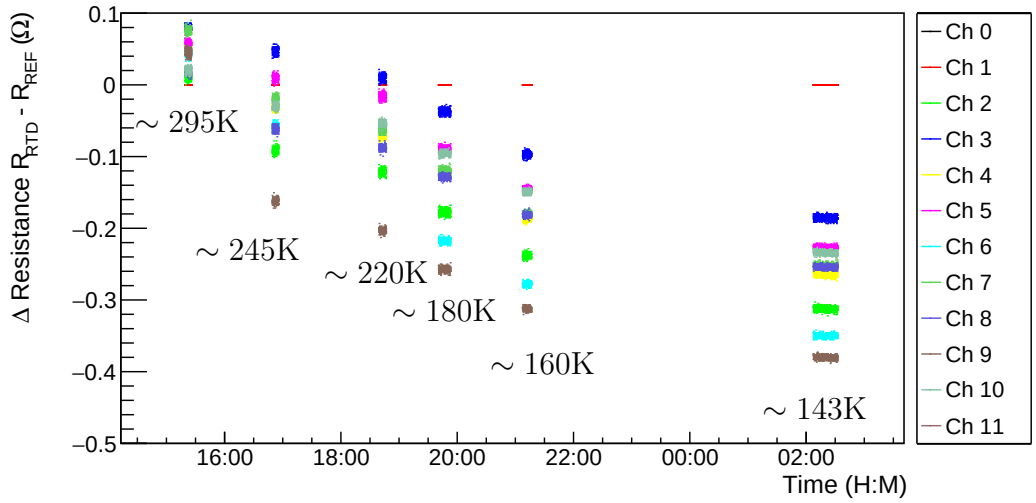

Figure 4.16: Difference in resistance for each thermometer to Lake Shore reference sensor (Channel 1) during the cool-down run on the 22nd January 2019 (see Figure 4.13). Data is shown for the points where the temperature is changing very slowly. Points where the system is not in equilibrium have been omitted for clarity, as these fluctuate significantly.

accuracy requirements. Figure 4.17 shows, this is still a good approximation, with all but 1 thermometers being within 1.5 K of the boiling point. The boiling point was calculated using the saturation curve for nitrogen based on atmospheric pressure data for each day [132][133].

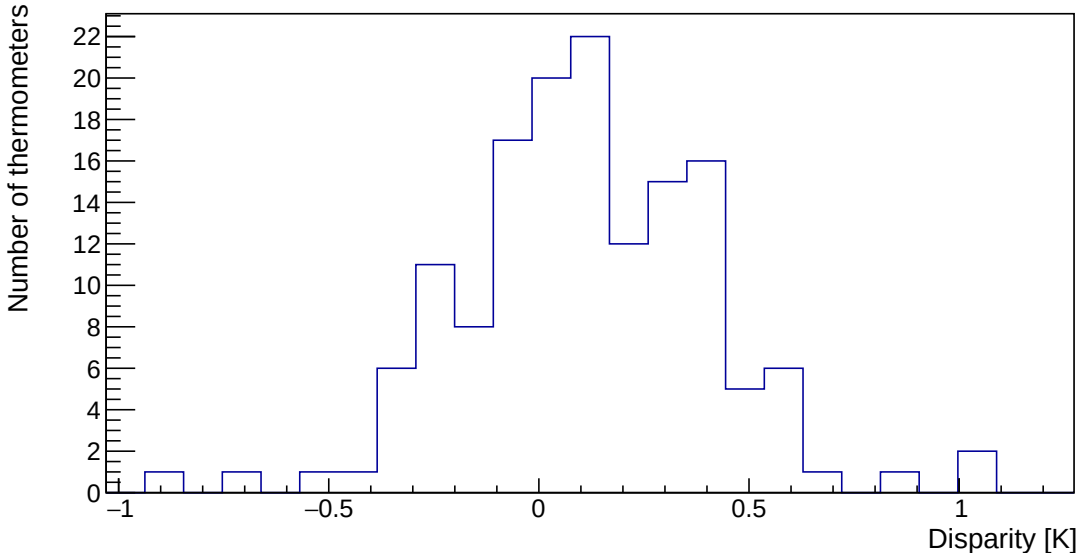

Figure 4.17: Difference for all calibrated thermometers between the value measured with the extrapolated calibration curve and the boiling point of nitrogen corrected for pressure variation.

The sensors were integrated into the detector systems. During cool-down (Discussed in the next Chapter) none of the sensors stopped working. Thus, illustrating the robustness of such a system.

4.5 Conclusion

The thermometry for the LUX-ZEPLIN experiment is necessary, based on the need to control the temperature within the cryogenic vessel, and to maintain a set of conditions that are suitable to collect data. The desire to do this drives the main requirements of such a sensor, and it is shown that the system of temperature sensors, the front-end read-out electronics in combination with the calibrations presented in this section, meets the requirements on accuracy and precision. The calibration procedure for the thermometers shows sub-Kelvin accuracy can be achieved, even after any systematics are estimated. The electronics are also shown to be sub-dominant in achieving the accuracy and precision requirements of the sensors themselves. In combination this evidence supports the use of these thermometers during detector operation.

Chapter 5

Understanding origins of position dependence in the ionisation signal

This Chapter discusses elements of the commissioning of the LUX-ZEPLIN detector. Specifically a focus is placed on the ability to understand origins of position dependence on the S2 signal response. This position response would make a search for Dark Matter impossible, if it were not to be addressed. A poorer energy resolution and nuclear-electronic recoil discrimination would be the main reasons for this. The position dependence on the signal is related to, amongst other things, the way the TPC cools down, the detector levelling, light collection efficiency effects and the xenon purity. A consideration of TPC sensors and the S2 response is used to examine these effects. Techniques are then developed to enable this position dependence to be calibrated out for the whole detector; an essential process for extracting a Dark Matter signal.

5.1 Motivation

5.1.1 Energy Resolution and background discrimination

It is vital to understand the origin of S1 and S2 position dependence in order to ensure optimimal energy resolution and the nuclear-electronic recoil discrimination.

The ability to reconstruct the energy of an event is an important consideration in two phase noble gas detectors. The reconstructed energy aids in the understanding of mono-energetic sources, calibrations of the detector, and is required to develop an understanding the xenon response [90]. As introduced in Section 3.1.1, the reconstructed energy is based on the magnitude of the position corrected S1 and S2 signals ($S1_c$ and $S2_c$) and is given by

$$E_R = \frac{W}{\mathcal{L}} \left(\frac{S1_c}{g_1} + \frac{S2_c}{g_2} \right), \tag{5.1}$$

where g_1 and g_2 are detector specific quantities related to the efficiency of detecting light for an S1 and S2 signal respectively. The factor $\mathcal{L}$ is the Lindhard factor and is dependent on the interaction type. Crucially, this relationship requires the position dependence to be removed from the signals to ensure that the energy can be reconstructed properly. By leaving in this dependence a worse energy resolution would result. Figure 5.1 shows how the position dependence can affect the signals and therefore broadening the distribution of a hypothetical 50 keV mono-energetic electronic recoil source in both S1 and S2. This is a typical recoil energy LZ has sensitivity to. The energy resolution is (by Equation 5.1) related to this broadening and hence illustrates the need to perform the correction. In the corrected case the remaining variance is largely due to the recombination fluctuations present in the xenon physics. Some signal searches within the detector make use of the recoil energy in statistical analyses. This would not be a reliable parameter to use should position dependence in the signals remain.

Additionally, the ability to separate electronic from nuclear recoils is also of primary importance. The majority of the background events within the detector are electronic recoils [112], whereas a WIMP would coherently scatter from a nucleus. Hence, remov-

ing electronic recoil events with 99.5% [61] efficiency significantly improves sensitivity to WIMP-nucleus interactions. From the xenon micro-physics the partitioning into the S1 and S2 signals is different. The S2 signal is disproportionately smaller for a given S1 due to the xenon microphysics. Therefore, the relative size of the S1 and S2 signals can be used to discriminate these two interactions. Should the S1 or S2 signals be position dependent then the nuclear and electronic recoil bands would have greater variance and hence the leakage would increase. As many effects result in the S2 signal being smaller, this is particularly problematic as it makes electronic recoil events look more like nuclear recoils.

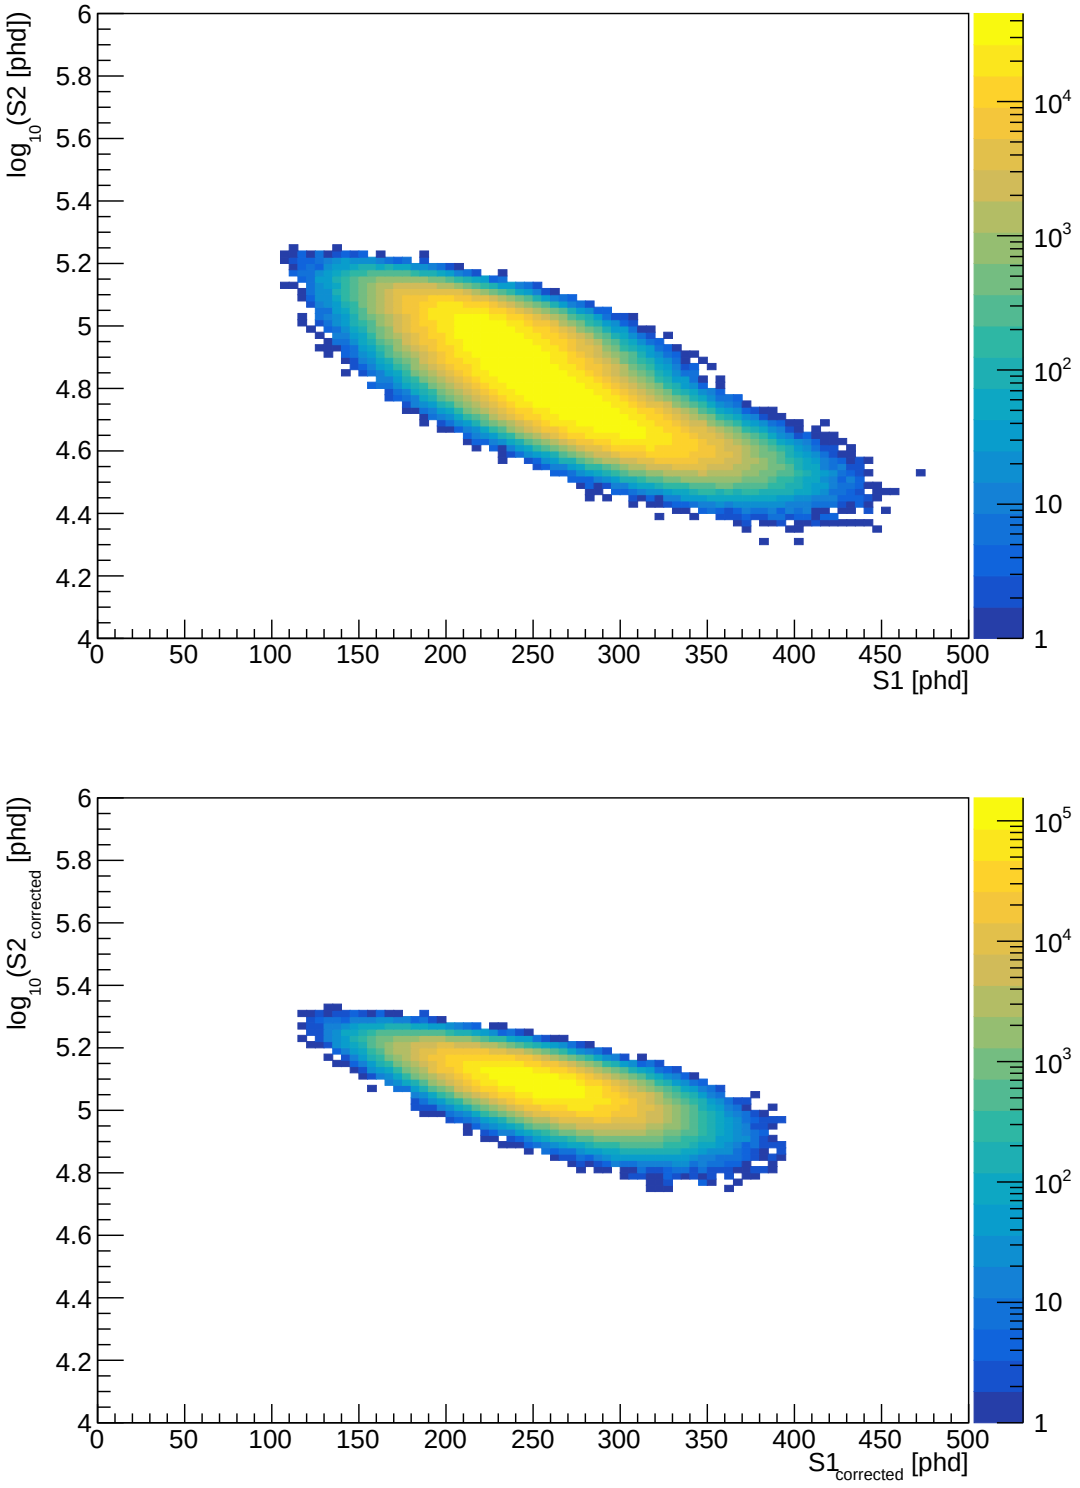

Figure 5.1: Expected signal sizes for 10^7 50 keV electronic recoils, for the projected LZ configuration. Recoils of this energy are typical. The top plot shows the raw number of photons detected for each signal, and the bottom shows the same events when the position dependence has been removed from both the S1 and S2 signal. Plot made using NEST simulations [99].

5.1.2 Origins of S2 Position dependence

There are many effects which potentially lead to the variation of the S2 signal within LZ. These effects can be divided into three main categories. Firstly the geometry of the extraction region and how that governs the motion of the electrons during extraction and electroluminescence. Secondly, the variation of light detection efficiency of photons produced within the extraction region (a function of the geometry and PMT response). Finally, the effects on the drifting electrons from xenon purity and electric field. The principles here are not unique to LZ and hence could more generally be applied to any two phase noble-gas experiment performing a similar search.

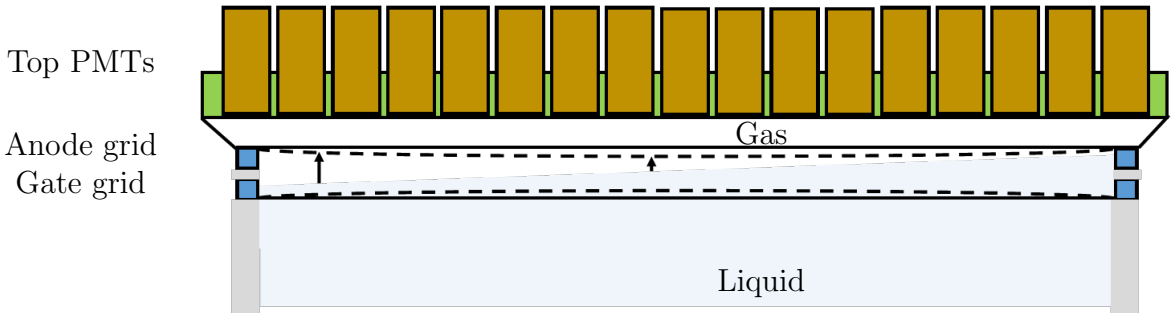

Figure 5.2: Diagram to illustrate the cross-sectional anode and gate bending due to electrostatic attraction (dashed line) in comparison to the zero field case (solid lines). The liquid level has been exaggerated here to not be parallel with respect to the zero-field grids. The black arrows illustrate the different electroluminescence paths generating different S2 signals which is then a function of the position. Diagram is not to scale.

The extraction region geometry determines the electroluminescence behaviour and is in general position dependent. This means slightly different electric fields and also a differing distance between the liquid surface and the anode grid wires can be present. The S2 response is parametrised by Equation 4.1 [30] and hence illustrates how these effects can change the signal yields and timing distributions for an S2 signal. These situations can arise in several ways. For example, due to electrostatic attraction the anode and gate are closer (by about 1 mm) in the centre than at the edge [103]. Hence, there is a spatially dependent path difference and electric field which leads to a position dependent response. Additionally, path differences can occur due to the grid wire alignment between the gate and anode. There are places where the anode and gate grid wires align, but also regions

where they do not [103]. Due to kinematics, the electrons sometimes do not have enough energy to be extracted into the gas region to generate the S2 signal. Therefore, part of the signal can be lost in a spatially dependent way. Furthermore, in the unlikely event that the detector moved and the liquid level became not be parallel to the cathode then an observable position dependence could be observed in the S2 signal. This would also be the case if the contraction of the TPC during cooldown was not completely uniform or a small slant remained after liquid levelling. From NEST simulations, the mean S2 signal size as a function of the gap between the liquid surface and the anode varies linearly by about 10% for a 1 mm difference in liquid level [99]. Hence, even a very small departure could lead to a notable difference. Figure 5.2 illustrates some of the mentioned effects.

Beyond the electroluminescence behaviour, the light detection efficiency from the produced light can vary due to the spatial placement of the walls and PMTs. This is linked to the light propagation of photons from the S2 generation. Photon paths can differ as a function of the electron extraction position, introducing a position dependence. Additionally, different PMTs can have different detection efficiencies, and there are some PMTs that are excluded from the analysis that will affect the light collection efficiency.

In terms of the xenon purity, electronegative impurities in the xenon, such as oxygen ions, can capture electrons during the drifting process. These impurities are in the gas when obtained from the supplier, and from detector component out-gassing and are largely unavoidable [30]. When circulating xenon, the getter (Section 3.2.5) can remove many of these ions. This means this effect is generally time dependent. The rate capture is proportional to the original number of electrons produced [134]. Mathematically this is written as

$$\frac{\mathrm{d}\mathrm{N}_e}{\mathrm{d}t} = -\,\alpha \mathrm{N}_e\,, \tag{5.2}$$

where α is a constant, and N_e is the number of electrons. This behaviour means that the number of electrons reaching the surface falls off exponentially with drift time, t.

All of these effects lead to a full three-dimensional position dependence in the S2 signal. Section 5.4 discusses how these effects can be calibrated out and hence largely

removed from the response.

5.2 Shrinkage of the TPC during cooldown

The various detector materials contracted when cooling from room temperature to about 185 K during detector commissioning in xenon gas. This fall in temperature of about 110 K was closely monitored using the LUX-ZEPLIN thermometry (considered in Chapter 4). It was possible that the detector could have contracted non-uniformly, whereby different regions could have cooled faster than others leading to thermal stresses. It was also possible that, despite extensive planning, a small permanent slant could have been introduced in the TPC stack. From Section 5.1.2 this deviation can be very small and still have an impact on the S2 signal. Knowledge of the TPC movement during this operation is important as it is for subsequent times to examine further movement.

Many of the effects leading to position variation in S2 response are degenerate and hence separating them is difficult. Position sensors, installed within the ICV lid can help resolve degeneracies that lead to variation in S2 response.

5.2.1 Position Sensors

Six position sensors are installed in the xenon gas space above the top array. These sensors measure the distance between two points connected by copper rods (Figure 5.3). One point is fixed to the ICV lid and the other is forced onto the top PMT array with springs and is kept in place with a guidance pin. Originally, this was designed to form a device known as a Hexapod (Figures 5.5 and 5.4). Such a configuration enables the 6 degrees of freedom to be resolved with the orientation of two hexagons. One such hexagon is fixed to the lid of the ICV and the other is formed with the guidance pins on the top array. As Figure 5.6 shows, the sensors are arranged into three triangular configurations. By arbitrary naming convention, Position Sensors 1 and 6, 2 and 3, and 4 and 5 form these triangles as illustrated in Figure 5.4. During installation the mounting point associated with sensors 2 and 3 did not engage properly, and as the later analysis

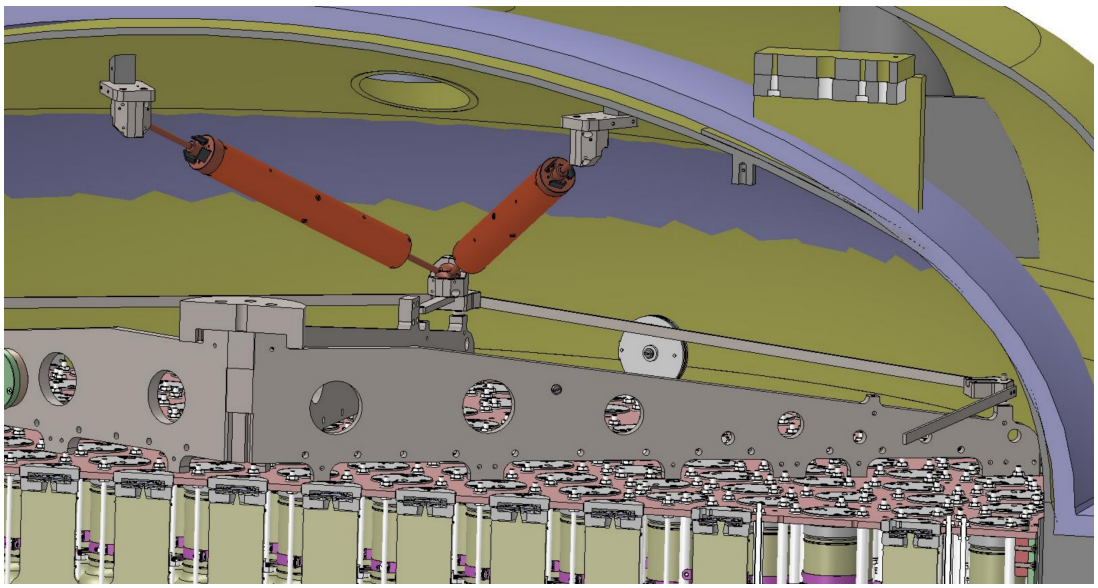

Figure 5.3: Computer aided design model of two position sensors within the lid of the ICV. One of the three corners is shown with the top ball joints mounted to the ICV lid and the bottom sitting on the top array truss. These sensors measure the distance between the lid mounting points and the top array truss.

will show, this means all sensors have to be considered individually to confirm if the movement is consistent with solely vertical motion.

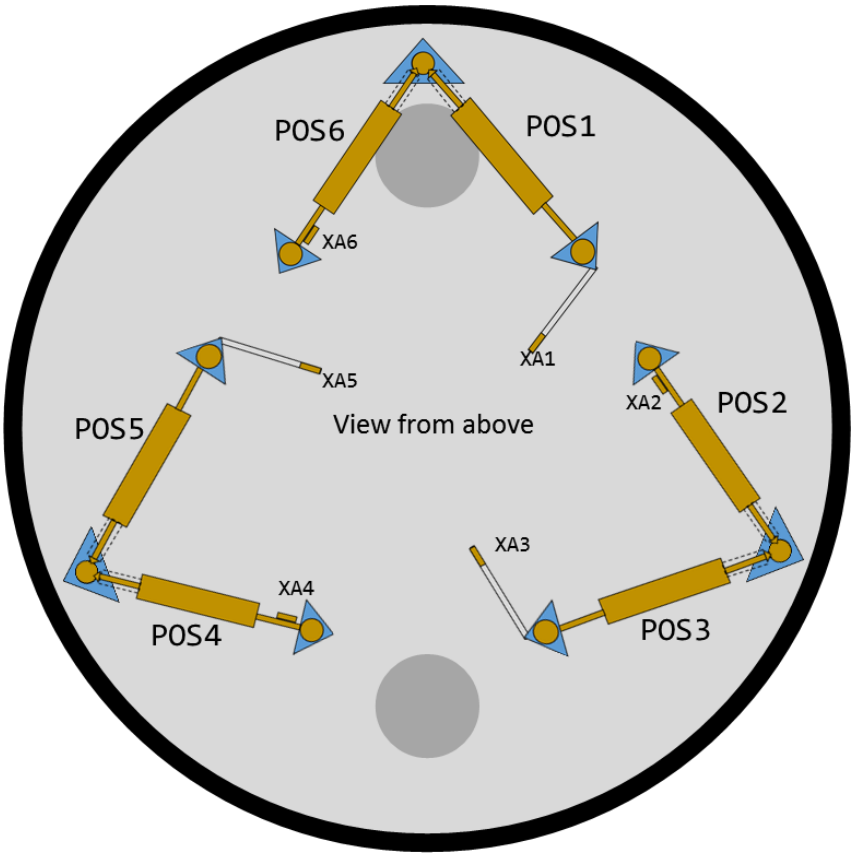

Figure 5.4: The orientation of the six position sensors (POS) within the lid of the ICV. Thermometers have also been included on the diagram and are labelled as XAx. This shows the arrangement within three triangular configurations. The two ports on the lid of the vessel are included to clarify the geometrical arrangement.

Figure 5.5: The Inner Cryostat Vessel lid (ICV) on its side prior to installation. The main sealing surface is shown here protected with polycarbonate pieces to prevent scratching. The position sensors can be seen within the lid, forming three triangular arrangements. The top array of the TPC (covered in cables) can be seen in the bottom left of this photograph.

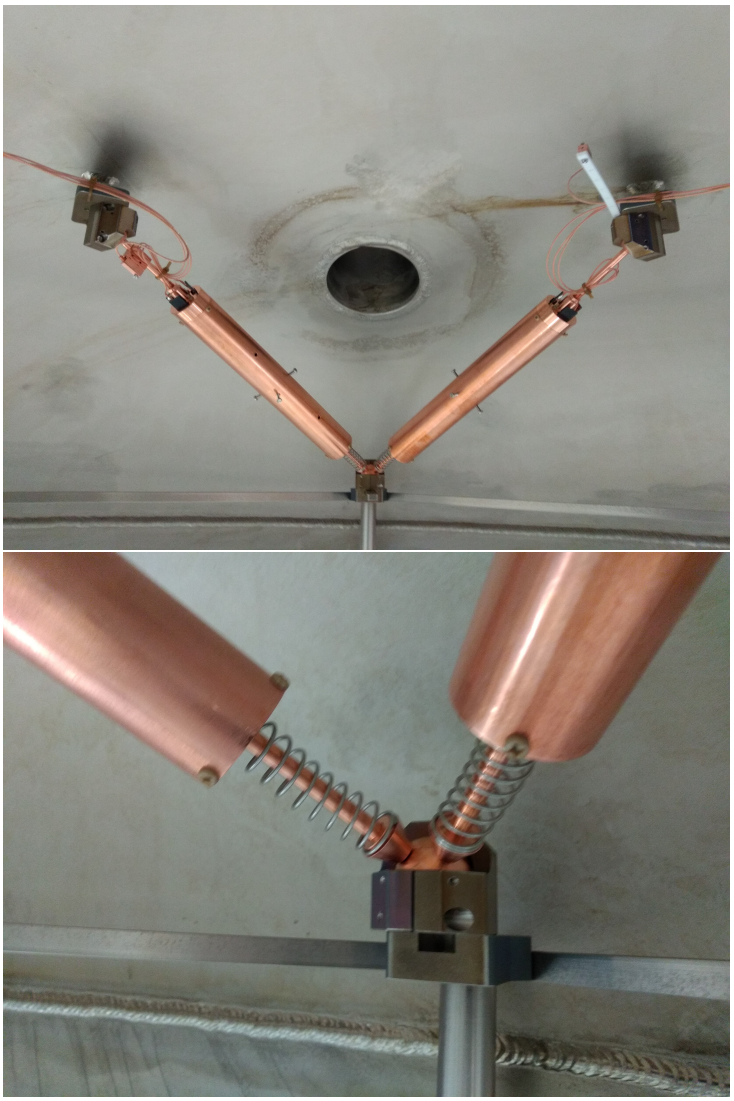

Figure 5.6: (Top) An example of a corner triangular arrangement. This image is taken with the vertical calibration rods installed (very bottom of the image) which mimic where the top array would sit during initial installation. (Bottom) Springs seen at the ball joint were added to ensure that the bottom hexagon maintains a good attachment to the top array.

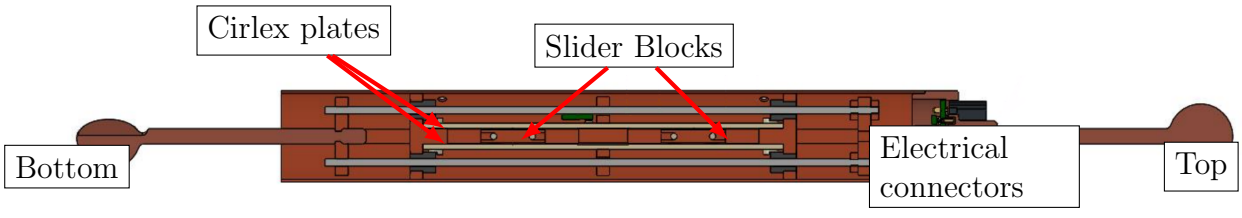

Figure 5.7: Cross sectional view of a position sensor connected to copper rods which fit into ball joints. The outer copper tube contains a dynamic arrangement of capacitors. The cirlex plates contain the electrodes which form these capacitors. Copper slider blocks move when the sensor is extended or compressed and form a different arrangement of capacitors. The spring attachment (not shown) fits onto the bottom half ball joint support.

The sensor itself is constructed of two capacitors inside a copper shell (Figure 5.7). The copper shell attempts to provide electrical shielding against external interference. The two capacitors each have a slider which sits between the plates. As the sensor extends and compresses, the slider moves and changes the capacitance of each capacitor. This forms a continuum of values when the block is moved spatially between the plates. Using a calibration (as discussed in Section 5.2.3) the extension of the sensors can be measured as a function of the capacitance ratio of the two capacitors.

5.2.2 Capacitance sensor electronics

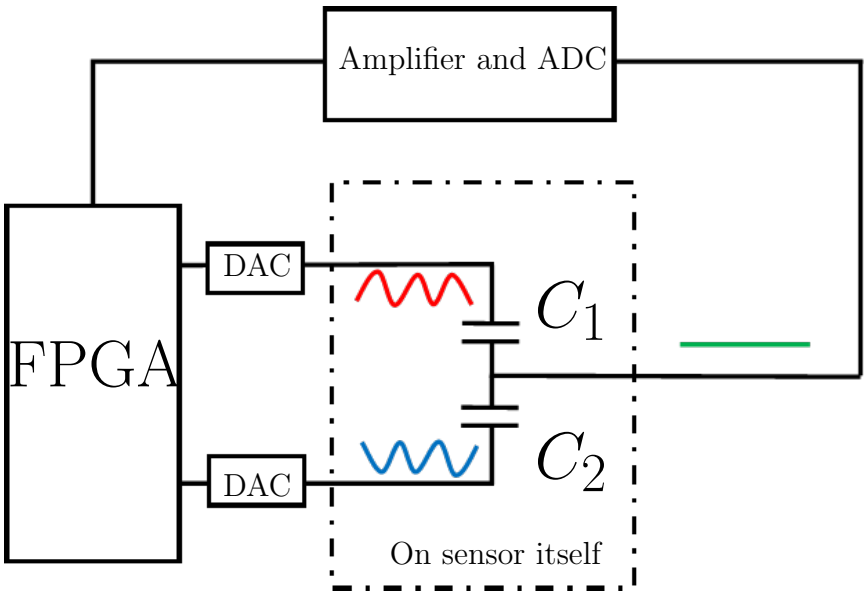

Figure 5.8: Schematic diagram illustrating the electronics configuration. The positive and negative excitations are fed in from the two Digital - Analouge converters (DACs) and then into the two capacitors (C_1 and C_2) in anti-phase (blue and red). The amplitudes are tuned so that their sum is zero (green). The output is fed into the FPGA to complete the feedback circuit.

The electronics used for the capacitance measurement are custom-built to perform this task. Developments of such electronics are discussed in detail within Reference [135], however to understand the arguments made in this Chapter the basic principles are outlined. The final design is a digital circuit based on a feedback loop. Two excitation signals are generated and are each passed to a separate capacitor within the sensor. The two voltage excitations are 180 degrees out of phase with each other as illustrated in Figure

5.8. The amplitudes of the excitations are tuned within the feedback loop so that the outgoing voltage from the parallel arrangement destructively cancels.

By considering the circuit characteristics, the output of the electronics is given by

$$f = 4096 \times \frac{C_2 - C_1}{C_2 + C_1}, \tag{5.3}$$

where f is an integer value between -4096 and 4095, and C_1 and C_2 are the two capacitors considered. The important feature of the method is that many of the external factors contributing to the measurement cancel out by considering a ratio of capacitances. Additionally, the electronics operate at about 195 kHz, which ensures that the capacitor impedance is low, to reduce the impact of interference.

5.2.3 Calibration of position sensors

In order to examine the extension of the sensor, a calibration against the ratio of capacitance (Equation 5.3) was performed. This maps the output of the electronics to the extension of each sensor. Therefore, during operation movement of each sensor can be tracked. In order to do the calibration each sensor was clamped into a vice to prevent movement. A bolt was attached to the moving plate and used to compress the sensor throughout the physical possible motion of around 25 mm. A commercially available low-power laser then measured the distance to the end of the moving section. Readings were taken several times, but the electronics varied by no more than a few integer units and so this variance can be ignored. At the centre of the operating region this corresponds to less than one-tenth of a percent variation in the ratio of capacitances. The resulting calibration curve follows a one-to-one mapping over the fitted range. The relationship is not linear as a result of edge effects and stray capacitance within the sensor. In all cases the data was fitted to the function,

$$f(x) = a + \frac{b + c(x+d)}{e + f(x+d) + g(x+d)^2}, \tag{5.4}$$

where x is the extension of the sensor and all other terms are experimentally fit constants. The format of this fit function is similar to previous studies with prototypes [135], but has been generalised to allow for calibration offsets. This offset is selected based on a point where the total sensor length was known following the installation in the lid. Hence, the length of the sensor is always a fixed length plus the value obtained from the calibration curve. Figure 5.9 shows an example calibration curve for one of the Position Sensors and the residuals to the fit are shown in Figure 5.10. This calibration illustrates the importance of being within the middle of the operating range of the sensor. Towards the centre of the range the sensitivity is of the order 2.4 µm per integer unit, whereas nearer the extremes of the extension and compression this is closer to 10 µm per integer unit. On the calibration curve at the extremes of the range the calibration becomes non-monotonic. This is believed to be due to a design error, when the copper blocks within the

position sensors move close to the end of the tube. At the extremes of compression and extension, where the data values ultimately flip on the most extreme points, the sensor reaches the end of the operating range. In order to place the sensors within the middle of the operating range, assumptions were made using engineering drawings which were then tuned after installation. The residual distribution shows the fit is not perfect. However, the errors lead to less than 25 µm within the fit, an acceptable value considering this is 1/1000th of the range of the sensor. There appears to be some correlation in the residuals to the fit. Although this is not ideal, the scale of these effects is sufficiently small that they will not impact the final result.

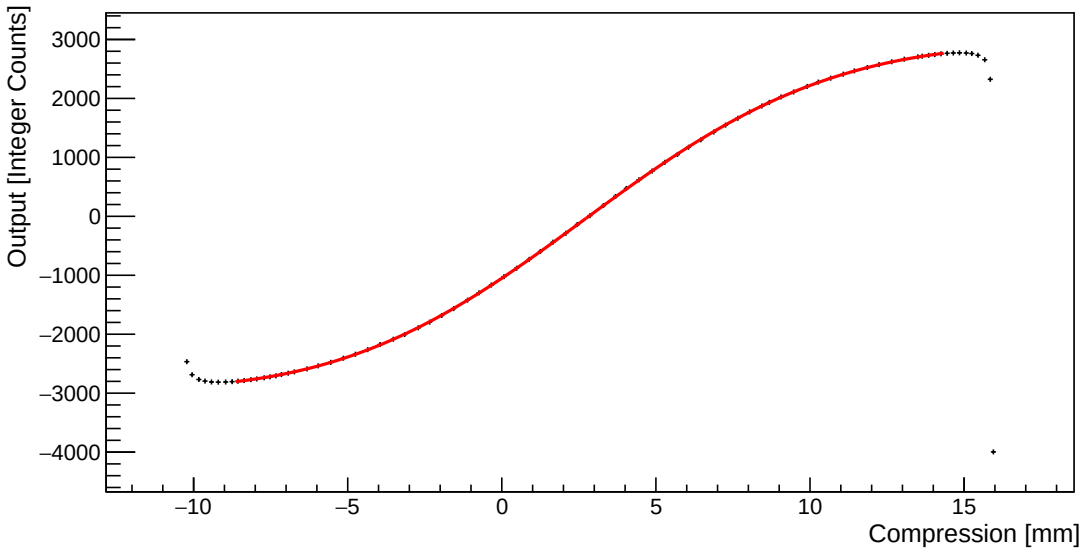

Figure 5.9: Calibration curve for Position Sensor 1 (POS1) which maps the sensor readout from the electronics to the compression. The red line is fitted to the data with Equation 5.4 for the range shown. The zero of abscissa is effectively arbitrary and is used in combination with installed measurements to calculate the length (see main text).

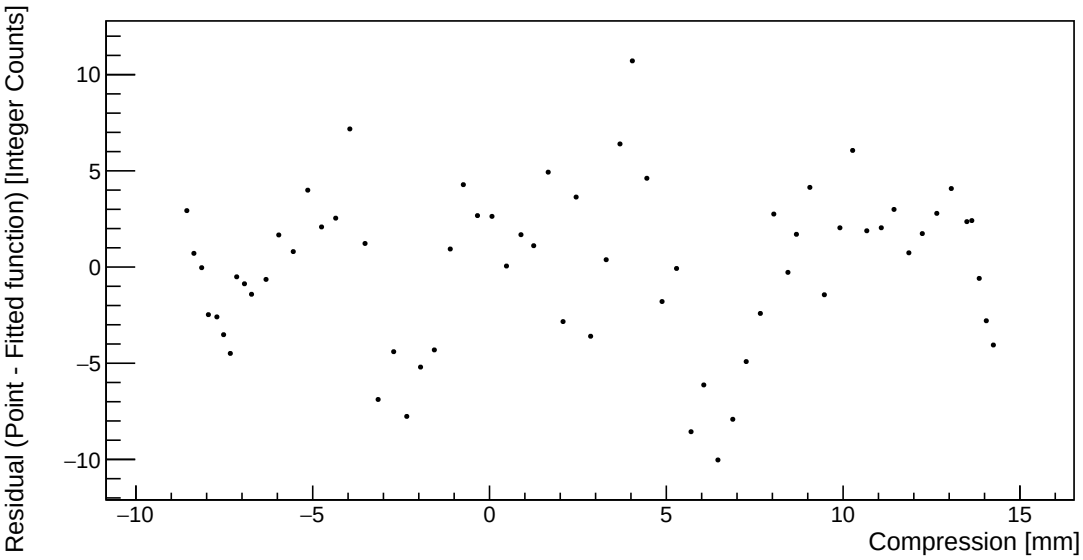

Figure 5.10: Residuals to fitting of the calibration curve. The maximum error is on the order of 10 counts which is of the order 25 µm in the centre of the operating range.

5.2.4 Cool-down of LZ in Xe Gas

Expected Shrinkage

The shrinkage of the detector can be estimated, based on information about the constituent components as shown in Figure 5.11. A simplified model is used for this purpose where these components are treated separately. From the top downwards, there is the gas region, with the position sensors (here it is assumed sensors are solely made from copper). The second region is the PEEK and titanium support structure on the top array. Thirdly, there is the main forward and reverse field region stack which is approximately 20% PTFE (Teflon), and the remainder titanium. The final region is the support structures at the base of the detector.

Considering literature for the shrinkage of Teflon [136], titanium [137], PEEK [138], and copper [139] an estimate formed in Table 5.1. This leads to a total change in height of the top array of 3.9 mm. This value is only approximate and based on the approximations made, an uncertainty of about 15% is appropriate on this value. This is due to the use of empirical data for the expected shrinkage, but also does not account for some potential intricacies of part dimensions. Changing the value of the expected shrinkage of the PTFE to 1.4% (not an unreasonable number) would increase the shrinkage to 4.5mm overall [135]. In calculating the contributions to the change in length of the position sensors, the shrinkage of parts made from titanium in the TPC stack cancel out with an equivalent shrinkage in the ICV wall. Hence, the dimensions of these parts do not require high precision. The values have however been included to illustrate the magnitude of the shrinkage of each component together.

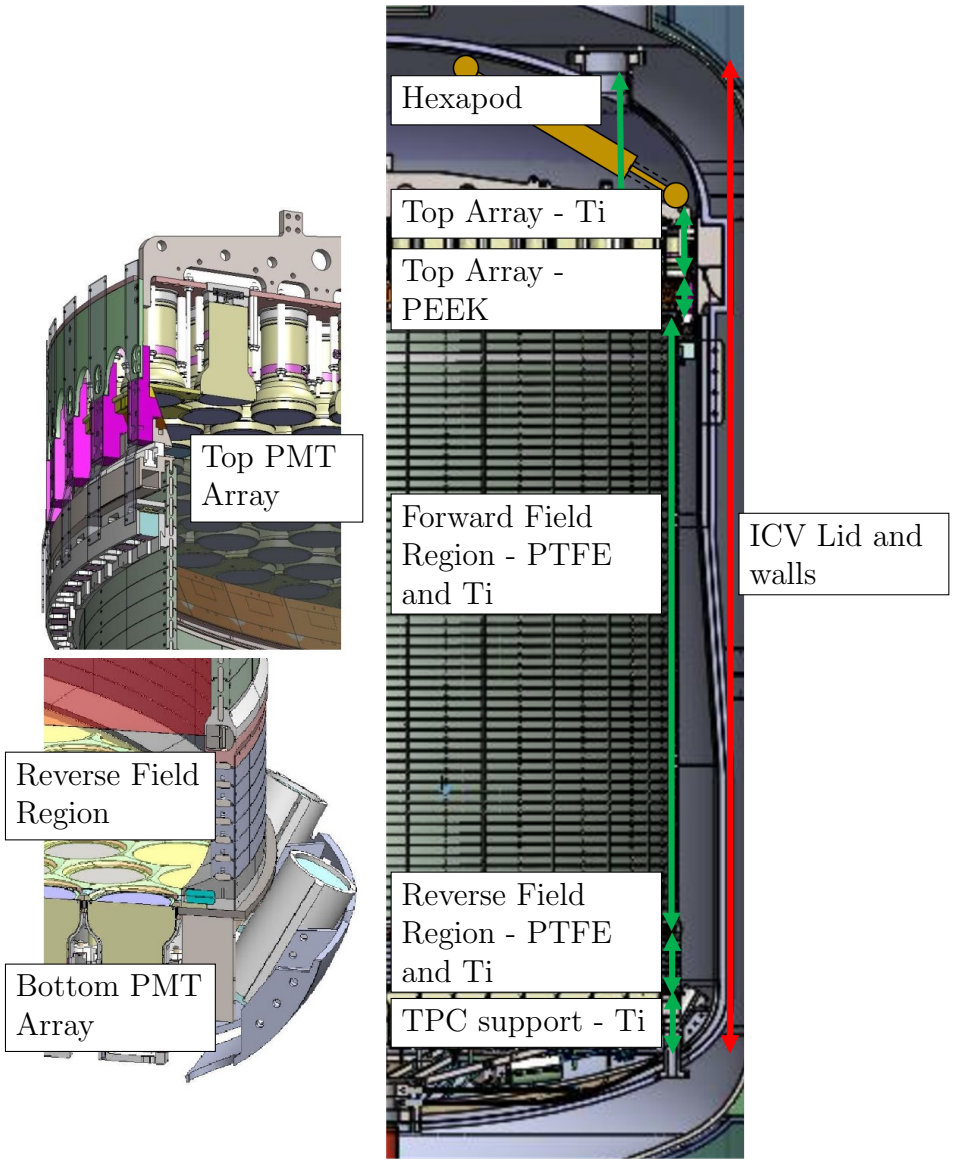

Figure 5.11: (Left Bottom) The bottom region of the TPC, and reverse field region (RFR) showing xenon skin PMT's, the bottom array and the cathode and bottom grids. In the RFR the titanium field rings make up less of the region than in the main TPC region. (Left Top) The top of the TPC and top PMT array. Here the (depicted as pink here) PEEK blocks can clearly be seen around the 'colosseum' structure on the top array. (Right) The full TPC and ICV structures as outlined in the calculations. Note the position sensor shows the approximate location of these sensors, but is not to scale.

Part of Detector	Material	Length (mm)	Expected Shrinkage (%)	Total Shrinkage (mm)
Hexapod Structures	Copper	166	0.17	0.28
Top array	PEEK	63	0.44	0.28
Top array	Titanium	200	0.08	0.16
Forward and Reverse field regions	PTFE	309	1.12	3.80
Forward and Reverse field regions	Titanium	1285	0.08	1.03
ICV Lid	Titanium	166	0.08	-0.13
ICV walls	Titanium	2156	0.08	-1.73
TPC support	Titanium	300	0.08	0.24
Sum	-	-	-	3.93

Table 5.1: Estimate of the shrinkage of the detector components that lead to the position sensor extension. Note here that the positive sign denotes that the component will extend the position sensors, and a negative sign will contract them.
‡ This is a projection into the z direction of the change in the length of the sensor.

Sensor Response

During the cool-down process the detector conditions were monitored with the LUX-ZEPLIN slow control system. The average temperature from the thermometry system in the vessel was monitored along with the output from the position sensors. Figure 5.12 shows a consideration of the cool-down of the detector throughout February 2021. The calibration curves have been used in performing the measurement of the length. Note that as with the calibration the absolute scale here is linked with the original setup and value of the nominal length.

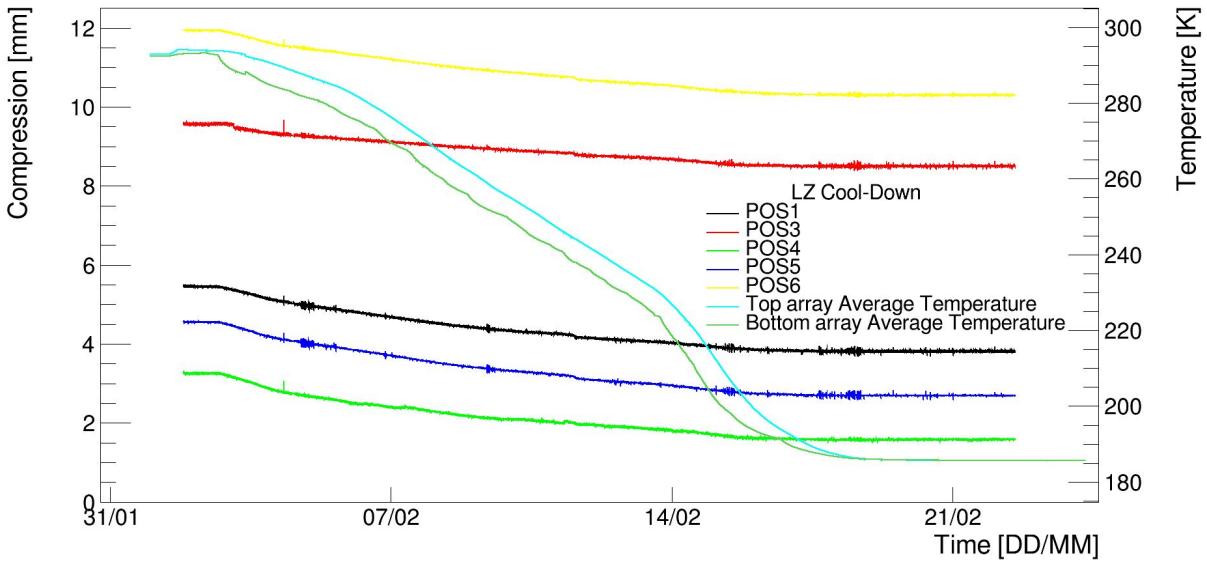

Figure 5.12: The movement of the 5 active position sensors during the cool-down in February 2021. The calibration curves have been applied to convert the response to a length. There is an obvious temperature gradient across the detector during the process between top and bottom PMT arrays and it is visible that these sensors move in correlation with the reduction in temperature.

Here, five of the six sensors can be observed in motion; correlated with the temperature falling. Four of the sensors, are within the middle quarter of the operating range, where the sensitivity is better, as expected due to pre-installation calculations. The remaining two sensors (2 and 3) are both on the same corner piece. During installation this failed to bind properly to the top array and so the corner piece was considerably higher and also has greater flexibility to move than was originally intended. A photograph of this pin is shown in Figure 5.13.

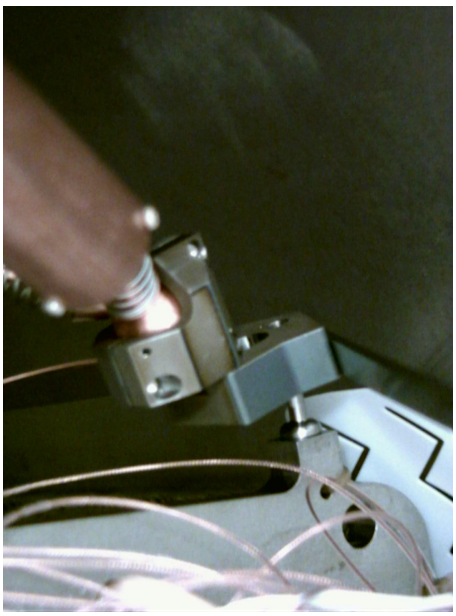

Figure 5.13: Figure showing the hexapod pin which failed to bind to the ball joint enclosure during installation. This ball joint includes both sensors 2 and 3. Image is taken using a camera inserted through one of the ICV ports of this pin which is effectively inaccessible.

The two sensors on the high pin moved differently than the rest. POS2 did not appear to move as expected and remains at the top of the range and so has been excluded from Figure 5.12. It is possible that this sensor is within the non-monotonic region within the calibration curve, or that some extra degree of freedom is allowed on the pin so as the detector cooled the motion was small. Regardless, the lack of knowledge about conditions on this pin means the system of sensors cannot easily be considered together. However, some information can be gained from the sensors when considered individually. If the assumption is made that the TPC contraction was all vertical and that no motion occurred in the radial plane, then the movement of each sensor can then be calculated and self-consistent checks made to ensure that this is a valid assumption. This is required, because the degrees of freedom need to be restricted to enable the problem to be resolvable.

Change in height

Based on the start and end point of the extensions the change in the sensor length was calculated. The different values for each of the sensors are shown in Table 5.2. These figures do not make any assumption about the contraction of the sensor components which

needs to be considered along with a conversion from the change of length of the sensor to the change in vertical height of the top array.

Position Sensor	Change in length (mm)
1	1.64
2	-
3	1.06
4	1.67
5	1.85
6	1.64

Table 5.2: Movement of the sensors inferred from Figure 5.12 which shows the change in extension of the position sensors in turn. Sensor 2 is excluded due to the lack of movement.

In order to convert the change in length of the sensor to a vertical change in height (Δz), trigonometry is used to give

$$\Delta z = \sqrt{P'^2 + r'^2} - \sqrt{P^2 + r^2}, \tag{5.5}$$

where P is the length of the position sensor and r is the projected length between the two ball joints, which is the theoretical length of the sensor if the two ball joints were at the same height. Primed coordinates denote the new lengths after cool-down. Within the original Hexapod algorithm some values were not required, and so measurements were not undertaken. Hence, uncertainties on the measured values are larger and lead to systematic variation in the measured height difference. Specifically, the measurement of the dimension, here labelled as r, has a high systematic uncertainty. The value itself has to be estimated from geometry, assuming regular hexagons and a common centre between them. This idea is illustrated in Figure 5.14, where the value of r can be extracted from knowing the side lengths of the hexagons. These were measured in the cleanroom during the installation. They were corrected according to the centre of the ball joints as illustrated on Figure 5.6. Hence, the dimension r is estimated by

$$r = \sqrt{\delta^2 + \frac{\zeta_{\text{inner}}^2}{2}}, \tag{5.6}$$

where δ is obtained by using that $\eta_{\text{inner}} + \delta = \zeta_{\text{outer}}$. To obtain r' a contraction factor is applied to r in accordance with values in Table 5.1.

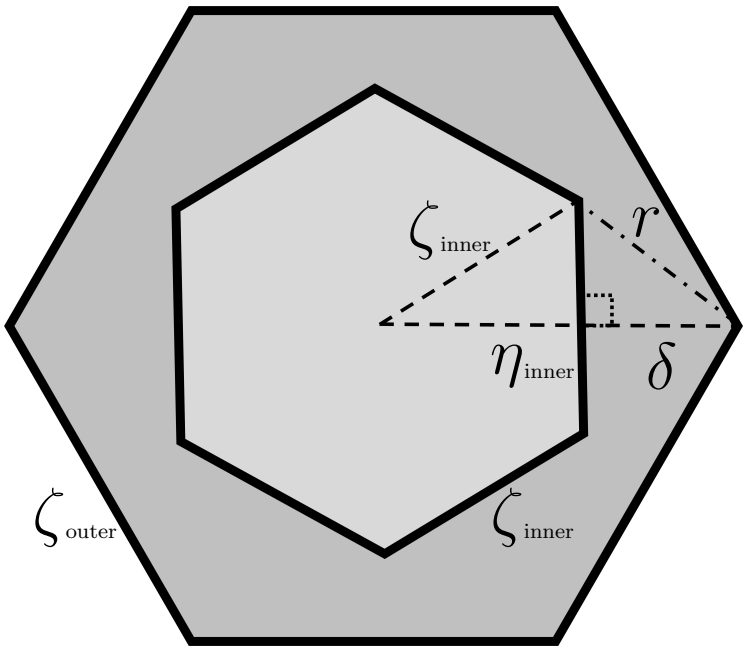

Figure 5.14: Hexagonal geometry of the outer and inner hexagons. The inner hexagon is mounted to the ICV lid and the outer hexagon is mounted to the top array. Diagram defines the variables for the regular hexagonal arrangement considered in the text.

This method leaves r with a sizeable systematic uncertainty of about 10 mm. In order to propagate this uncertainty the values of ζ_{inner} and η_{inner} are varied and propagated through Equation 5.5 to give the resulting bounds on the difference in height. Hence, the final movement in the vertical direction is given in Table 5.3 and summarised in Figure 5.15. The statistical uncertainty from the measurements is sub dominant to the systematic uncertainty quoted here, based on the variation in the electronics being only a few integer counts ($\sim$ 10 µm). In considering the sensors together the error is based on the average of the 4 sensors that are giving consistent results and their uncertainties. Errors here are not combined and just the maximum and minimum taken, due to correlations between the errors being unknown. These correlations are present as a large value of r is likely true on all sensors and affects all values consistently.

The contraction of sensor components leads to a small systematic on the value, as the capacitor plates within the sensor contract differently to the slider block. From previous studies [135] this is expected to be about a 0.6% variation on the output value f. For a typically larger value of 1000 integer counts this corresponds to only 6 counts variation. This is much less than the measurement considered here and hence is ignored.

Position Sensor	Inferred Vertical movement	Uncertainty - Lower [mm]	Uncertainty - Higher [mm]
1	4.14	0.41	0.66
2	-	-	-
3‡	3.13	0.37	0.64
4	4.53	0.52	0.90
5	4.90	0.54	0.93
6	4.32	0.47	0.78
Average‡	4.47	0.74	1.36

Table 5.3: Inferred movement of the sensors in the vertical direction. Average values have just taken the maximum and minimum values for the errors rather than propagating through to the average value.
‡ This sensor was on the higher pin and hence the results here should be taken with caution. Average does not include Sensor 3.

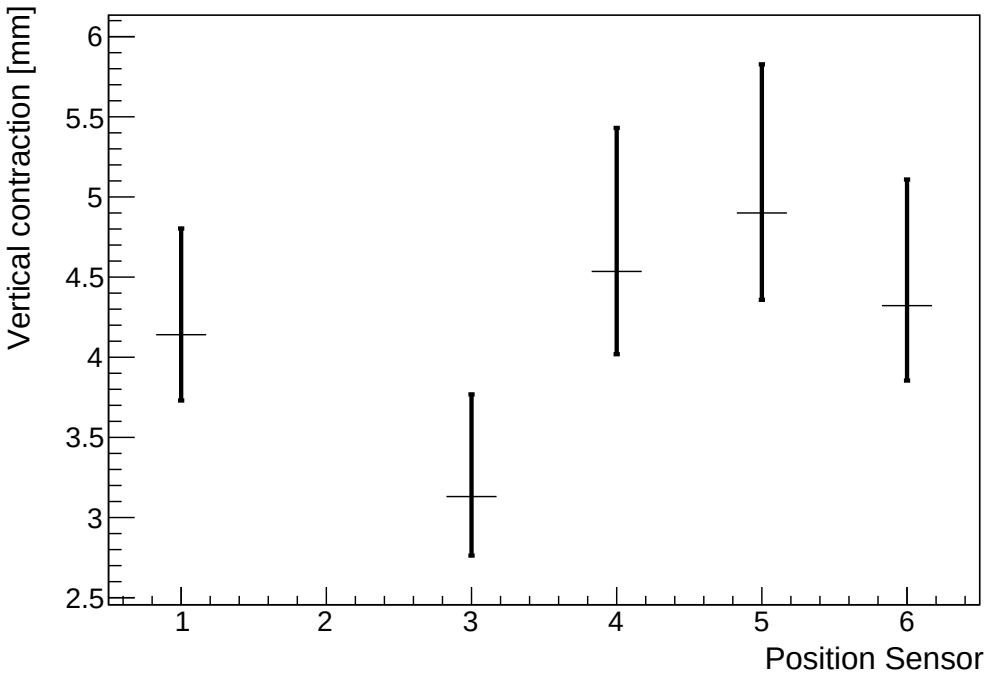

Figure 5.15: Translation of the motion into the vertical direction. This includes the movement of the top array relative to the ICV lid as well as the effective motion caused by the sensors contracting. Error bars are purely due to the systematic uncertainty in the distance r.

Based on the information presented, it is highly likely that the motion of the TPC was purely vertical. The relative difference between the two sides is likely < 1.5 mm which is small compared with the over 1.5 m TPC. Should the TPC have tilted during

the cooling process then each corner of the sensors would have moved differently. If the motion had been rotational then output from each sensor within a pair would have been greatly different. Comparing sensors 1 and 6 together as well as 4 and 5 does not show evidence of this motion. The vertical motion is also in fairly good agreement with the expectation from theory, despite the uncertainty.

5.3 Liquid Levelling

5.3.1 Level Sensors

Of the effects mentioned, the liquid level stability and initial levelling are important. It is essential to get the liquid level parallel to the zero-field configuration of the gate and anode grids and to ensure that this is maintained as a function of time. The liquid level determines the size of the S2 signal. Hence, if the liquid level were to change during operation then variation in the signal size could arise. Such a process could occur with surface waves for example. To help level the detector and ensure data is discarded in the case of surface disruption, level sensors are installed at the weir. The sensors sit on the outside of the top array between the anode and gate grids and are shown within Figure 5.16. These sensors are capacitive in nature and hence use the same electronics as the Position Sensors. A reference capacitor sits within the shell of the sensor, and another capacitor is allowed to fill with liquid and change the dielectric inside, thus changing the ratio of capacitance.

The sensors have to be calibrated in-situ as they are highly sensitive to the xenon response, which is hard to predict a priori. Theoretically, the response should be linear, as the capacitance is proportional to the dielectric. Hence, a two-point calibration method is preferred. This does not require any prior knowledge about the dielectric of liquid xenon as a function of frequency, for which literature is somewhat poor. One sensor is located below the bottom array and hence was filled with liquid xenon before the ones at the weir. The design of these sensors is identical. The calibration is shown in Figure 5.17 which shows that the ratio of capacitances changed by a factor of 1.81 between the empty and

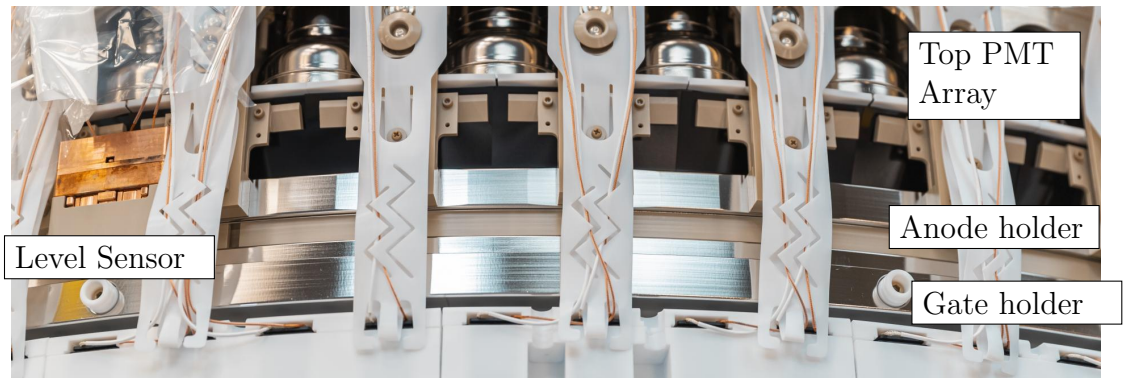

Figure 5.16: The outside of the weir region on the TPC. The level sensor spans the anode and gate planes and hence is sensitive to liquid within these regions. Image was taken by Matt Kapust, SURF

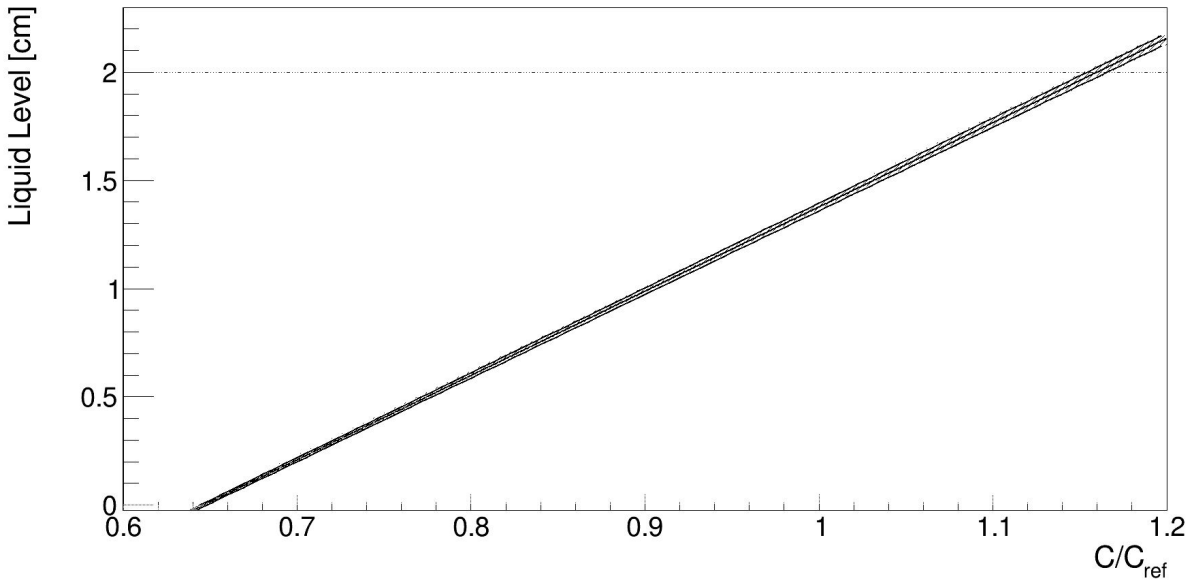

Figure 5.17: Two point calibration performed with a level sensor located below the bottom array.

full levels.

The difference in capacitance corresponds to a difference in the raw electronic readout of over 1080 integer counts. Hence, as the sensor has a 20 mm range then 20 µm precision can be achieved. Comparing this obtained precision with the possible deviations as stated within Section 5.1.2, then the effects of liquid level changing the S2 signal can be tightly constrained. Where disparities do occur, such as ripples and surface waves, these may define periods to be excluded from the analysis.

5.3.2 S2 Pulse Width and levelling procedure

In combination with the S2 response, level sensors were used to level the detector. Hence, this combines the hardware contributions with xenon data analysis techniques, which continue into later Sections.

It is important to ensure that the gate grid is fully submerged to allow the grids to be at a high voltage and not risk damage from electrostatic breakdown. Additionally, a level detector limits some spatial dependence in the S2 signal, as discussed previously. The levelling process was incremental and combines information from different sources, three of which are considered here: the initial liquid immersion times, the absolute level sensor calibration and the S2 signal response.

The time delay between level sensors commencing immersion was considered during the final stages of detector filling. These time offsets can be observed in Figure 5.18. This provides a good approximation as to the detector tilt but relies on the sensors being coplanar. Following estimates based on the geometry and inflowing xenon rate the tilt was estimated to be about 5 mm. This tilt is significantly larger than the uncertainty in the coplanarity of the sensors, and so the results can be used fairly reliably.

To calibrate the level sensors situated at the top of the detector, a different method was required to the bottom level sensor. This was because these sensors were not overfilled. The response was measured during the filling, to check for consistency of readings under the same stimulus. The lower plot in Figure 5.18 shows that the sensor readout is similar when the flow rate of xenon is stable. Apart from 5104LT, then the response per unit time (hence approximately per unit liquid height increment) is consistent within 1% when performing linear regression in each case. Therefore, the calibration relationship shown in Figure 5.17 is similarly used with a correction for the different outputs of the empty sensor. From these values the tilt was also estimated to be approximately $4-5$ mm. The level sensor response provided almost immediate feedback in the levelling progress. Once the first adjustments were made and checked with the level sensors, the level was finely

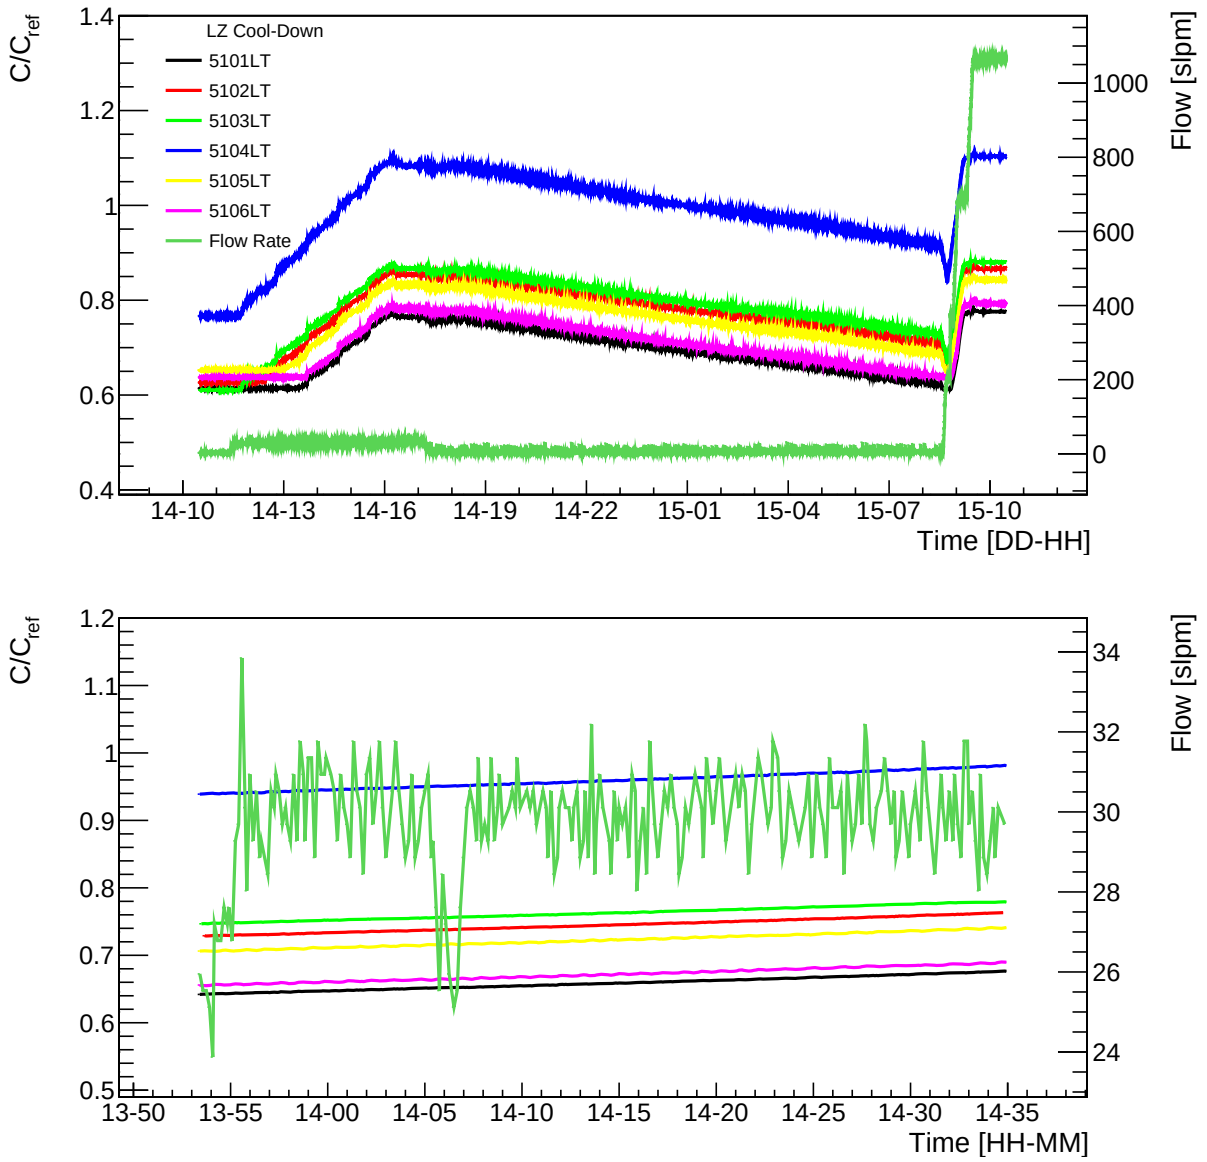

Figure 5.18: The un-calibrated level sensor output over a one-day period, showing initially the filling of the top few cm of the detector extraction region (ER). The flow of xenon was stopped at about 17:15, after which the liquid level dropped. When larger flow was resumed the detector re-filled considerably quicker. The bottom plot shows a smaller time period of about 40 minutes where the flow rate is approximately constant, and shows the response of the level sensors is similar.

tuned according to the S2 response.

The S2 response was measured when the detector was not level and supported the conclusion of a similar level of tilt. The method shown here is for the detector post adjustment, though the principle is the same in the pre-adjustment state. The S2 response is highly sensitive to the tilt of the extraction region and so can be used to check it is level.

In addition to the S2 light yield variation discussed, the effects of the liquid level also have an impact on the S2 temporal pulse width. If the liquid to anode distance varies then there is a difference in the temporal pulse width across the extraction region. The S2 pulse width is preferred over the S2 pulse size because it is less susceptible to variation from the light detection efficiency effects and should be independent of the energy deposit, providing the number of photons detected is large ($>10^3$). The pulse width is also dependent on the depth of the event, as diffusion of the electrons causes temporal broadening. This is considered in extensive detail within Chapter 6. Therefore, to level the detector only the extraction region was biased, restricting the volume where S2s are detectable to the very top of the detector, limiting the diffusion broadening. The anode and gate were set to have a potential difference of 5kV, to ensure the risk of sparking with unsubmerged gate wires was negligible. The gamma event rate was further increased in this region by deploying ^{57}Co sources in the calibration source deployment tubes.

To use the S2 pulse width some basic data quality cuts were applied to remove certain event topologies that are not indicative of detector tilt. The S2 pulses from the extraction region are rectangular. This is expected from the lack of diffusion broadening of the electrons, and the continual emission of light between the liquid surface and the anode wires. The first event in Figure 5.19 illustrates when the electrons near the anode, some amplified emission sometimes occurs at the end of the rectangular pulse. Despite the first event having a good S2 pulse for levelling, all the events in Figure 5.19, lead to problematic measurements of the pulse width based on topology. Firstly, events that were indicative of a multiple scatter topology, or have split reconstructed pulses, were removed. Additionally, a cut to remove events that have a drift time larger than 4 µs was used. This limits the number of events that occur in the gas region that have pulses not useful for levelling. Many events were reconstructed with the S1 and S2 pulses merged, despite them being physically separable, for example, the first event shown in Figure 5.19. This unfortunately biases the measurement of the pulse width if the S1 is large compared with the amplitude from the S2 pulse. The width metric chosen is the full width at half the pedestal height. This was chosen to limit the effects of the amplified

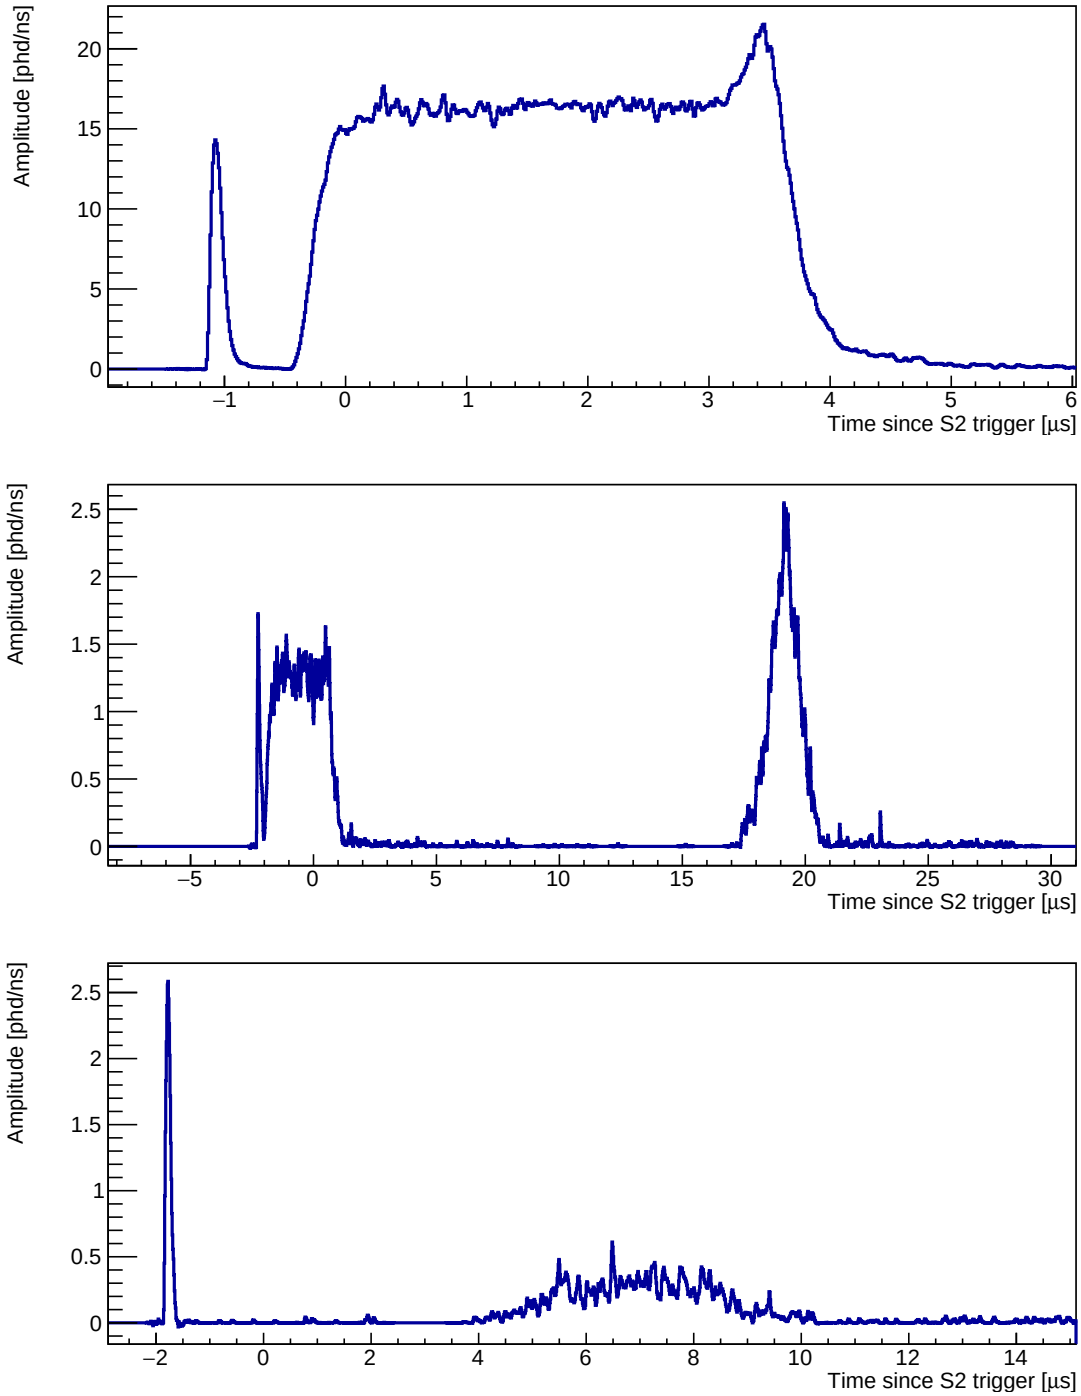

Figure 5.19: Waveforms for various events seen during levelling. In the first case the large S1 pulse preceding the S2 pulse leads to a badly reconstructed value for the pulse width. The main S2 pulse however is typical of the events used within this study. In the second case, a multiple scatter interaction can lead to the wrong S2 pulse being used. For the final case, likely an above anode event, the pulse is not indicative of the liquid to anode distance. This topology can be removed by a drift time cut.

response observed when electrons are closer to anode grid wires and any tails contained within the reconstructed pulse. To remove any populations not covered by these cuts,

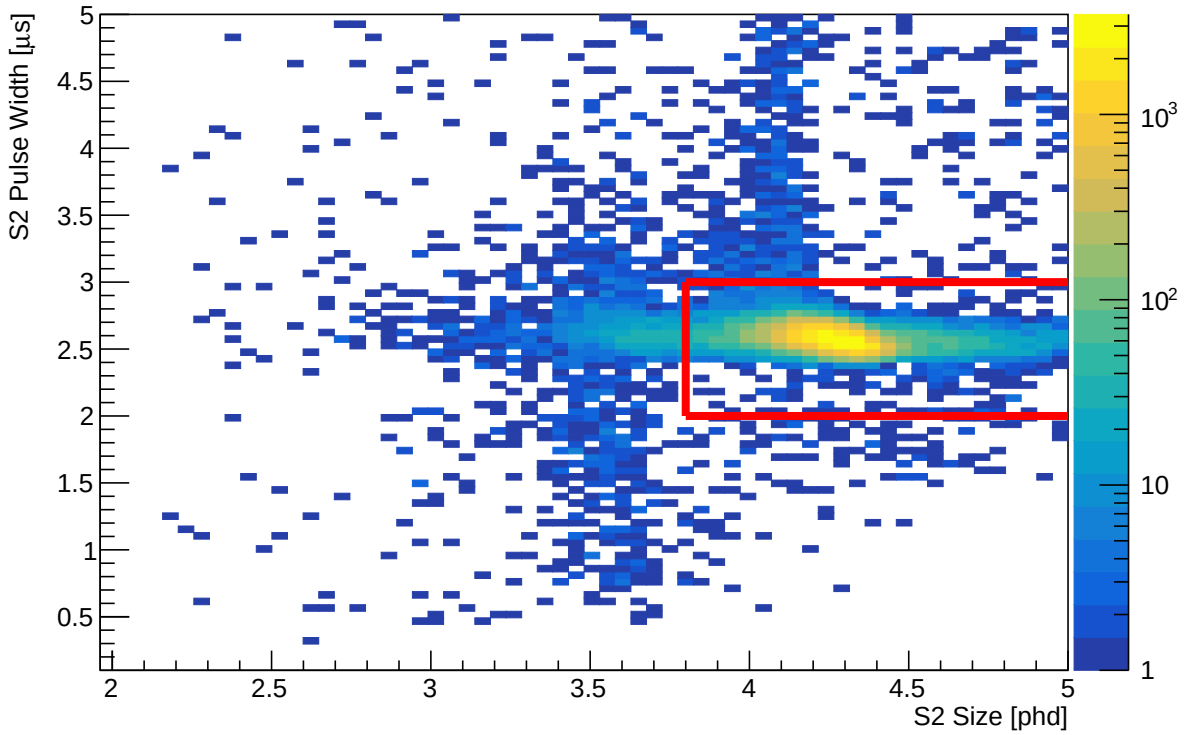

Figure 5.20: Event selection for pulses considered in this study within the red box. The events at smaller S2 sizes are dominated by gas events. Those at lower width tend to be badly reconstructed and those at higher widths have merged S1 pulses that are present.

events were selected in an S2 pulse width-size parameter space as shown in Figure 5.20. Events selected are contained within the red box. Events with a lower S2 size are more likely to have Gaussian pulses, indicative of above anode gas events. Those events with smaller widths, below the selected population, are often badly reconstructed, whilst events above the selected population have a biased pulse width measurement from the merged S1 signal. When performing this selection, the purity of events was the largest driving factor. An efficient selection, whilst preferable, was not the goal for this analysis. The remaining events were considered for position dependence, as shown in Figure 5.21. Here, the remaining tilt is shown to be a few hundred micrometres across the detector. Effects of the grids bending due to electrostatic attraction lead to a smaller pulse width towards the middle of the detector. Therefore, based on this analysis the detector is effectively level, as effects are smaller than the impact of the grids bending. The expected liquid to anode distance is derived using interpolation from a data-driven relationship as implemented

in NEST [99]. This shows some systematic deviation from the expected value of 8 mm. The origin of this systematic is difficult to identify but is possibly due to the choice of pulse width metric, and some broadening from stochastic variation in electron arrival times. The offset could alternatively be due to modelling uncertainties. As the liquid is restricted by the weir height then this systematic is not a concern.

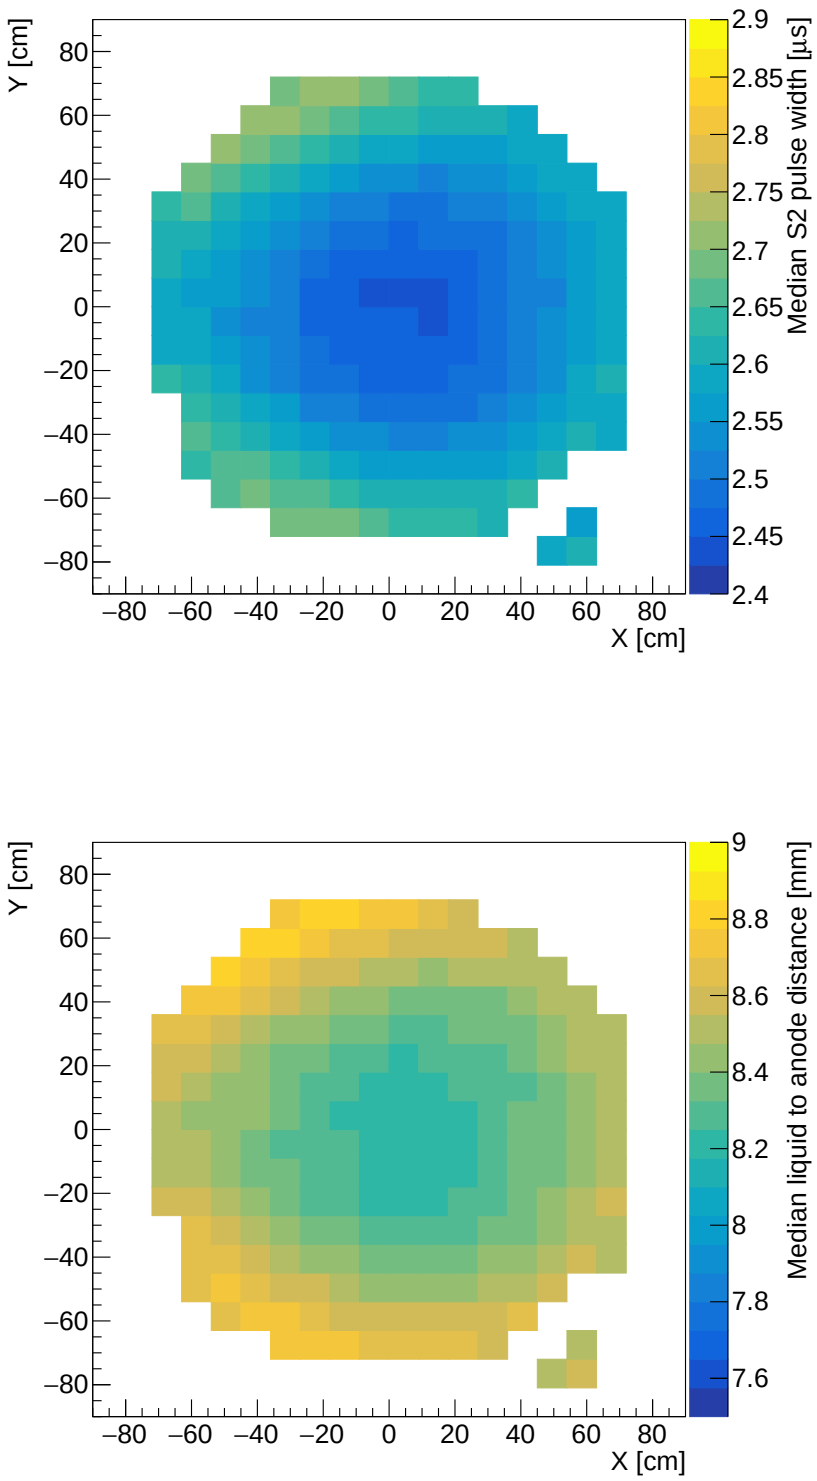

Figure 5.21: (Top) The median pulse width distribution as a function of position after the levelling. (Bottom) The pulse width translated into a median distance between gate and anode. In both cases the larger trend is to observe the effects from the anode-gate electrostatic attraction rather than the slight mismatch in the levelling of the detector.

5.4 S2 Corrections for position dependence

The position dependence on the S2 signal can largely be removed by calibration. Due to the origins of the position dependence previously identified, it is evident that the effects as a function of drift time should be separable from the those in the plane perpendicular to the drift (XY). This will be shown to hold in Section 5.4.3. This therefore means that the corrected S2 ($S2_c$) can be expressed in terms of the measured signal as

$$S2_c = f_{\text{purity}}(t) f_{XY}(X, Y) S2. \tag{5.7}$$

The values of the two factors are obtained as a function of position in the following Sections. All signals are corrected to the S2 response at the top centre of the TPC where the signals are the largest for a given energy deposition. Therefore, $S2_c \geq S2$ for all signals.

5.4.1 Xenon Purity Corrections

The xenon purity can be highly time varying and can greatly affect the measurement of S2 signals. As introduced in Section 5.1.2, the purity depends on the material outgassing and residual xenon purity. Impurities are removed in the circulation system, so the quantity of impurities varies as a function of time. Typically, there will be initial improvement in the purity until an equilibrium is reached. There can be cases where the circulation system has to be taken offline and the purity may fall again during these periods. A need for a standardisation with each day to enable the results across a variety of xenon purities to be used in the same analysis is required.

Mathematically the number of electrons is given by Equation 5.2. As the size of the S2 signal is proportional to the number of electrons extracted then the S2 response also follows a similar relationship and has a solution as a function of drift time t as

$$S2(t) = S2(t = 0) \times \exp(-t/\tau), \tag{5.8}$$

where τ is a characteristic survival time, or electron lifetime and is time varying depending on the purity levels of the xenon. In order to calibrate this effect there is the need to use

a calibration source which can produce mono-energetic energy deposits homogeneously across the xenon detector volume. This therefore means that energy deposits from the source can be used as a 'standard candle' since the yield should be almost identical across the volume. This is because the electric field variation is expected to be small, and hence the recombination fluctuations will dominate over this variation.

It is necessary to be able to select a population of events from a mono-energetic calibration source. As illustrated in Figure 5.1 mono-energetic sources manifest as populations in S1-S2 space. The xenon purity is an effect which is factored in when mapping uncorrected quantities to corrected quantities. Hence, the selection needs to be performed on the uncorrected population. The xenon purity depends on time, so in theory this cut needs to be applied as a function of time as the population changes within this space. It can be partially avoided, by selecting a population with the time independent S1 photons detected, as is discussed later in this Section. There are several mono-energetic sources that would be suitable to perform such a calibration. For example using the ^{83m}Kr and ^{131m}Xe sources would be possible, as well as using alpha decays in the radon decay chains. For the purposes of this calibration the ^{131m}Xe is chosen. The longer half-life of the source (11.8 days) means that activity rates of the source remain high over several days. This allows things to mix adequately, and additionally reduces the need for the injection process to be undertaken regularly. The source gamma emission at 164 keV is above the nominal WIMP search energies.

To illustrate the feasibility of these algorithms the method is applied to simulated data. One month of data has been simulated with accurate backgrounds. In order to select populations various cuts need to be applied. To ensure that the energy deposit is from the photo-peak absorption of the gamma ray, a cut is applied to select single scatter events. These are events that have one identifiable S1 pulse and one S2 pulse within a drift window of about 800 µs. This process relies on the LZ event classification algorithms through which all simulated data passes through. The events passing the single scatters cut are shown in Figure 5.22. To select the ^{131m}Xe source, the 'brightest' population, a

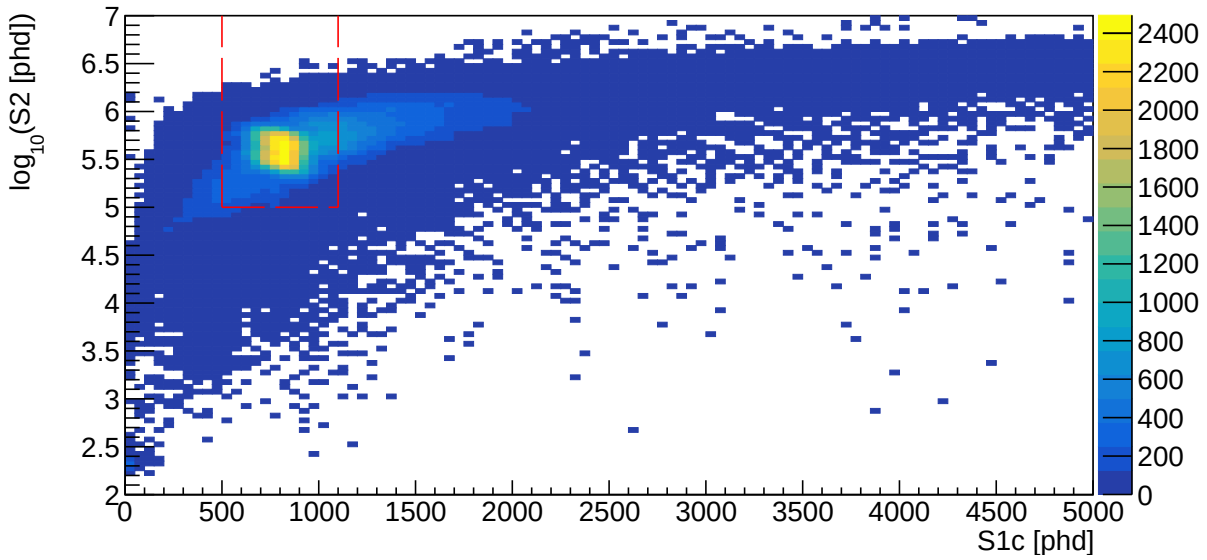

Figure 5.22: Populations of simulated events as a function of the S1 and S2 signal size. The red box shows the events selected for the study considered here. These events are predominantly ^{131m}Xe decays, which form part of the highly populated region within the red box.

cut on the S1 variable is used. The S1 variation with position is mainly related to light propagation efficiencies, which are not time dependent, so the response should be stable. This therefore means that the algorithm can be used with minimal user intervention. The S2 variable, which will change with time has a loose lower cut applied. These cuts are illustrated within Figure 5.22. The events within this population will still contain events that are not from ^{131m}Xe decays. Therefore, there is the need to apply further cuts to increase the purity of the event selection.

Figure 5.23 illustrates that the population of events after those classified as single scatters is not homogeneous in position. This also illustrates the successive application of cut including the radius at $R^2 < 4750$ cm^2 and a drift time cut of $80 < t < 785$ µs. This leaves the remaining events that are approximately homogeneous in the detector. The artefacts shown come from systematics in the position reconstruction software Mercury [105]. With higher statistics (as is planned during detector commissioning) then these effects could be removed by greater tuning of the algorithm. The algorithm must be robust against any events remaining that are not from ^{131m}Xe decays.

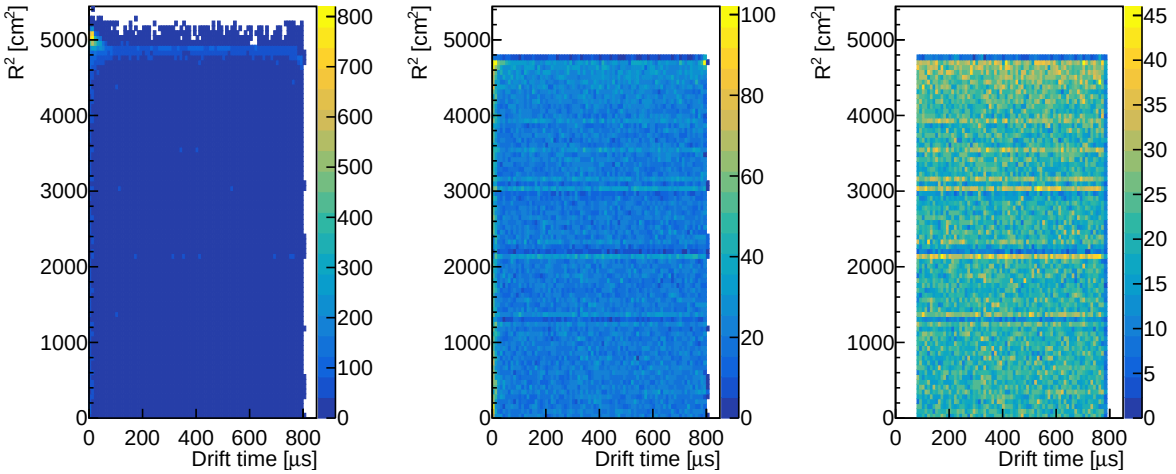

Figure 5.23: Effect of applying position cuts upon the simulated event population to remove the populations of events close to the detector walls and grids. The selected population is shown as a function of the radius from the centre R in the plane parallel to the liquid surface and the depth. Left plot shows all single scatters for one day, middle plot is after applying radial cut, and right plot is after application of a drift time cut ($80 < t < 785$ µs). The banded populations are a position reconstruction artefact and are not physical.

Finally, the issue related to photomultiplier (PMT) saturation requires a novel approach. The PMTs have a maximum voltage that can produce discrete values on the digitiser (Section 3.2). As the S2 signal is highly localised over a few PMTs (Figure 3.2) in a short period, some PMTs are likely to exceed this range for a higher energy deposit such as ones from $^{131\mathrm{m}}$Xe decays. These saturated sensors therefore will measure less light than expected and thus lead to more variation in the S2 size than from the expectation from xenon purity. This is also a drift time dependent effect because the larger uncorrected signals for a given energy deposit at the top of the detector are larger. To mitigate the effect of this, the bottom array contribution to the S2 size is used that does not have this issue. The amount of light in the bottom array is highly correlated with the total S2 signal size.

In order to measure the effect of xenon purity the logarithm of the S2 size is examined against drift time (Figure 5.24). Equation 5.8 can then be transformed as

$$\log_{10}(S2(t)) = \log_{10}(S2(t=0)) - \frac{t}{\tau}\log_{10}(e), \tag{5.9}$$

where the convention is to take logarithms in base 10 on the experiment is used. The

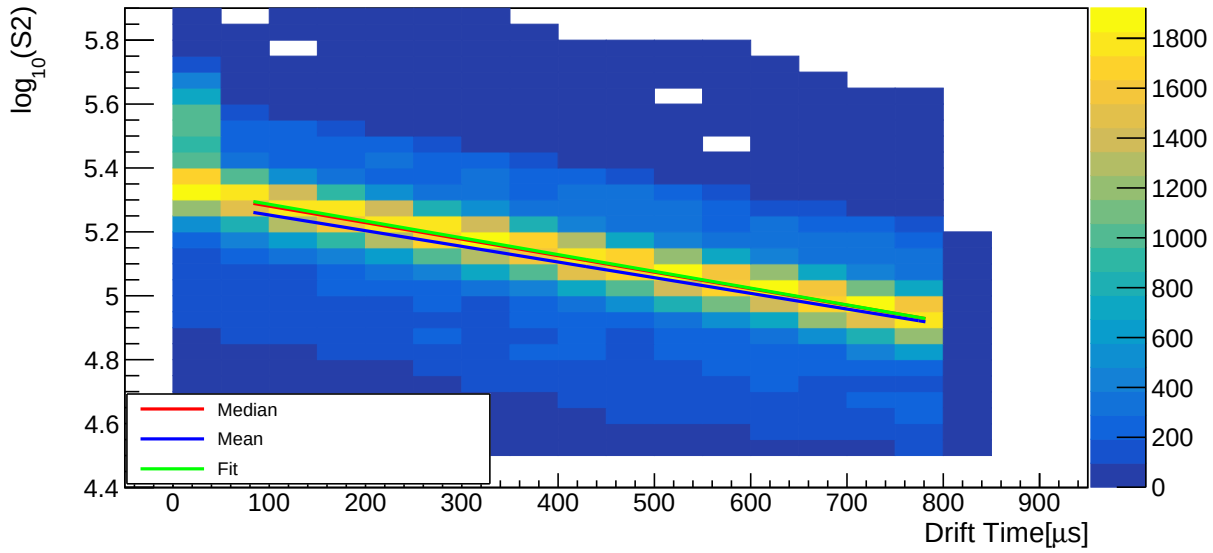

Figure 5.24: Simulated S2 signal size variation with drift time due to xenon impurity attachment. The three lines show the fits in the case of the mean or median of each drift time bin, alongside of the central region is fitted with a Gaussian and then mean is taken. This hence shows a small deviation due to the presence of outliers.

resulting relationship should be linear. In order to be able to fit the trend lines in Figure 5.24 the drift time is split into slices and each slice is fitted with Gaussian distribution. The S2 response is characterised by the product of various efficiencies, hence the logarithm is the sum of these numbers, which is approximately Gaussian under the central limit theorem. The mean of the Gaussian and the error in the fitted mean is used in Figure 5.24, which also shows the results if the mean and median for the whole distribution are used. These metrics are more susceptible to the outliers, which can be drift time dependent. In the case of the mean the values are shifted by populations above and below the trend line and thus increase the measured value of τ. In the case of the simulation considered the value measured is 848 $\pm$ 25 µs, which compares well with the input value of 833 µs. This is much closer than the mean and median estimations. Here only the statistical uncertainties are considered, but the value is clearly consistent. Systematic variation is considered below with the commissioning data. To extract the lifetime τ then the gradient α is given by

$$-\frac{\log_{10}(e)}{\tau} \tag{5.10}$$

and hence the statistical error in the value of τ, Δ_τ is given by:

$$\Delta_\tau = \frac{\log_{10}(e)}{\alpha^2}\Delta_\alpha. \tag{5.11}$$

This process can be repeated over several periods of time and then interpolated for all events to obtain a correction factor for each day. The factor $\exp(t/\tau)$, is then applied to all S2 signals.

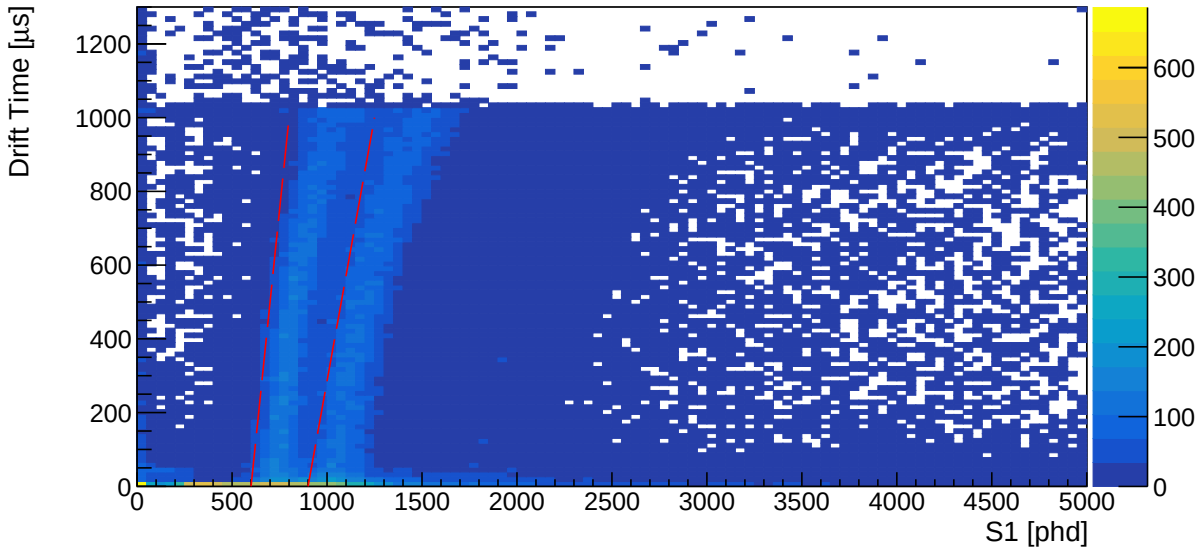

Figure 5.25: Event selection during detector commissioning. The ^{131m}Xe events form a 'band' in this parameter space and are selected based on the depth and S1 signal size. The other population corresponds to events that are from other xenon activation products and are not used within this study.

During early commissioning the algorithm was used to measure the electron lifetime. The anode and gate were biased to ±4 kV and the cathode was set to −20kV. Although this setup is not the final configuration, it does enable an understanding of the xenon purity to be obtained. The purity could have been very poor at this point, whereby the full drift time would not produce S2 signals. In that case there would be a significant proportion of events with an unpaired S1 signal, which would have required a modification to possibly examine only a short drift time region. From looking at the data it was quickly evident that this was not the case, and therefore a method more similar to the one described above could be utilised. At this early stage the S1 signals have not been corrected and so it is not possible to perform the event selection with a cut in S1c.

The ^{131m}Xe population was however still used for this study, which was present from cosmogenic activation during xenon storage above ground. Figure 5.25 shows how the event selection was performed with a cut in both S1 and drift time. The spatial dependence on the S1 signal is due to variation in the light collection efficiency, and is corrected in the quantity S1c. This cut should be time-independent for a given electric field configuration.

Running this data through the same algorithm as the simulated data recovered a similar dependence of the S2 signal with drift time, as shown in Figure 5.26. Again, only the bottom PMTs were considered when examining the S2 size to limit the effects of saturation on the signal. Only the drift time region between 200 and 1000 µs was used during the fitting process. The same binning and radial cuts were used as with the simulated data. These cut variables were extensively verified by allowing them to vary and examining the impact on electron lifetime measurement. This also enables an estimation of systematics to be made. The measured value of 2001 µs is confidently known with a $\pm 5\%$ uncertainty. This uncertainty is dominated by the choice of event selection over those from the fitting.

The analysis critically informed the collaboration that the detector is in a good condition with better than expected purity, so the commissioning process could continue.

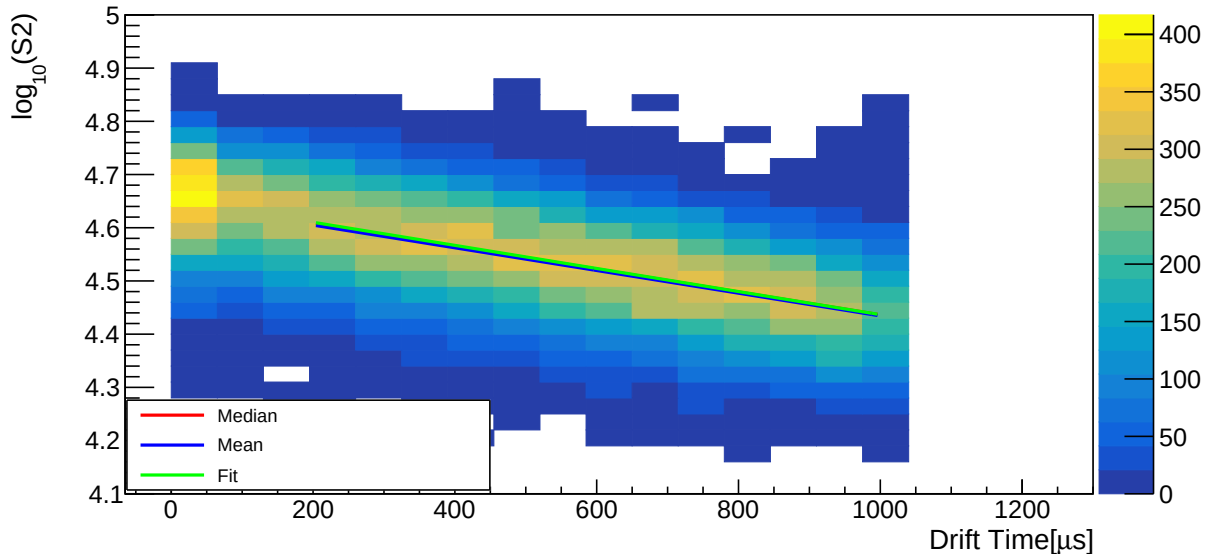

Figure 5.26: The drift time dependence of the S2 signal on 7th October 2021 with the trend lines from the different fit methods. Here the contributions to the S2 signal are only considered from the bottom PMTs to remove the effects due to saturation. This plot is equivalent to Figure 5.24 which is produced with simulated data.

5.4.2 Corrections of signals in the XY Plane

There are additional variations in the XY plane (perpendicular to the drift field) as discussed in Section 5.1.2. As with the xenon purity, a correction factor needs to be applied so that the response is uniform. The method to perform this calibration is very similar to the one discussed in the last section for the electron lifetime. However, because the XY effects should not vary as a function of time, these calibrations can be done in a short period of about one day. The ^{131m}Xe decays could be used to do this; however because the correction factors depend on the geometry of the top and bottom PMT arrays individually, saturation of PMTs becomes an issue, not resolvable by using the bottom PMT array only as before. Additionally, a high resolution map requires many more events than ^{131m}Xe can provide. The ^{83m}Kr source is chosen preferentially and a simulation is performed to examine the feasibility of such a calibration. This is done at about 10% of the events that will be present in real data, and so the resolution should ultimately be an order of magnitude better. The source is discussed in more detail in Section 3.3. The lower energy of the decays reduces the impact of saturation of PMTs. Additionally, due to the short half-life the activity falls quickly over a few days, so the rate can be high for a short period, thus ensuring any time dependent effects are sub-dominant.

Similarly to the ^{131m}Xe population, the ^{83m}Kr source forms a population in S1-S2 space. Events are selected within that population with a data cut. In the case of the electron lifetime both a radial and drift time cut was applied. This is less well motivated with the XY corrections because events can still contribute to the effects from all positions. It is also necessary to correct events close to the walls for some background rejection studies. By having a short calibration period however the contamination from background events should be significantly sub dominant to the ^{83m}Kr population.

Considering the ^{83m}Kr population, a considerable amount of variance in S2 response comes from the xenon purity effects. These effects should be separable from the XY effects (Section 5.4.3) and hence the S2 response can be corrected for the purity before considering the XY effects. This is illustrated in Figure 5.27 which shows the response

uncorrected and then corrected for purity effects. On the right-hand side of Figure 5.27,

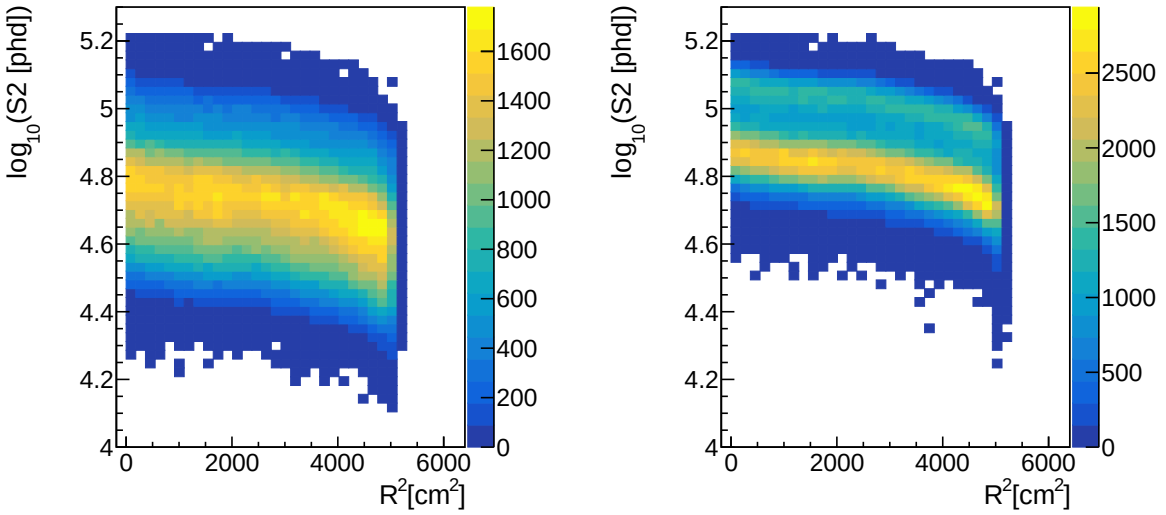

Figure 5.27: The simulated radial dependence of the S2 size. Left plot shows the total detected signal, and the right plot shows the same events with purity effects removed. In removing purity effects it is possible to see that the ^{83m}Kr population can now be resolved into two separate populations.

two separate populations are visible. The more complicated recombination models can produce varying yields from ^{83m}Kr, due to a short half-life of the second decay. Within the NEST [99] simulation, the response is different depending on the time between decays and is based on empirical data. The peak with the highest number of events is chosen to perform the corrections in the XY plane. As Figure 5.27 shows the radial variation of the S2 signal size is dependent on the radius and falls closer to the walls, where the light detection efficiency is lower. The method is the same as in the purity code, whereby each spatial region is fitted with a Gaussian distribution and the error is obtained. Again, this is to remove the effects of biasing with outliers.

In addition to the total S2 size, various other S2 derived quantities can be similarly examined. Figure 5.28 considers the correction factors for different S2 variables. The four sources are various measures of the S2 response. The S2 size is the total light observed following an S2, and is corrected for the purity effects. The S2 size from a population of single electrons can also be considered. The single electron population only considers extracted single electrons so the response does not account for any extraction

efficiencies. Therefore, the radial dependence is, as expected, similar but not identical to that of the ^{83m}Kr population. The position reconstruction algorithm, Mercury, is capable of correcting the S2 response using the light response model. As is evident from Figure 5.28, the algorithm may not produce a dependence consistent with a flat distribution. Potentially more tuning of the algorithms with a greater number of events could improve this. This is however beyond the scope of this work. Additionally, the bottom array is considered separately. This shows that the distribution is smoother, and is a result of a more average light response across the array that is less spatially correlated with the event position. The bottom array can be used for higher energy analyses which would otherwise have issues with PMT saturation.

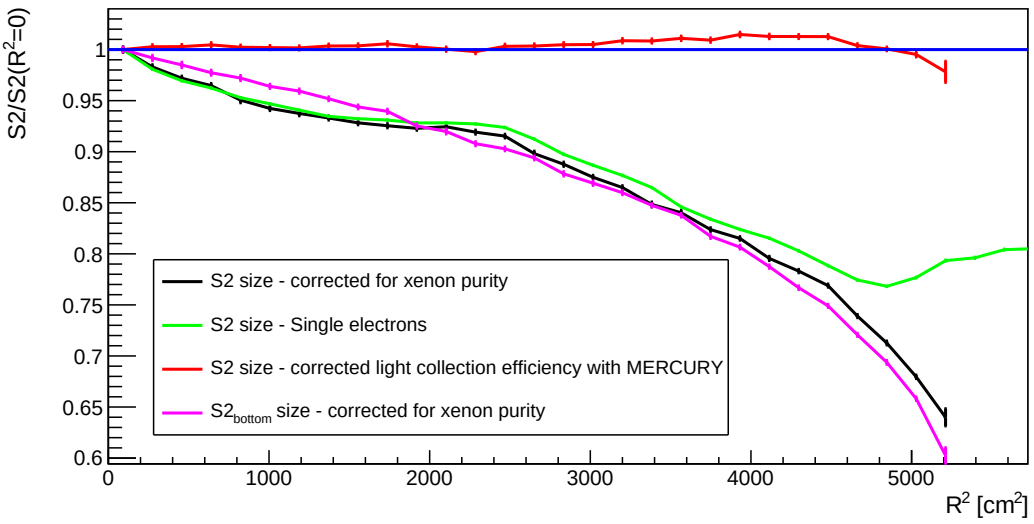

Figure 5.28: Simulated variation of the average S2 response as a function of the radius for different S2 derived variables.

With a higher event rate, the response can be extended to a full two-dimensional map of correction factors. Figures 5.29 and 5.30 show two possible configurations to perform such a correction. The first is using cartesian coordinates, and the second polar coordinates. The Cartesian coordinates are easier to visualise, but this has the issue of lower numbers of events in the outer bins, so fitting is poor. Conversely, the polar coordinate map can model the edges better but leads to ambiguities for defining the centre and also in understanding the physical origins of such effects.

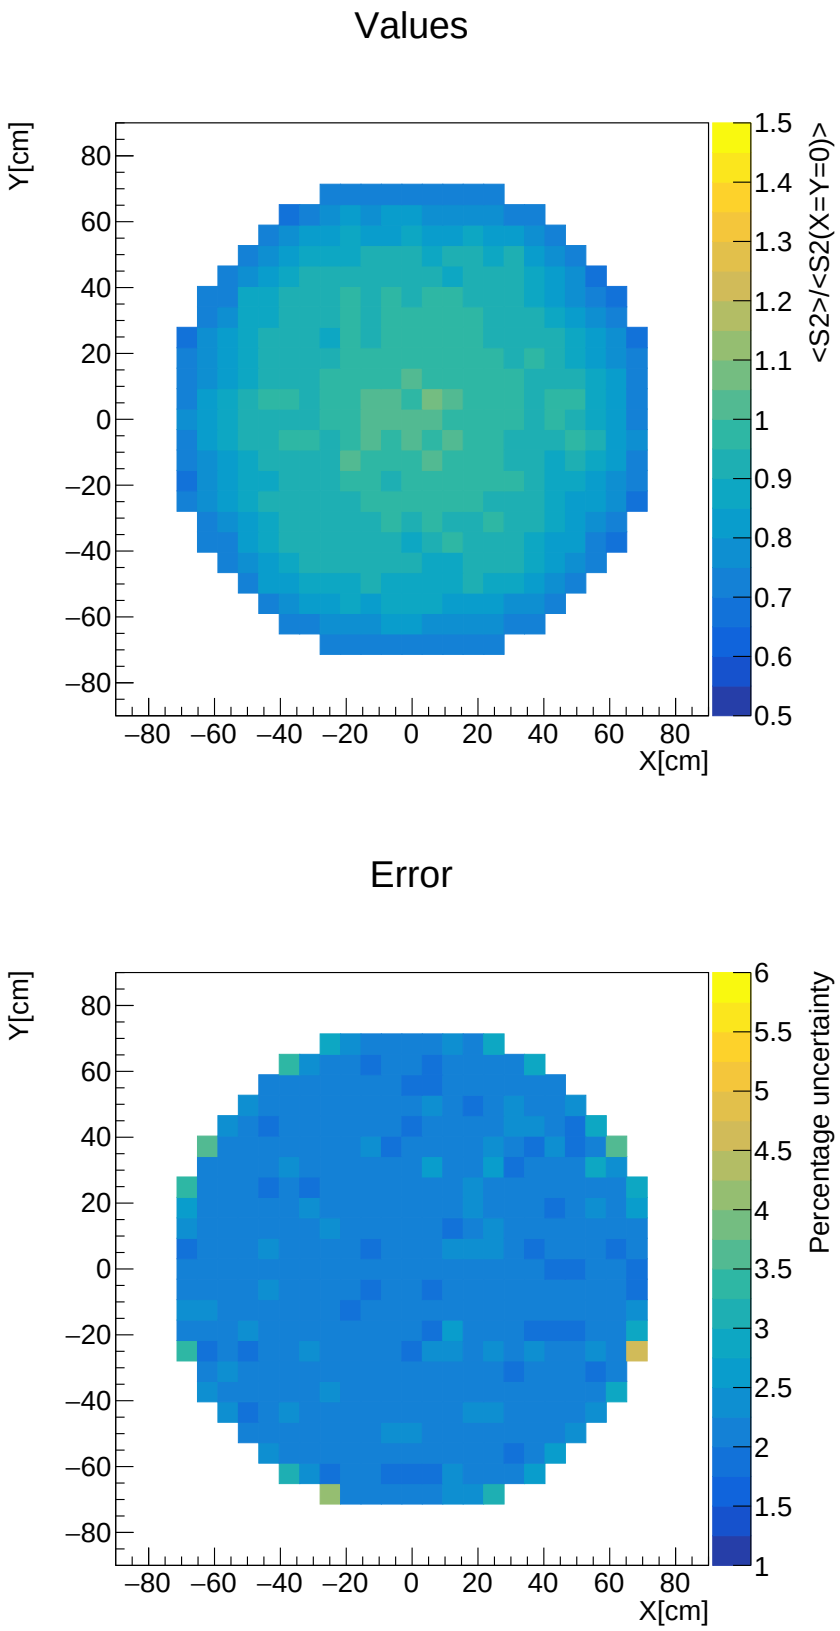

Figure 5.29: Upper plot shows the simulated variation of the average S2 response as a function of the position in the XY plane. The lower plot shows the statistical uncertainty on the fitted parameter.

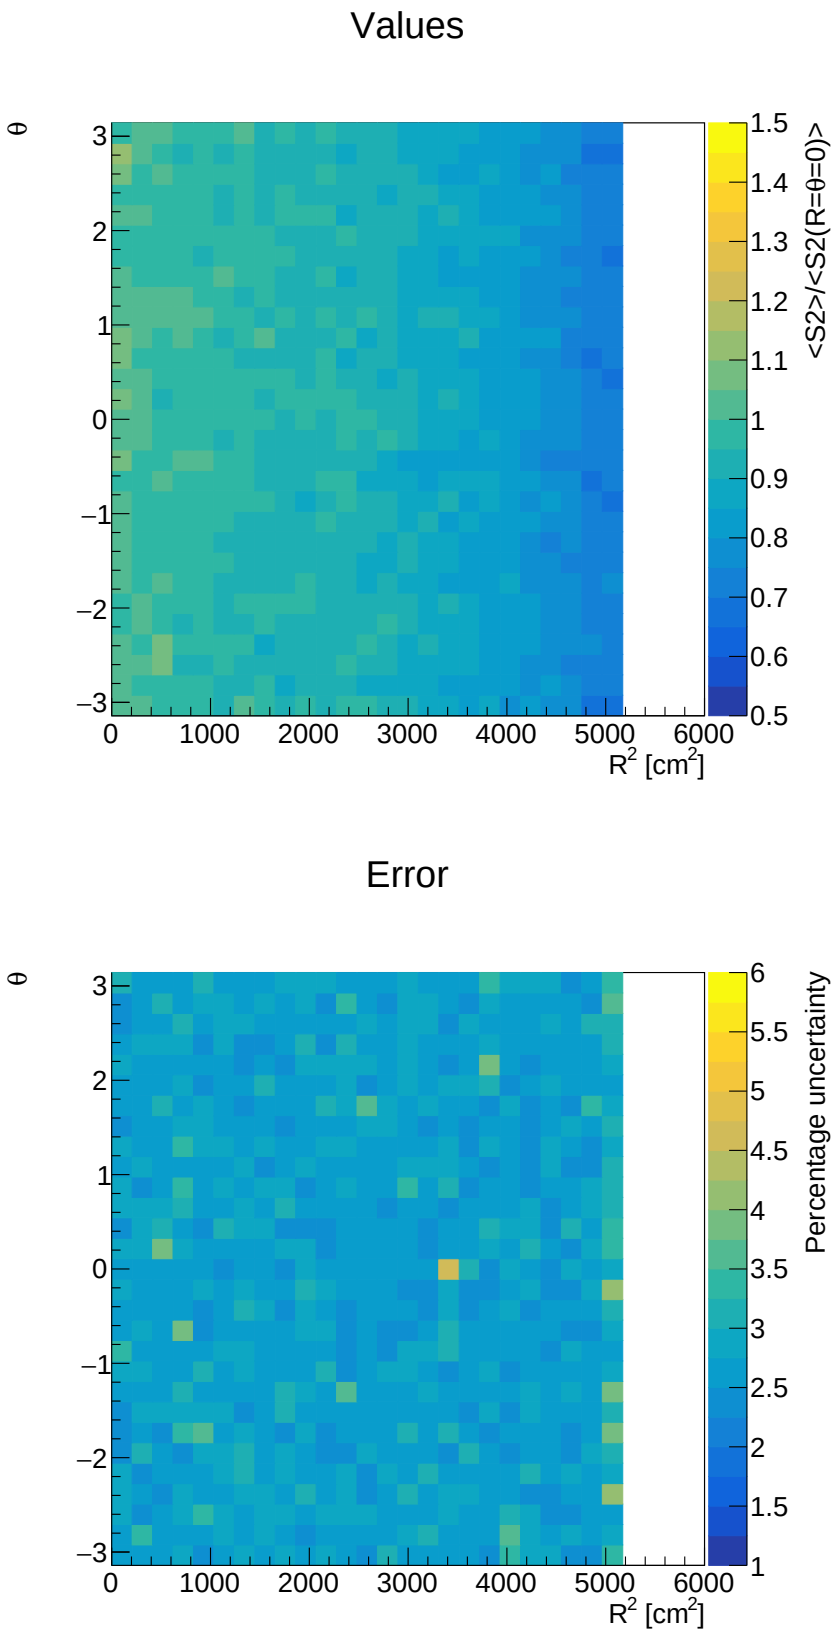

Figure 5.30: Upper plot shows the simulated variation of the average S2 response as a function of the position in the $R\theta$ plane. The lower plot shows the statistical uncertainty on the fitted parameter.

The single electron response, as shown in Figure 5.28 is similar to that of a mono-energetic source. Single electrons, released from photoionisation, low energy interactions and from grid emission are continually emitted within the detector. Hence, during periods where a ^{83m}Kr calibration is not present, the single electron response can track the longer term behaviour of the XY response. This includes specifically the average number of photons detected per single electron S2, and its variation as a function of XY position. The single electron S2 response is dependent on the light detection efficiency and the yield from the secondary scintillation in the extraction region. Hence, a time varying electric field could be tracked with this method.

5.4.3 Considering corrections together

In order to confirm that the corrected quantities are now uniform in drift time and radius then the dependence of the corrected S2 can be examined as a function of drift time and radius. Figure 5.31 shows the resulting simulated distributions where it possible to see that the corrected quantities are now flat as a function of the drift time and radius. This therefore validates that the method works. Such plots form part of the validation procedure when the system is automated and run regularly to monitor the electron lifetime as data is collected.

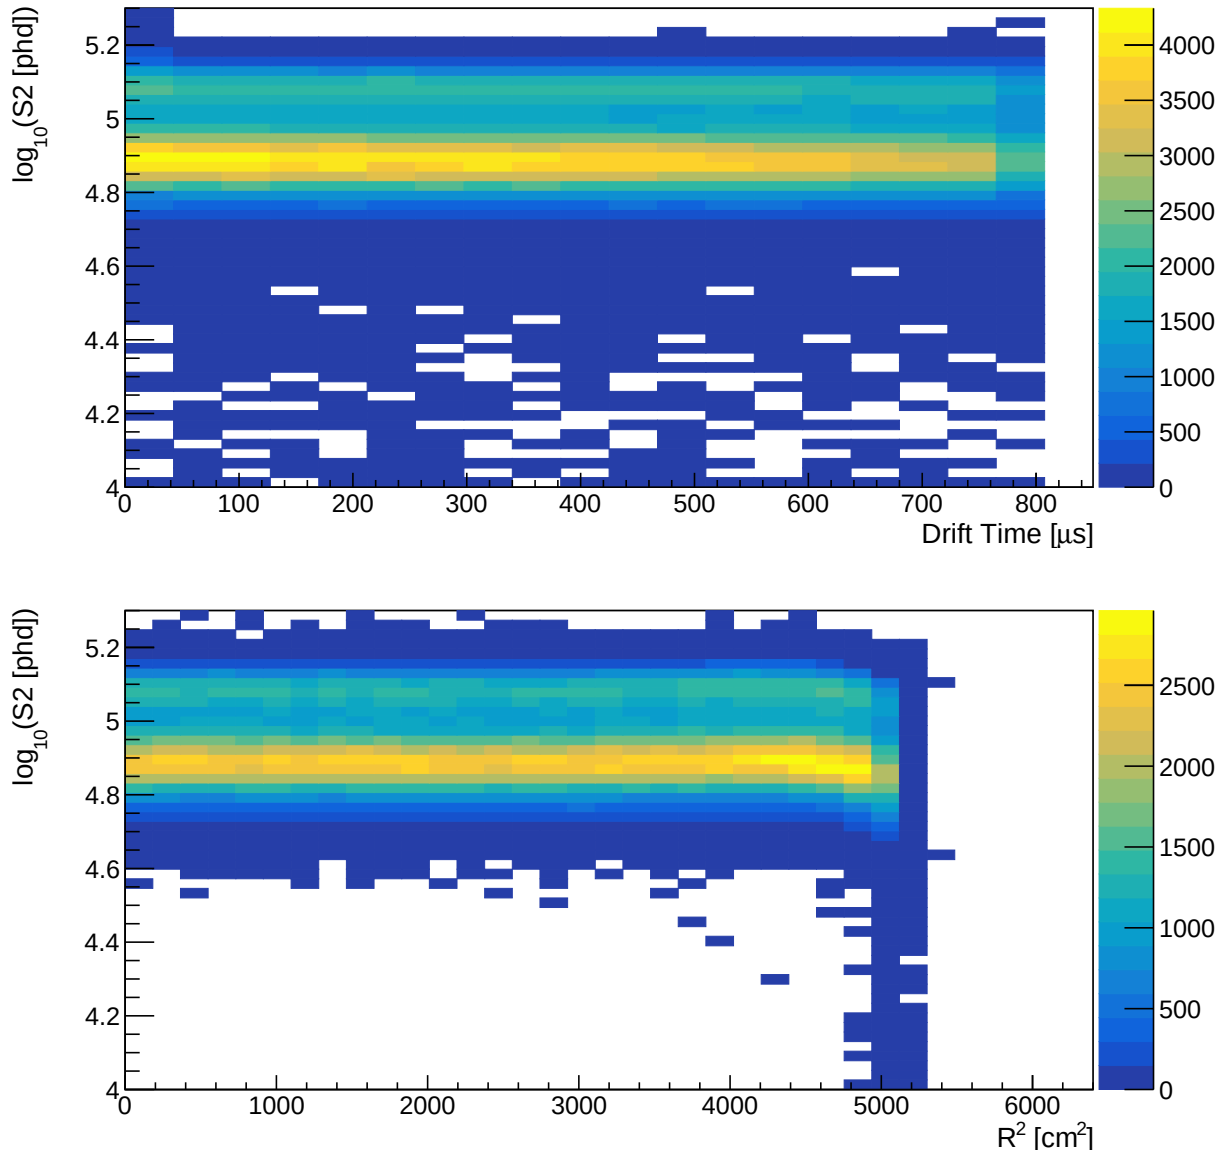

Figure 5.31: Corrected S2 size as a function of radius and drift time for the ^{83m}Kr simulated calibration.

The purity corrections should be, to first order, independent of the XY effects. This is because the purity corrections are largely driven by the drift time, whereas the XY variation is driven almost solely on the extraction region behaviour. Potentially some variation may arise according to non-uniform fields close to the walls where the electron trajectories may become curved. Additionally, this may lead to some charge yield variation in the detector. To estimate how large these correlations are a radial examination of the S2 response for several drift time slices can be evaluated with the ^{83m}Kr simulated data. Effectively this then can be used to produce a higher dimensional treatment where both the purity and extraction region effects are considered together. As Figure 5.32 shows, the radial variation in the S2 response is very similar for each drift time slice. Additionally, Figure 5.33 shows for different radial regions how this then varies with the drift time. Other than for the largest radii, the data shows that any dependence is negligible. Even in the case of the outer most radial bin, effects are at the percent level and so can largely be ignored. It is evident that the variance in each radial bin leads to a much larger effect than this.

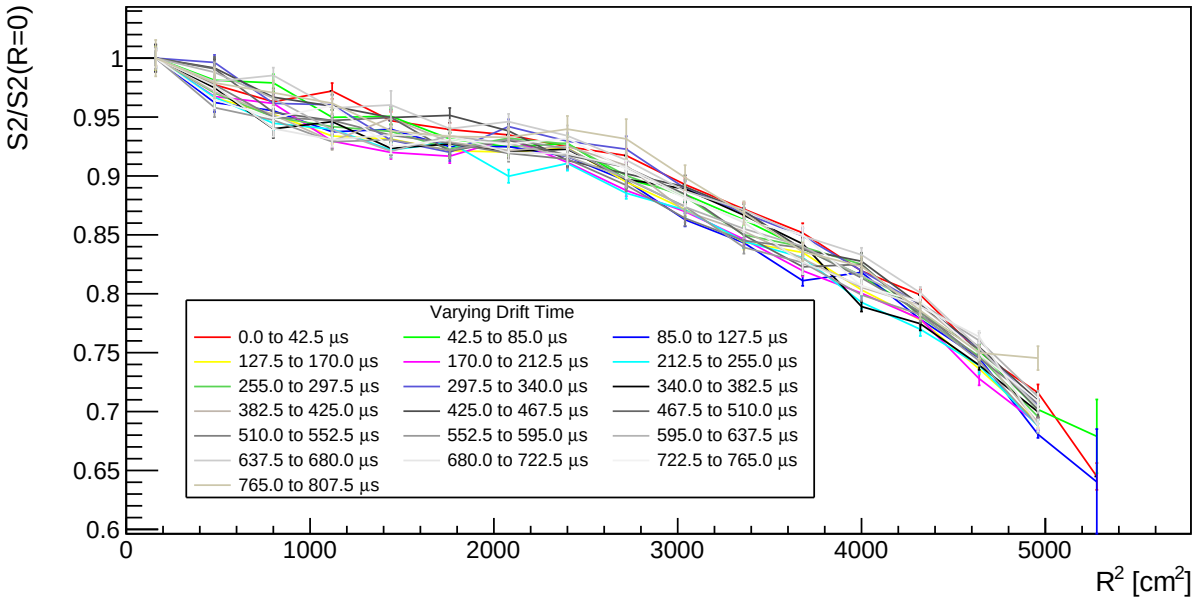

Figure 5.32: Radial variation of average S2 size for different drift time regions. This shows any dependence on the drift time is negligible.

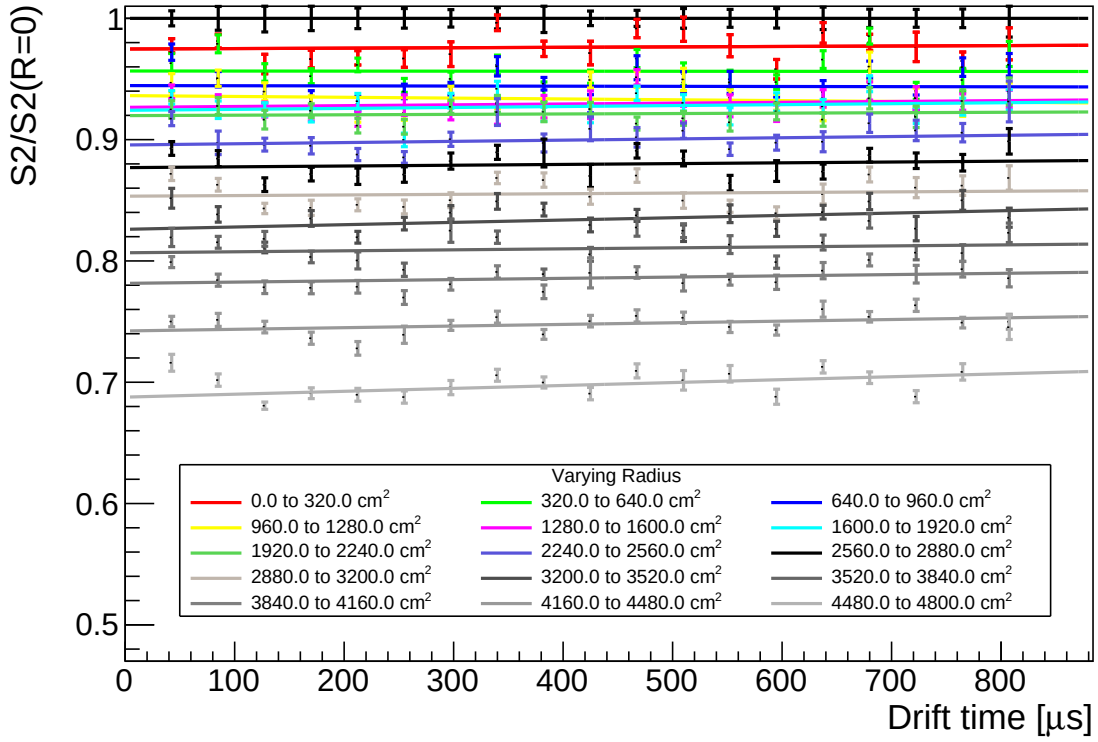

Figure 5.33: Drift time variation of average S2 size for different radial regions. This shows any dependence with drift time is negligible.

5.5 Conclusion

The variation in S2 response across the detector presents an issue for both the energy resolution and the nuclear-electronic recoil discrimination. If these features remained a search for Dark Matter would not be possible. The position sensors mounted within the xenon gas space showed that the contraction of the TPC during the cool-down from room temperature is uniform within experimental uncertainties and so should not be a dominant component to this variation. The level sensors are used to examine the liquid level and can identify disparities leading to signal variation. This Chapter also includes a consideration of the levelling process with both the level sensors and the S2 response. A calibration of the detector to remove the position dependence in the S2 response is also performed on simulated data illustrating the feasibility of this analysis.

Chapter 6

Improving the sensitivity of LZ to low energy signals

This chapter discusses the development of a method designed to improve the sensitivity of the LZ experiment to lower energy physics signals. This includes ^{8}B neutrinos and sub 10 GeV/c^{2} Dark Matter. Signals produced from interactions of such particles are small; in many cases the S1 is undetected, either as any detected photons are below the three fold coincidence requirement or are not detected at all. Therefore the drift time cannot be calculated using the traditional time between S1 and S2 signals. This means that rejection of background events from the top and bottom of the detector close to the gate and cathode electrodes is lost. The temporal distribution of electrons extracted from the liquid is, through diffusion, dependent on depth. The width of the S2 pulse proves to be a key measure of the extent of diffusion broadening, and hence can be used as a proxy for the depth of an event, which can help with background discrimination. A study is performed into the capabilities of doing this within LZ where the low-energy sensitivity is considered.

6.1 Motivation

6.1.1 Low energy physics

Whilst LZ considers mainly WIMPs as a Dark Matter candidate, there remains the possibility that a signal exists just below the nominal reach of the experiment with a mass of around 10 GeV/c^2. Theoretically, candidates of this mass are slightly less well motivated due to their failure to satisfy the WIMP miracle, but still provide an interesting region of phase space to consider which is as yet unexplored.

Additionally, signals from low energy neutrinos, specifically those within the solar interactions of ^{8}B in the sun, manifest themselves through coherent neutrino recoils around the nominal threshold of LZ [112]. These are doubly interesting; they provide validation of the analysis methods for low energy signals, and with a low energy search could provide insights into solar physics [140]. This would also constitute the first conclusive observation of coherent nuclear scattering with solar neutrinos.

As with any detector, when observing energy deposition processes, a point is reached when examining increasingly small signals, that the events can no longer be detected. Specifically this is when the scintillation light or ionisation charge is not produced or they are in small enough numbers as to be indistinguishable from background processes or noise. Specifically this occurs around the $1 - 10$ keV point for nuclear recoil interactions in the nominal LZ search. Here, the limiting factor is usually the S1 signal due to the low probability of light being produced and detected. Detection efficiency is of the order 10%. Additionally, dark counts originating from intrinsic properties of the PMTs mean a three-fold coincidence requirement is needed on the S1 signal to reduce the probability that they accidentally pile up and create a 'fake S1' [141]. This requirement needs 3 photomultiplier tubes to register a signal above threshold within about 150 ns. Hence, close to the threshold there is a probability that the scintillation signal is not large enough to meet the requirement to be classified as an S1. So the event cannot be classified as a single scatter with 1 S1 and 1 S2 signal. However, as the single electron detection

efficiency is much larger, 5 to 20 electrons within the S2 signal can be observed for such events and are hence called 'S2-only'. It is possible to have some events that have a measurable S1 and S2 signal, where the S2 signal is within this 5 to 20 electron range. In analyses considered here, these events are included, but are considerably sub-dominant.

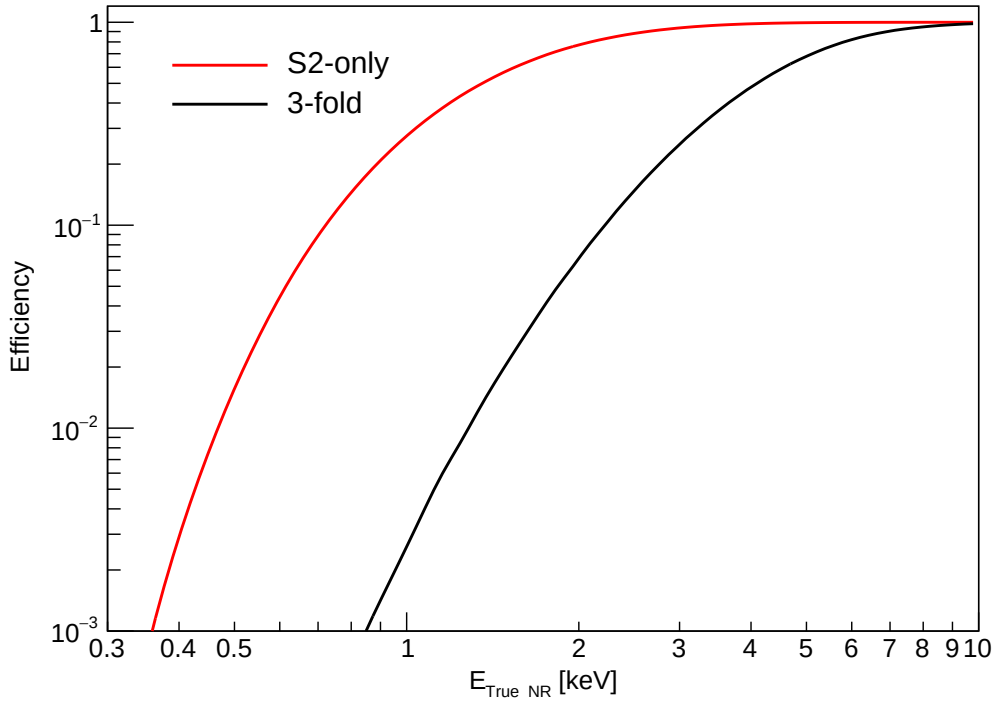

Figure 6.1: Efficiency of detecting nuclear recoil scatters from simulated data as a function of the true recoil energy. Mean charge yields are assumed and a cut is placed at 420phd in S2 size to ensure adequate radial position resolution.

In considering such S2-only events the energy threshold of the experiment is improved, and hence more of the signals from ^{8}B neutrinos and lower mass Dark Matter are then within reach. This can be seen in Figure 6.1, which shows that for lower nuclear recoils, the acceptance improves for the S2-only case over the nominal search. The only cut that has been applied on the data is to require 420 phd such that the radial position resolution is better than 1 cm. This is discussed in detail within Section 6.2.5.

6.1.2 Disadvantages of lowering the threshold with an S2-Only search

Lowering the energy threshold with an S2-only search does have disadvantages, however it is still considered a worthwhile endeavour. For example, by removing the S1 requirement, new backgrounds are introduced into the signal searches. For S2-only events, the drift time for the electrons forming the S2 signal is unknown. This means that the detector cannot be fiducialised in the direction parallel to the drift field. Hence, on top of the standard backgrounds as discussed in Section 3.4 which include the ^{220}Rn and ^{222}Rn decay chains, ^{85}Kr, and neutrinos, additional background events originate from decays on or near the gate and cathode grids as well as liquid surface events. The decays near to the grids, typically originate from radioisotopes which are embedded within, or upon, their surface. When these decay, they can cause a energy deposit in the detector. This is particularly problematic due to the high electric fields close to the wires which can accelerate charges. Several mitigation strategies are implemented to try to reduce these contributions, such as by constructing the detector within a low radon environment and grid passification in acid [103], but some remain from material procurement. Liquid surface events originate from x-rays which can travel longer distances within the gaseous xenon and can interact within the top few centimetres of liquid. All of these backgrounds are doubly problematic as the usual nuclear-recoil electronic-recoil separation that arises from the knowledge of the S1 and S2 signal magnitude (as discussed in Section 3.1.2) is also lost. Hence, the electronic-recoil background is superimposed on top of the nuclear-recoil signals from low mass Dark Matter and ^{8}B neutrinos. Within this Chapter the issue of depth determination

is discussed at length, outlining attempts to reduce grid backgrounds. Specifically, this uses pulse shape parameters, which are sensitive to the diffusion properties of the electron cloud, that leads to the production of the S2 signal.

6.2 Developing an S2-only model based on pulse width

6.2.1 Physics of temporal width of S2 pulse

In order to look at pulse shape characteristics, the physics leading to pulse shaping needs to be considered. The production of quanta (photons and electrons) occurs following an interaction with a xenon atom. The region from which these quanta emerge is highly localised to the vicinity of the atom ~ 0.1 µm [92]. Hence, initially the electron cloud can be considered point-like compared with the meter scale time projection chamber. The electrons produced will, under the influence of the electric field in the bulk, drift upwards towards the extraction region. During this drift period the electrons will diffuse outwards with respect to a co-moving coordinate system propagating at the mean drift speed. In this regime the diffusion behaviour is constrained by the continuity equation [142],

$$\frac{\mathrm{d}\rho(\vec{\mathbf{x}},t)}{\mathrm{d}t} = -\boldsymbol{\nabla}\cdot\vec{\mathbf{J}}(\vec{\mathbf{x}},t), \tag{6.1}$$

where ρ is the electron density, dependent on the position ($\vec{\mathbf{x}}$) and time (t), and $\vec{\mathbf{J}}$ is the vector current density of electrons. Combining this with Fick's law, which is separable into three dimensions [142]

$$J_i(x_i,t) = D_i\frac{\partial\rho(\vec{\mathbf{x}},t)}{\partial x_i} \quad (i=x,y,z), \tag{6.2}$$

where D_i is the direction-dependent diffusion coefficient (assumed here to be constant in space and time), in the dimensions x, y, and z, gives

$$\frac{\partial\rho(\vec{\mathbf{x}},t)}{\partial t} = D_i\frac{\partial^2\rho(\vec{\mathbf{x}},t)}{\partial x_i^2} \quad (i=x,y,z). \tag{6.3}$$

This remains separable if the diffusion in one dimension does not depend on the other dimensions, an abiding principle of Fickian diffusion. In the longitudinal dimension (parallel

to the electric field), the density of electrons can therefore be modelled as

$$\rho(z,t) = \frac{N}{\sqrt{2\pi D_L t}} \exp\left(-\frac{z^2}{4D_L t}\right) \quad t \geq 0, \tag{6.4}$$

where N is the original number of electrons. Equation 6.4 is a Gaussian distribution with standard deviation on z position of $\sigma_z = \sqrt{2D_L t}$. Longitudinal diffusion manifests itself in the temporal delay of electrons originating from the event. Hence, it is more appropriate to consider broadening in time, $\sigma_t = \sqrt{\frac{2D_L t}{v_d^2}}$, where v_d is the drift speed. The S2 signal originates from the electroluminescence of these electrons. The temporal distribution of the resulting photons is a convolution of electron arrival times with the electroluminescence region effects. To first order, the latter can be approximated to be a continual emission of light while each electron transverses the region between the liquid surface and the anode, meaning the light has a probability distribution that is a boxcar function. This is developed later in this chapter. The S2 signal distribution is therefore a combination of diffusion and extraction region behaviour. These effects should be uncorrelated with one another and hence the pulse width can be estimated mathematically as has been derived in Reference [143],

$$\sigma_e^2 \simeq \sigma_t^2 + \sigma_0^2 = \frac{2D_L t}{v_d^2} + \sigma_0^2 \, . \tag{6.5}$$

Here, v_d is the drift speed, nominally ~ 1.8 mm/µs, D_L, the longitudinal diffusion coefficient is $\sim 15 - 30$ cm/s, and the σ_0 is a constant, typically ~ 300 ns that originates from the single electron width, the smallest constituent of an S2 signal. Equation 6.5 provides a functional formulation that links the variance, σ_e^2, of the S2 signal and the drift time that can be utilised to provide some depth information. The relationship between the drift time and depth of the signal is highly correlated and is discussed in detail within Section 6.2.3.

6.2.2 Determining the pulse variance, σ_e^2

The variance of an S2 signal can be extracted from the waveform itself. For Gaussian waveforms, as discussed in the previous section, the procedure is relatively simple as the waveform can be fitted or the width calculated based on a pulse metric such as the full width at half maximum. For example, the simulated pulse in Figure 6.2 shows an approximately Gaussian waveform from electrons that have drifted over most of the TPC length. This Gaussian behaviour is more generally true for S2 signals where the interaction point is far enough from the gate that the diffusion behaviour dominates over the approximate box-car distribution. This is the case for electron clouds drifting for more than about 150 µs.

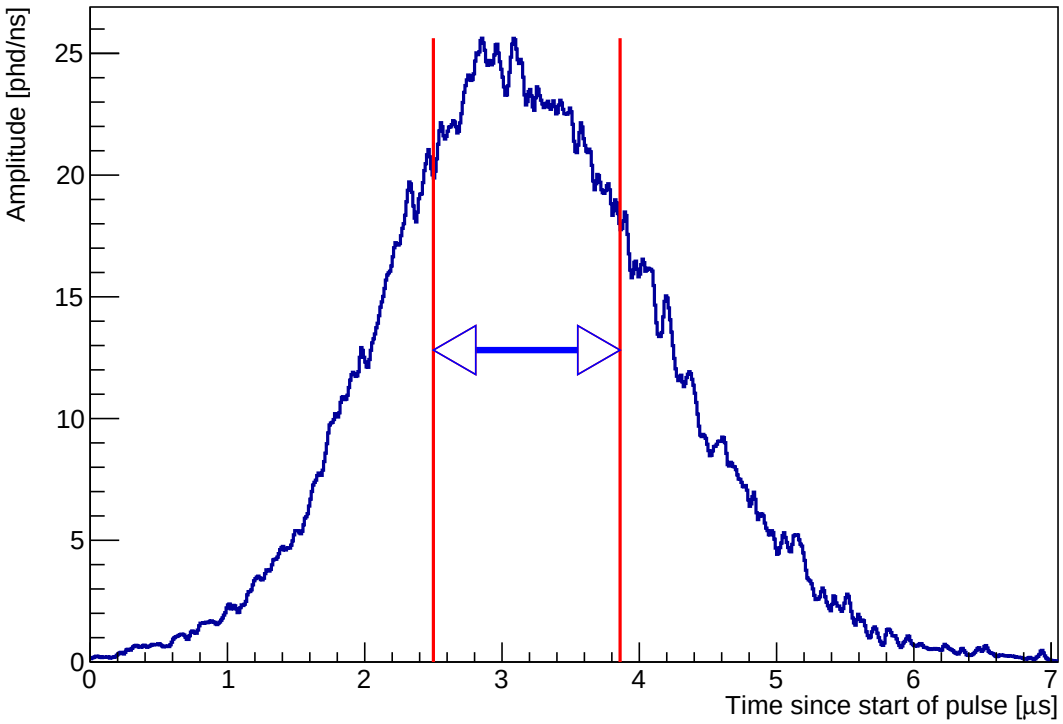

Figure 6.2: Gaussian S2 pulse from a simulated ^{83m}Kr decay 10mm above the cathode. Vertical lines illustrate the central 50% of the total charge accumulated on PMTs in the pulse, which can be used to estimate the width of the pulse.

Calibrating the drift time - width relationship

As part of the suite of calibrations performed in LZ, the ^{83m}Kr isotope is well suited to characterise this diffusion behaviour. The S2 signals originate from more than 1000 electrons, so the electrons fully populate the diffusion Gaussian. Events also have an associated S1 that can be used to determine the depth of the event through the drift time. This is defined to be the difference between the start times of the S1 and S2 pulses. The injected ^{83m}Kr decays away after a few days, hence the calibration produces high-statistics in a short period. This has the advantage that time-dependent detector effects, such as the time dependent attenuation of the signal due to the level of xenon impurities, can be neglected (more details in Chapter 5). This calibration also has the advantage that saturation effects are expected to be small, however this remains possible for pulses originating from interactions at the top of the detector, which are temporally denser and have little attenuation in the Xe bulk due to impurities. To mitigate against this effect, low-gain waveforms can be used for the width calculations for the calibrations. Some disadvantages of using this source exist and need to be addressed. Firstly, the pulse may have some small non-Gaussian behaviour as the events are made up of two successive decays rather than just the one. Due to the short ($\sim$ 200ns) half-life between successive ^{83m}Kr decays, this effect is expected to be small compared with the pulse width which is typically an order of magnitude larger. Additionally, delayed extraction electrons (electron trains) can be included within the main pulse definition and hence bias the width. This therefore requires a cut to be applied to remove events where this happens in the calibration dataset. Using a measurement of the width that is less affected by these effects can also be considered. A study using simulated Kr events within the LZ detector, using the LZ framework, is performed to illustrate the possibilities with real data.

Events are selected that corresponded to the ^{83m}Kr source by applying cuts in the space governed by the S1 pulse size and the logarithm of the associated S2 pulse size. The standard procedure that has been used for placing a cut on mono-energetic source

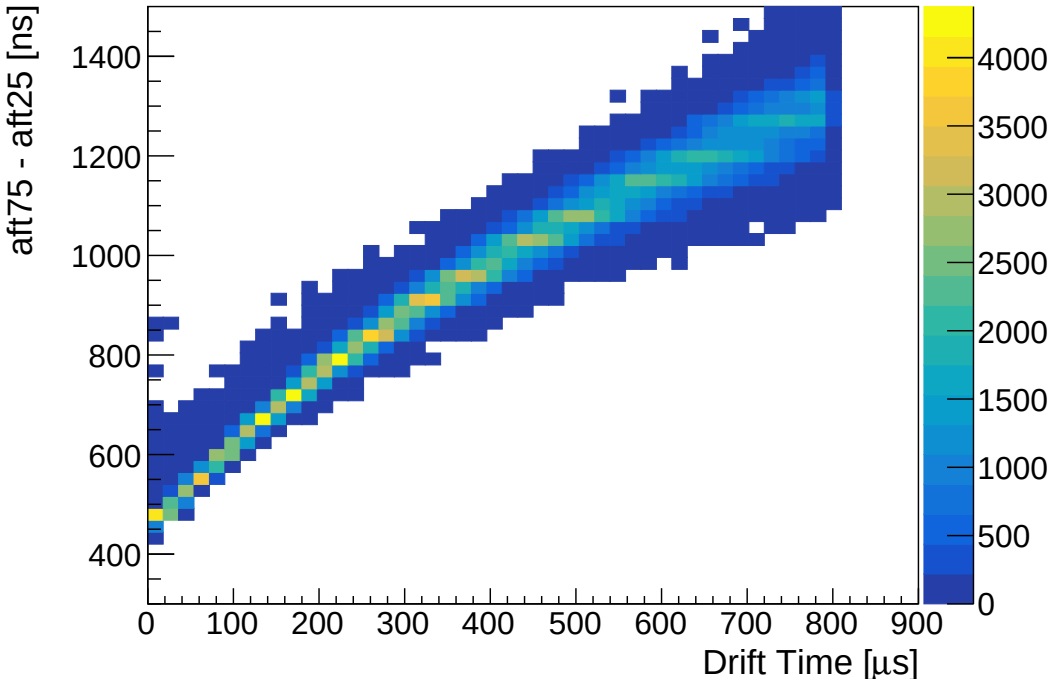

Figure 6.3: S2 width as defined using the difference in the times of the bounds of the central 50% of the S2 pulse as a function of the drift time for the simulated ^{83m}Kr population. The aft75 and aft25 denote the time at which 75% and 25% of the total S2 signal is measured.

data is discussed in Section 5.4. A drift time dependent cut is also applied to the pulses to remove long tails where electron trains have been merged into the main pulse. Only the central region of the detector is used to perform this calibration ($R^2 < 4000$ cm^2). This is due to a discontinuity in the temporal pulse width that arises in the distribution at larger radii, and is believed to be an artefact of the electric field simulation. It remains to be seen if this will be the case with real data.

In order to parametrise the diffusion effects, the temporal pulse width is calculated based on the time difference between 25% and 75% of the size of the S2 pulse (known as aft75-aft25, where aft refers to the 'area-fraction-time'). The reason for this is discussed later within this Chapter, but is used to mitigate against the electron trains and additional complexities in the data. The pulse width as a function of drift time is shown in Figure 6.3. Assuming the pulses are approximately Gaussian in nature, the σ_e of the pulse can then be calculated from the difference in area-fraction-times, using the property that

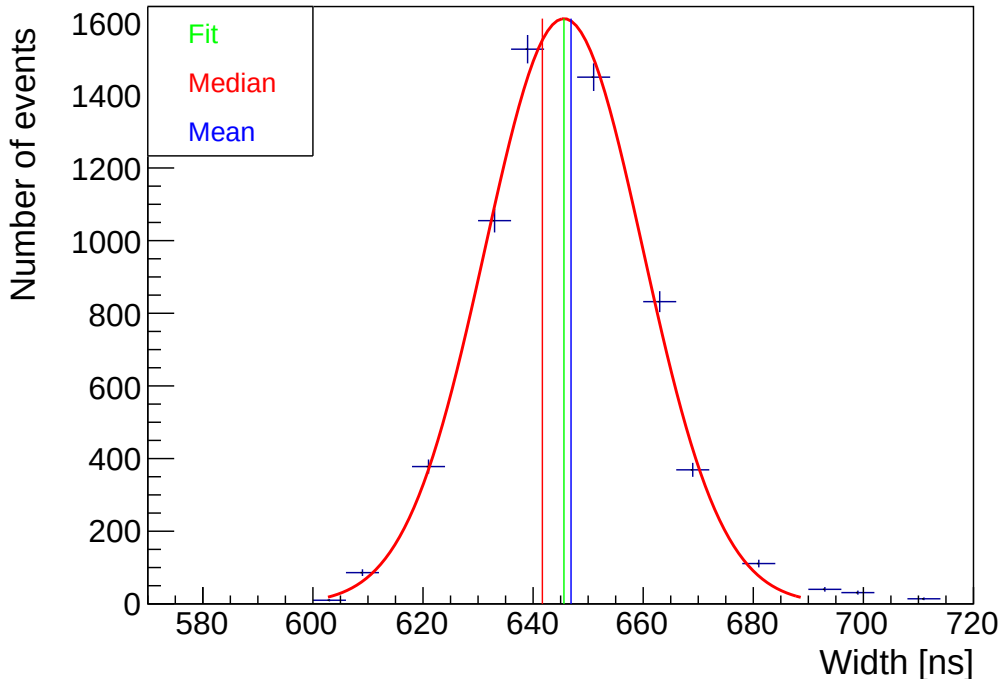

Figure 6.4: The $90-106$ µs drift time bin from Figure 6.3 to illustrate the fitting method. Values obtained from the mean of the Gaussian fitting are compared with the mean and median of the distribution to illustrate biasing due to outliers.

the central 50% of the pulse size is at 1.34 standard deviations from the mean for the aft75-aft25 metric. In each drift time slice the distribution is taken and a Gaussian is fitted (Figure 6.4) to remove statistical effects of outliers. The mean and median are also included to illustrate this bias. The mean and error are extracted from the fitted Gaussian and propagated forward to enable Equation 6.5 to be fitted to the higher drift time regime, as shown in Figure 6.5.

The values of the parameters D_L and σ_0 are reconstructed to be 17.2 ± 0.1 cm^2/s and 355 ± 43 ns respectively. The 'true value' of D_L in the simulation is field, and hence, spatially dependent [144]. Above a field of about 150 V/cm the measured value asymptotes, likely when the diffusion behaviour dominates over the coulomb repulsion effects within an electron bunch. LZ is comfortably within this regime and so treatment here is valid. The variation with field is small enough within this higher field regime, that the true value can be taken as a constant $\sim$ 17.1 cm^2/s, where variations occur in the second decimal place. Therefore, the input value is in agreement with the one

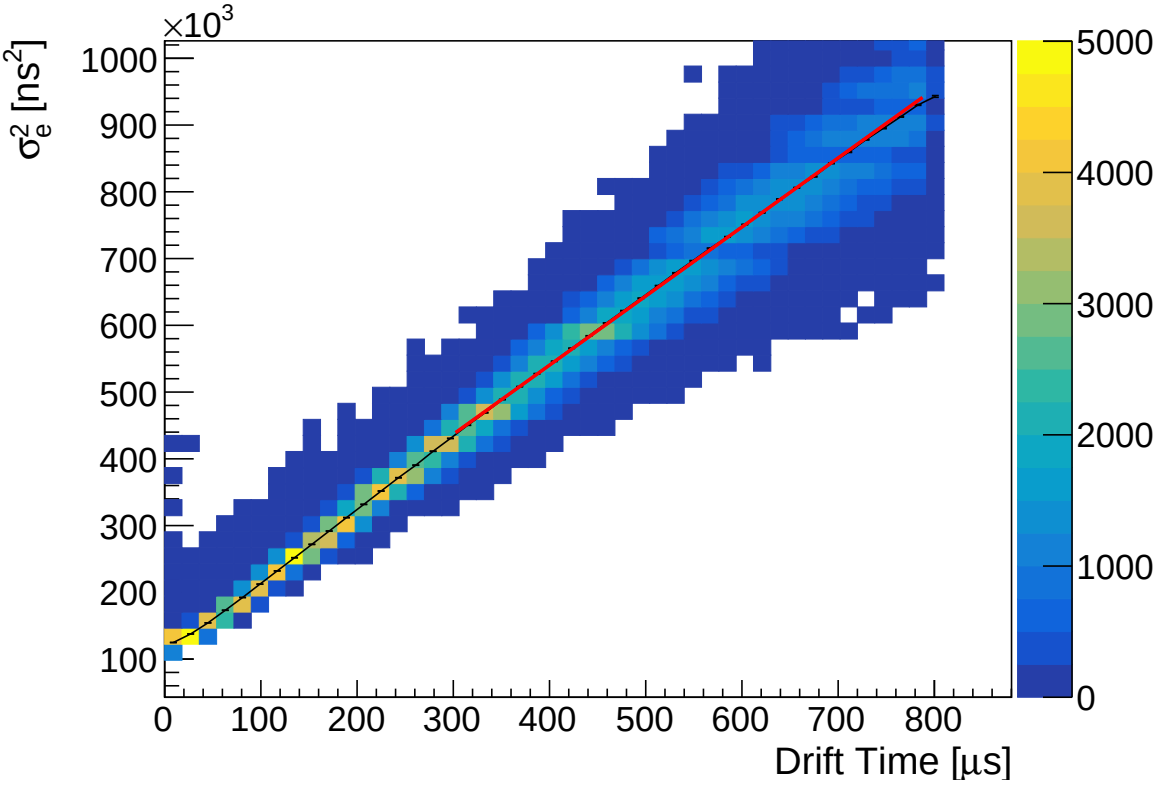

Figure 6.5: Variance of simulated ^{83m}Kr pulses, σ_e^2, as a function of the drift time. Black points show the mean and error from Gaussian fits to each drift time bin. The line is a linear fit to determine the diffusion characteristics in a region where the diffusion behaviour is dominant.

reconstructed. In the simulation, σ_0 is not a defined parameter and depends on many factors. Although, a true comparison is difficult, this value is consistent with a boxcar width of about 1.2 µs. This includes expected variance with pulse finding propagating through to variation within the pulse parametrisation and effects related to the drifting of electrons discussed in the next section.

6.2.3 Subtleties regarding depth and drift time

A subtlety arises when examining the vertical displacement from the extraction region within the TPC volume. Commonly this is referred to as the depth of an event. As part of the TPC operation, drift time is used as a proxy for this depth. However, in regions of the detector where the electric field has a non-zero radial component, the drift time can be longer than is implied from a certain depth. This is particularly more likely towards

the edge of the detector where fringe effects can occur, despite the field cage mitigations. Additionally, disparities can arise when reconstructing the pulses which are used to define the drift time. Fluctuations can be seen when examining a plot of the true depth of mostly ^{83m}Kr events against drift time as shown in Figure 6.6. In regard to the corrections of the S2 signal, the drift time is the important quantity, whereas with fiduciualisation it is the true depth.

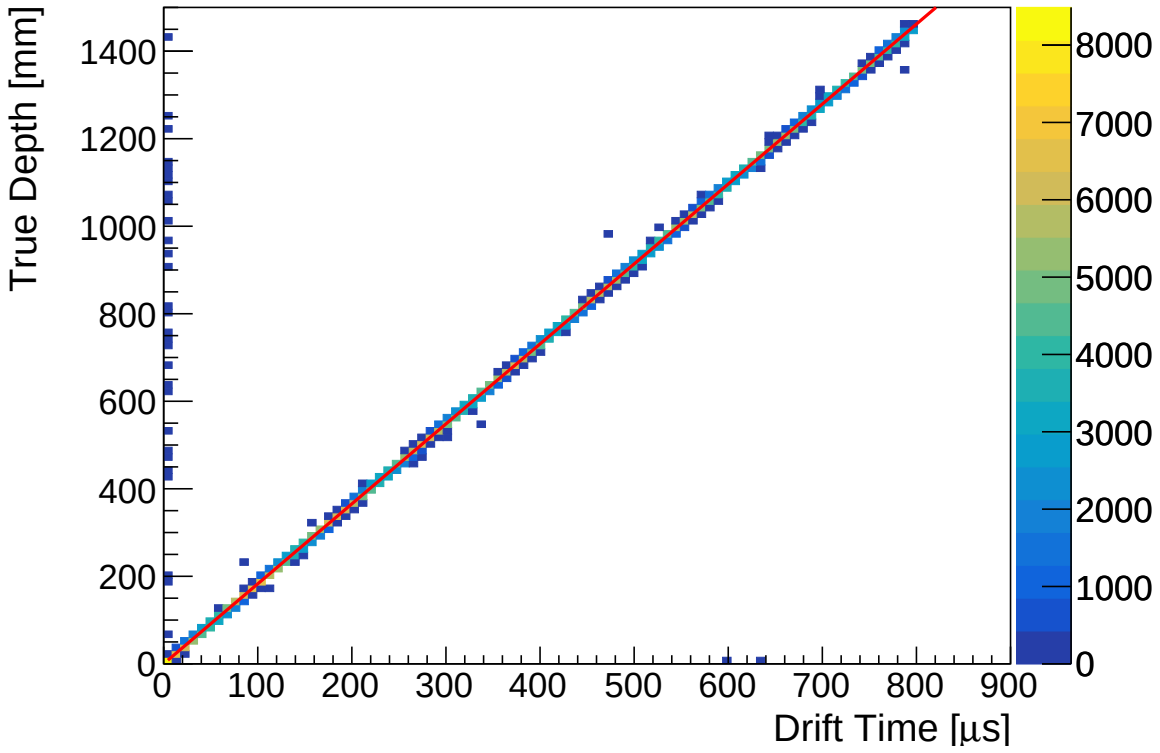

Figure 6.6: Using Monte-Carlo truth for mostly Kr events to show the mapping between drift time and true depth arising from non-uniform electric fields and reconstruction effects. Here events that have low drift times but high true depth are reconstruction failures.

The gradient of this line defines the mean drift speed. Using this simulated data this is determined to be 1.83 mm/µs using the nominal electric field configurations predicted in LZ software. This means that the drift time is a good proxy for the depth of an event in LZ.

6.2.4 Improving understanding of small signals in LXe TPCs to combat background rejection

Despite being able to show the relationship holds for higher energy calibration sources, the main benefit of using the S2 pulse width as a proxy for the interaction depth comes for events that are close to the nominal analysis threshold, with topologies lacking a measurable S1. Hence, this is a technique which can be examined at higher energies and then applied to an S2-only search. Other than using a similar technique, it is almost impossible to obtain depth information about S2-only events. Using the pulse width should enable the background from the grids to be at least partially separated from bulk events such as neutrinos and potential Dark Matter signals. Within this section a model is developed to examine small S2 signals based on the S2 pulse width. Using simulations, this model is used to examine the feasibility of improving background rejection with pulse width. Due to the complexities of many phenomena affecting the pulse width, effects are added gradually. Firstly a toy model is used, based on the first to last photon time with ideal pulses. Due to the limitations of this method, events are then simulated using the full LZ software and analysis framework and pulse widths are calculated based on the produced waveforms. At this point the study is considered with the addition of many phenomena that complicate this basic picture.

Toy model

In order to get a basic understanding of the measurement of the pulse width and performing an analysis, a toy model is first used. Within LZ, an introductory sensitivity study has been performed using such a model (Reference [141]). In order to build on this foundation, the study is repeated making the same assumptions to act as validation. The procedure followed is illustrated diagrammatically in Figure 6.7.

To simulate the physics leading to the pulse width, electron extraction times from a diffusion Gaussian are simulated as per Equation 6.5. A diffusion constant of 25 cm^2/s is used. In addition, photon PMT arrival times for each electron are approximated by

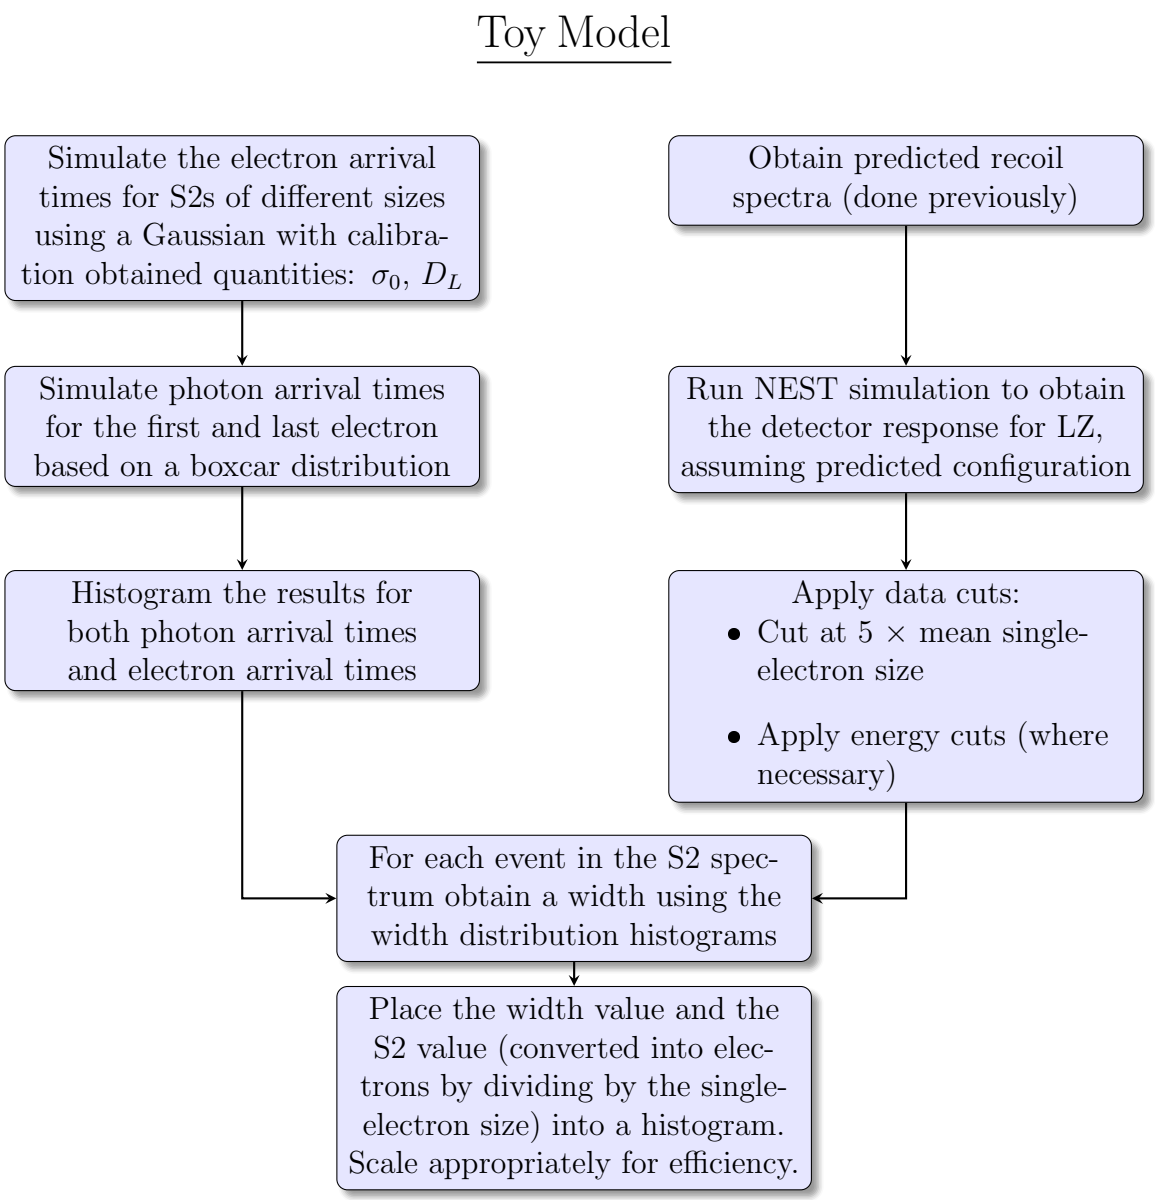

Figure 6.7: Flowchart explaining the steps in generating the toy model for the pulse width.

sampling from a boxcar distribution. This boxcar has a standard deviation of σ_0. From these distributions, the width is calculated as the time separation between the first and last photon from the event. Hence, this toy model contains the diffusion physics alongside a basic parametrisation of the extraction region effects as detailed in Section 6.2.1.

The simulated widths are then combined with the background and Dark Matter events, which are either simulated using NEST [99][145] or, in the case of grid backgrounds, using a scaled version of the LUX S2-only spectrum. In the case of the former the predicted energy recoil spectrum of events for each source is used. These are the same background

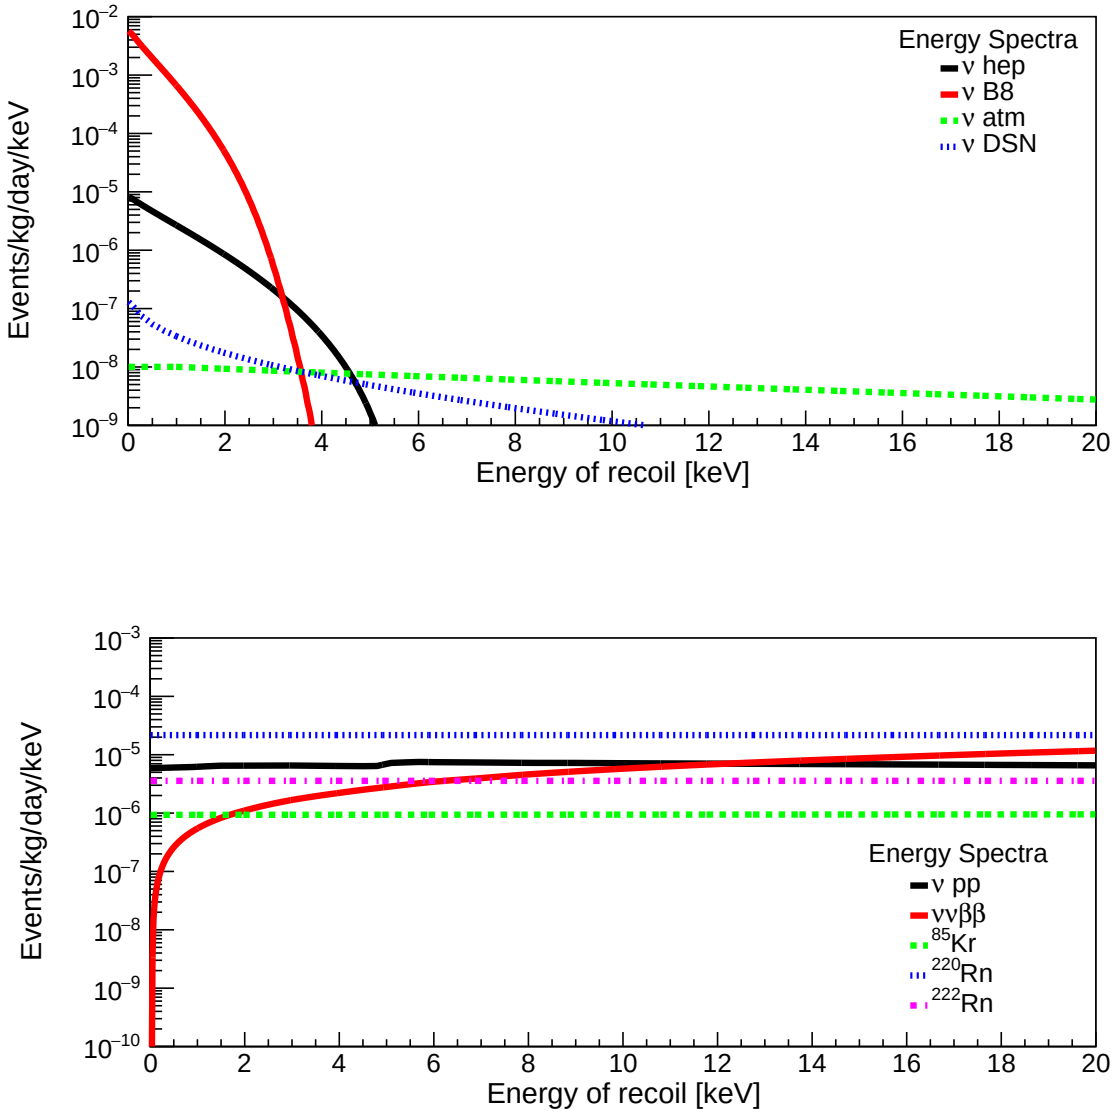

Figure 6.8: Energy spectra of dispersed sources considered within this study before efficiency cuts are applied. Hence, some events will either be out of the range of the nominal S2-only search or not produce detectable quanta. The top (bottom) plot shows the nuclear (electronic) recoil backgrounds. These sources are explored in more detail in Section 3.4.

sources as detailed in Chapter 3.4. Normalised to the expected rates (as shown in Figure 6.8), the rate is obtained from analytic calculations and estimates of activities within the experiment. The NEST simulation then converts this energy spectrum into quanta by Monte Carlo simulations which fold in the detector response and efficiencies. Before a more detailed scheme is developed with data, the background from the materials that make up the grids is difficult to estimate and previous experiments have seen very varied rates [146][147]. This background includes radio-isotopes which are embedded in or on the surface of grid wires, liquid surface events, and accidental coincidences of single-electron pulses. To obtain an estimate, the total LUX S2-only spectrum is hence scaled to a measured gate activity of 0.5 mHz, and the cathode similarly whilst building in the predicted detector effects of LZ such as electron attenuation from xenon impurities. Throughout an electron lifetime of 850 µs is assumed. To perform this scaling, detector effects, parametrised by g_2, are unfolded from LUX signals and then re-applied using the expected LZ value. This is done by multiplying by the ratio of g_2 values for each experiment. A 50:50 ratio is assumed for these backgrounds at the gate and cathode heights as was used in the previous study.

Minimal cuts need to be applied within the study. Only pulses above 5 electrons in size are considered (420 phd), to ensure that the radial position resolution is better than 1 cm. This is discussed in more detail within Section 6.2.5. Additionally, events below the calibrated energy range of the charge yield models are removed. In the case of electronic recoils only events with a recoil energy above 0.186 $\mathrm{keV_{ER}}$ are considered. This is also discussed later within this Chapter when the model is updated to include such events. Once those passing cuts are placed into a histogram, the result is scaled to ensure the rate is correct based on the efficiency and the initial rate of events for grids or dispersed sources.

The distributions of the reconstructed number extracted electrons and the resulting widths are shown in Figure 6.9. Here electronic recoil (ER) sources such as those of ^{220}Rn, ^{222}Rn decay chains and ^{85}Kr decays have been grouped together and to aid in

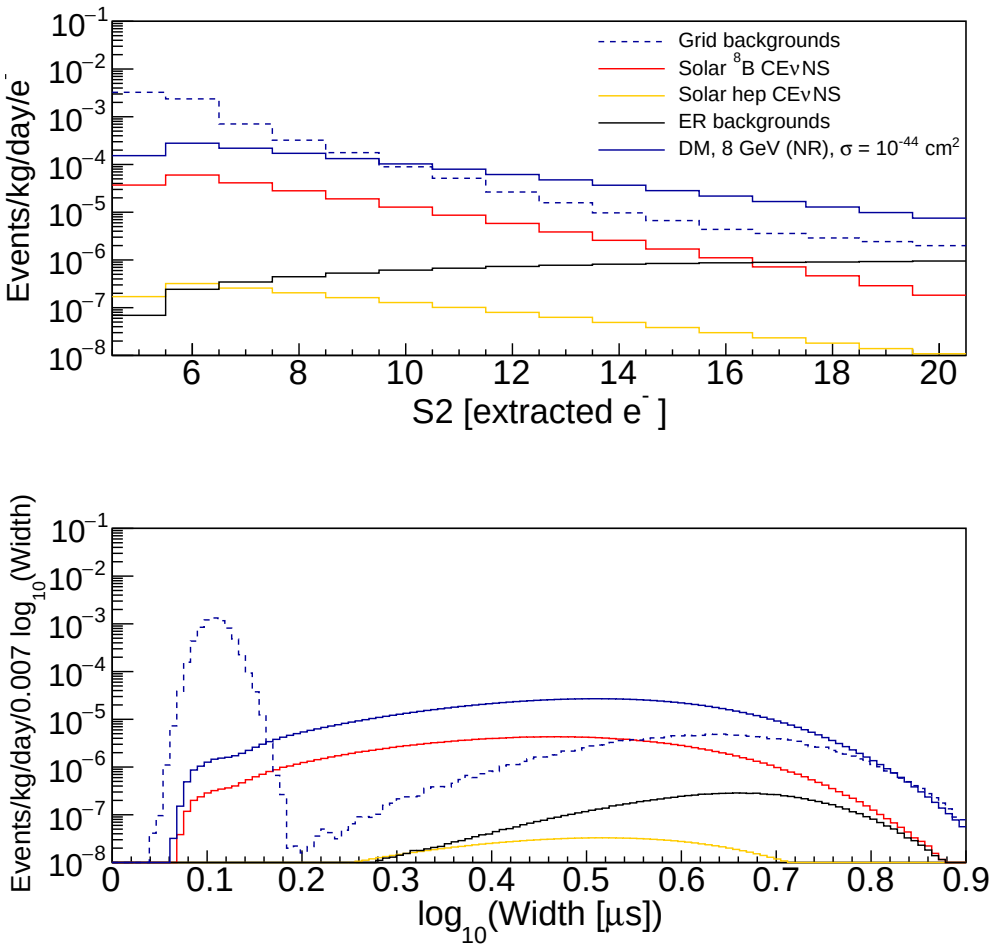

Figure 6.9: One-dimensional probability distribution functions (PDF) using the toy model for pulse width metrics for the principle background components. The grid backgrounds dominate over the other dispersed sources (top), whereas there are regions where the ^{8}B dominates over the other backgrounds in the width parameter (bottom). This figure is comparable to Figure 2 in Reference [141].

clarity of presentation. In order to reconstruct the number of electrons, comparisons are made between the mean and standard deviation of different numbers of electron values and the most compatible is chosen. This study shows good agreement with the study performed in Reference [141]. This is expected, given the same assumptions have been made and validates the use of this model as a comparison reference. Grid backgrounds are seen to dominate over those that are dispersed in terms of S2 size, however there are regions in the width distribution where this is not the case as pulse width is a powerful discriminatory parameter.

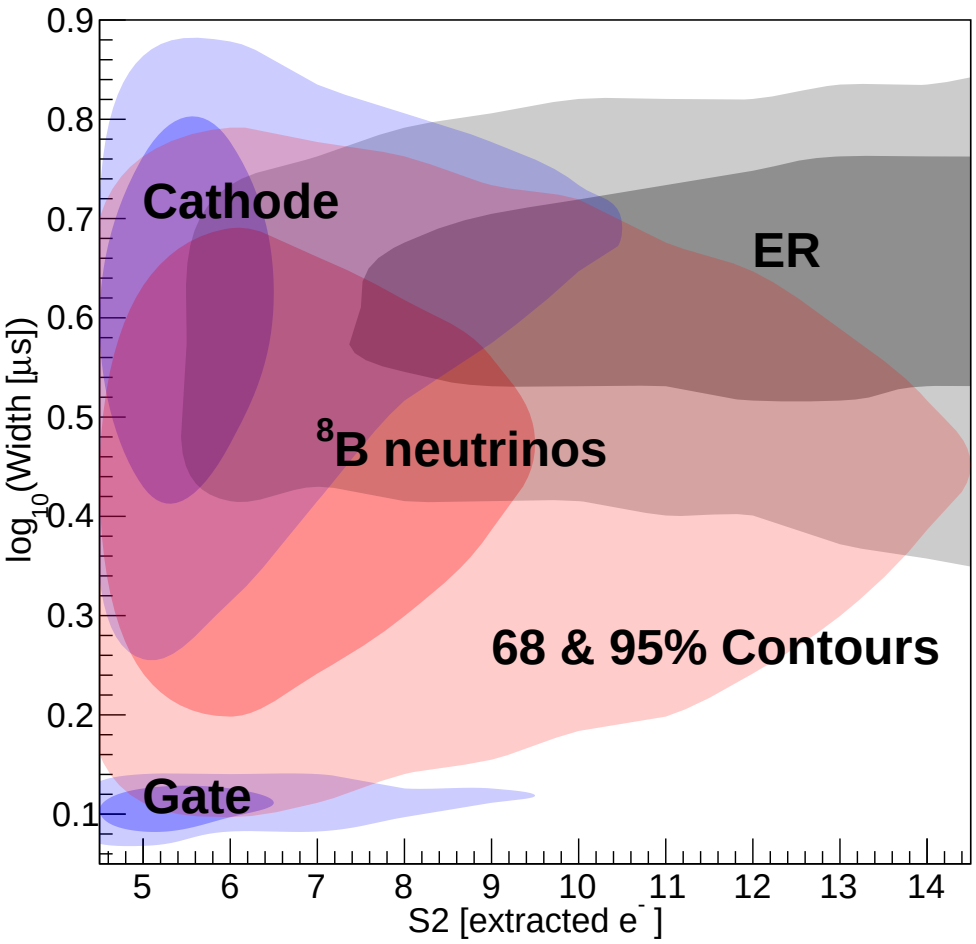

Figure 6.10: Two-dimensional probability distribution functions (PDFs) using the toy model for pulse width metrics for the principle background components. This shows some separation of cathode and gate events from those in the bulk. As there are intermediate values in the width, where the grid background is sub dominant, the sensitivity to Dark Matter is improved. Plot comparable to Figure 6 from Reference [141].

The true power of this study is when these two parameters are considered together in a two-dimensional probability distribution, as shown in Figure 6.10. This is due to the presence of some regions not dominated by grid backgrounds. The gate population is constrained tightly in width and size, whereas the cathode has a much broader width distribution. This shows that the position reconstruction ability in depth is particularly poor from width alone, and gets worse from events lower in the detector. This is especially the case for very small signals where width variance is greater. Despite both ^{8}B and ER sources being homogeneous in origin, once S2s are constrained to the range of 5-20

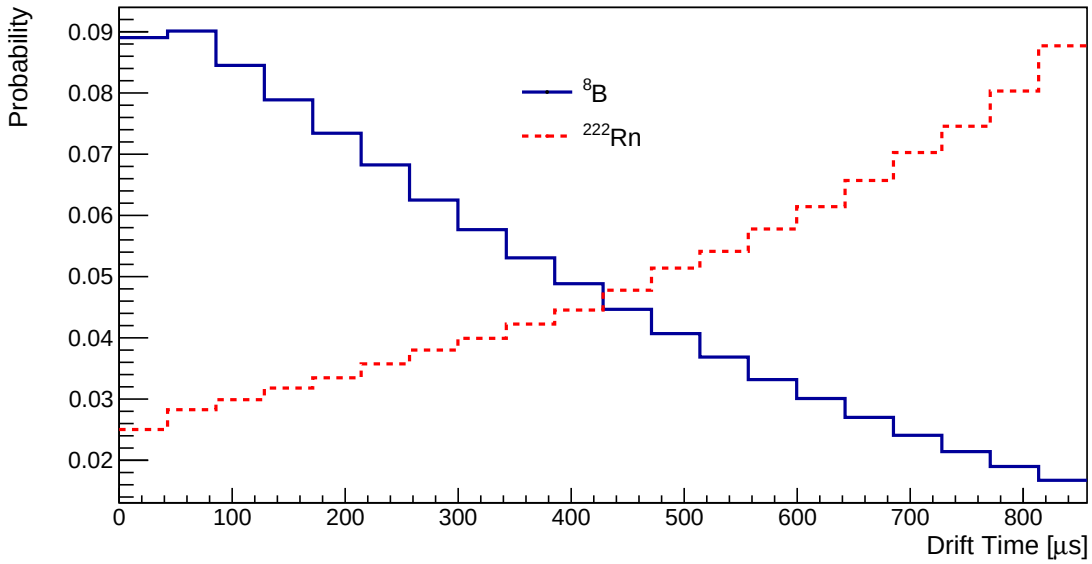

Figure 6.11: Position of ^{222}Rn and ^{8}B events that contribute to extracted electrons within the region of 5-20 electrons. Here, despite the homogeneous interaction spread, applying the cuts produces a spectrum which illustrates that most ^{8}B background events to the S2-only WIMP search occur in the top half of the TPC. This explains why the mean of the ER and ^{8}B populations in width space is different, and aids in the discrimination behaviour of this study.

electrons extracted, the distributions differ (Figure 6.11). Hence, the mean value of the width in each case is different. Due to the purity attenuation effects, the ^{8}B spectra falling exponentially with energy means that most events are of lower S2 sizes. Hence, some of the S2 signals are pushed below the 420 phd cut. Many of these events occur at the bottom of the TPC and have so have smaller S2 signals after electrons drift and are lost on impurities. A similar phenomenon occurs for the ^{220}Rn source but given the flat energy spectrum, higher energy events are observed in the 5-20 extracted electron range. This toy model provides a good comparison for more realistic cases, considered later in this chapter.

However, the toy model considered here has limited scope. The time between the first and last photon time with small S2 pulses is not the best choice of metric to use. This is because it is difficult to determine which photons are from the event and which ones are not. Additional photons, which can be delayed within the S2 signal due to light reflections, in combination with dark counts from the photomultiplier tubes would make this metric

too unreliable. There are also further physical phenomena present in the data which can lead to a poor estimation of the pulse width. Firstly, a small percentage (1-5%) of electrons are delayed by anything from several microseconds to seconds during the drifting and extraction process. Currently this is believed to be a mixture of electrons becoming stuck at the liquid surface and then extracted later, photoionisation, and impurity trapping [134]. This phenomenon creates something known as an electron train. Secondly, pulses can also be visibly split, both due to individual electrons being distinct quanta and from reconstruction. Thirdly the picture is complicated by PMT after-pulsing. This is a phenomenon observed whereby small vacuum impurities of gas ions contained within the PMT can travel to the cathode and produce a delayed pulse. The variation of light collected in the XY plane is another factor, and is discussed in Section 5.4. Additionally, merged pulses, further light collection effects, and reconstruction artefacts all can also influence the measurement of pulse width. Examples of some of this behaviour are shown in Figure 6.12 from bunches of about 12 extracted electrons. Crucially this does still show that there is some discriminatory power between events originating from the top and bottom of the TPC despite these extra effects. Compared with the main S1-S2 analysis many of these effects are of greater potential impact due to the small nature of the S2 signals. Hence, for future considerations the toy model should be discarded and a more accurate representation of real data considered.

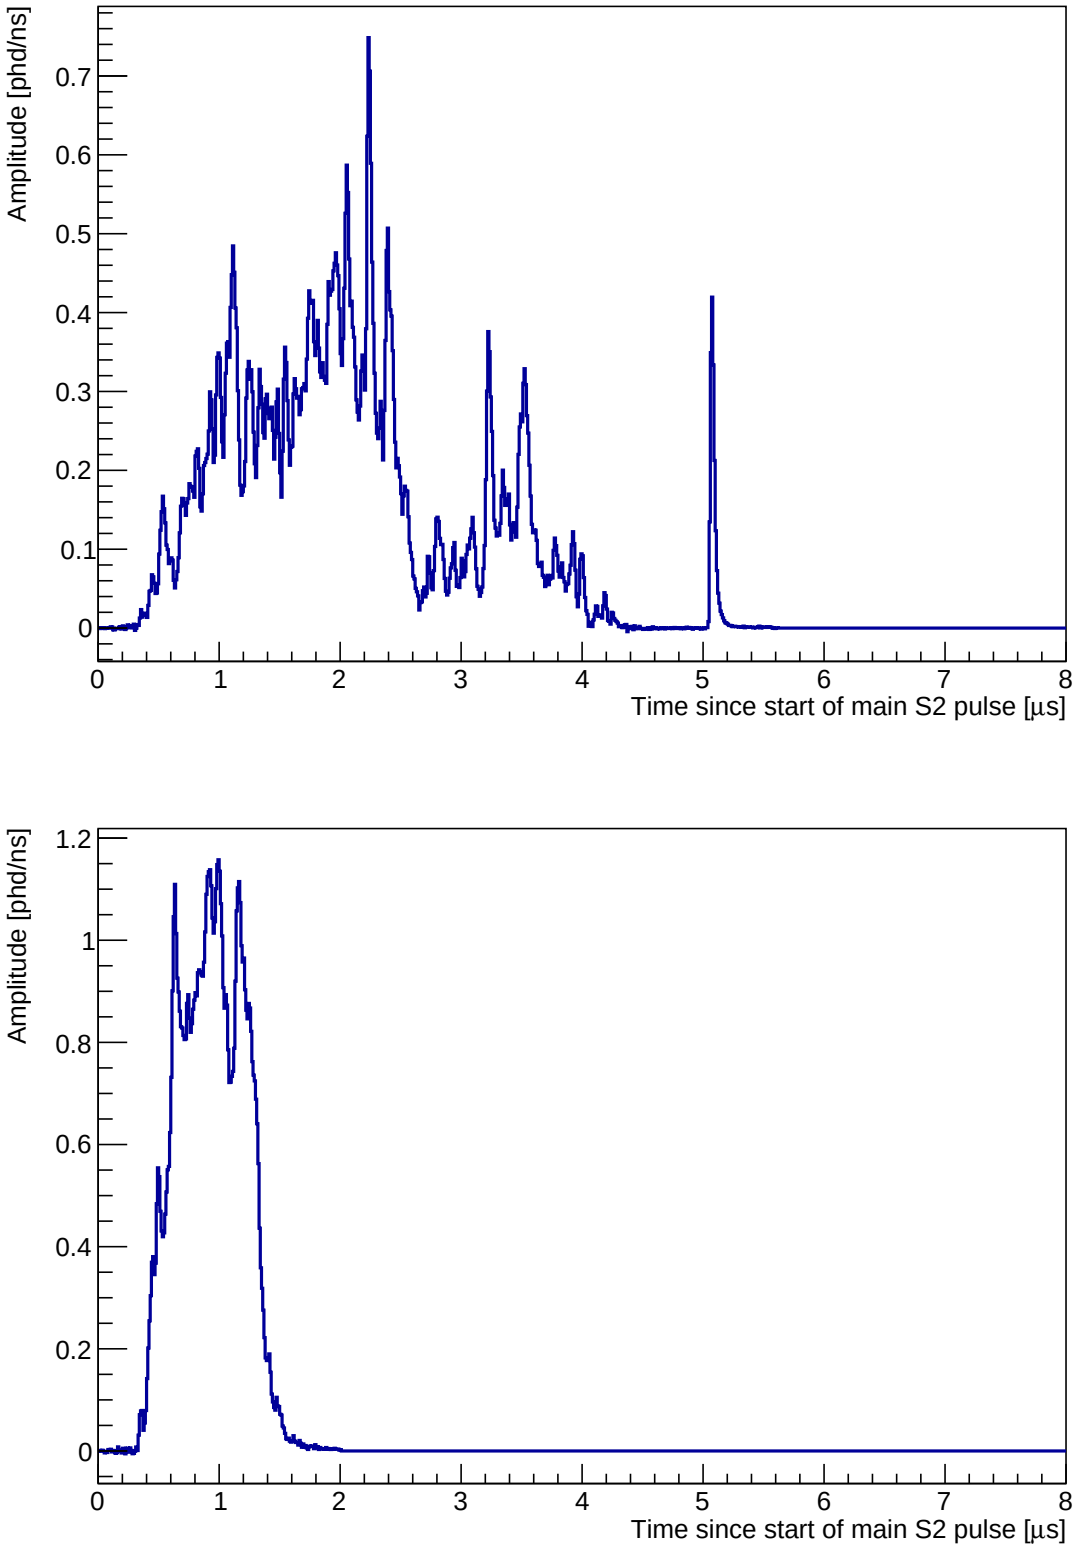

Figure 6.12: Some waveforms of S2 signals from about 15 extracted electrons. The simulated events occurred at the cathode depth (top) and at the gate depth (bottom). It shows that the pulses can be considerably non-Gaussian and also can contain additional effects such as in the first event where PMT after-pulsing (the final peak) and have split pulses. However, crucially this shows there is some variation of the pulse shape with the drift-time of the event.

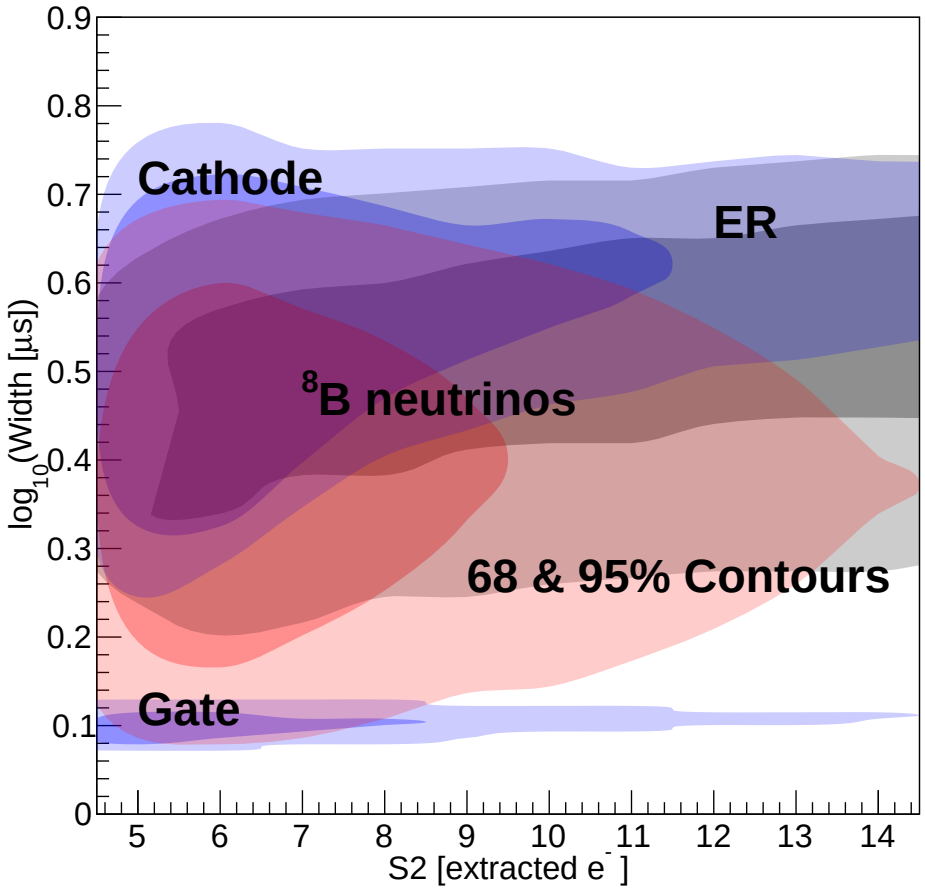

Figure 6.13: Two-dimensional probability distribution functions (PDFs) using the updated toy model for pulse width metrics for the principle background components. This includes changed assumptions as detailed in the text: removing energy cuts, changing drift speed and diffusion constant, and changed ^{85}Kr rate.

Before considering a more accurate model, in order to validate what changes originate from the additional phenomenon, the toy model needs to be modified to provide a more valid comparison. This is to ensure that the physics simulated by the toy model is similar to what is in the more accurate simulations. The updated probability distribution functions are shown in Figure 6.13. A comparison between the toy models is shown in Figure 6.14, which shows good agreement except where distributions are expected to vary. The modifications can be summarised into six main changes. Firstly, within the LZ simulations a different value for the diffusion constant, D_L has been used - 17.1 cm^2/s instead of 25 cm^2/s in the original toy model. The latest EXO data has yet to be implemented in the simulations and both values would be reasonable given measurement uncertainties [144]. A smaller value of D_L has the effect of shifting the entire population closer together in the width dimension. As this reduces the discrimination power between gate and cathode events, it should be considered the most conservative case of what can be achieved with LZ. The differences in the mean of each distribution are consistent with such an effect as can be seen when comparing the updated model in Figure 6.14 with the previous version in Figure 6.10. This should not affect any of the conclusions within this study. Secondly, the electronic recoil spectrum, consisting of the ^{220}Rn and ^{222}Rn decays and the following decay chains alongside ^{85}Kr decays and electronic neutrino scattering extends to lower energies. An energy cut at 0.3 keV_{NR} and 0.186 keV_{ER} has been applied above to achieve comparability with Reference [141]. This is included due to it being the lowest point calibrated with real data. NEST simulations confirm that the nuclear recoil energy cut does not affect the S2 range considered here whereas the electronic recoil cut does. Therefore, by applying the cut the electronic recoil backgrounds are artificially reduced at the lower S2 sizes. Hence, applying this cut leads to a more optimistic scenario for a sensitivity study of nuclear recoil events. For studies here, despite the larger charge yield uncertainty, the predicted response has been used below these energy values from the NEST models. Thirdly, the ^{85}Kr background is also now expected to be a factor of 20 times larger than the original rate shown in Figure 6.8 (page 153). This is due to the requirements changing within the krypton removal strategy. Fourthly, a small

correction to the grids spectra has meant that the rate does not fall as quickly as was previously expected with increasing S2 size. Fifthly, the drift speed is updated from 1.7 mm/µs to 1.83 mm/µs which also impacts the width distribution. Finally, the version of NEST is updated. Within these updates the nuclear recoil fluctuation probability is reduced to provide better agreement with experimental data. In doing so this also reduces the fluctuations in the ^{8}B spectrum in Figure 6.14 and, due to the exponentially falling distribution reduces populations in the higher S2 sizes.

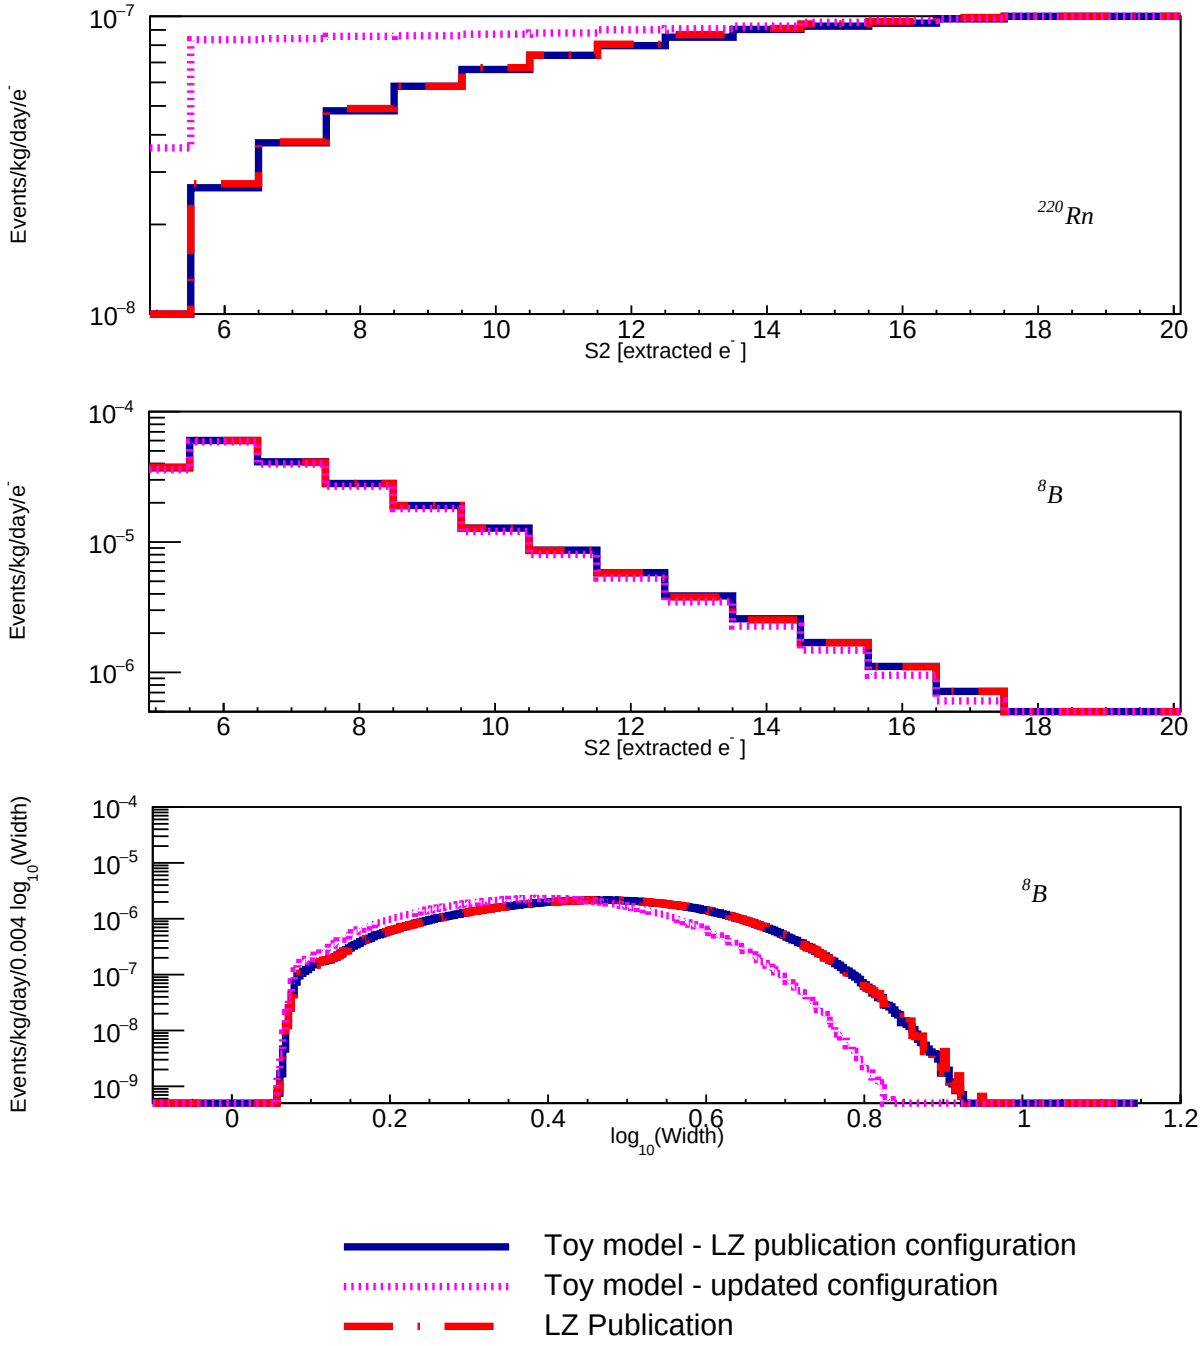

Figure 6.14: Plots illustrating the effects of modifying the assumptions made within the toy model for ^{8}B and ^{220}Rn. There is good agreement with the LZ publication with the toy model study, using the same assumptions. The main text discusses the changes made to update the toy model shown here.

Developing a more realistic model

In order to perform the generation of PDFs with more realistic physical features, the method is modified from the toy model study to one using more realistic simulations. In making this change the method by which the PDFs are created also must change. Instead of reconstructing the number of electrons within the S2 signal which are extracted into the gas region (as is done above with the toy model), the number of electrons produced in the liquid is instead used. For events occurring in the bulk of the liquid xenon, the same NEST simulations as before are used to obtain this number. For grid events the LUX experiment spectrum, expected to be composed of mostly grid and surface events, is scaled similarly as in the toy model study by dividing the spectrum by the detector specific quantity g_2 from the LUX experiment (12.1 ± 0.8 phd/e$^-$ [96]) to obtain the number of electrons in the liquid. This assumes that any effects from the attenuation are negligible in LUX data, which is true when the electron lifetime is large compared with the maximum drift time. Similarly, as before, a 50:50 ratio has been assumed for cathode and gate events. This number of liquid electrons then interfaces with the more accurate simulations (Section 3.5.3) which handle many of the physical processes contributing to pulse width. These simulations incorporate the propagation of the electron cloud within the bulk, its extraction and subsequent light detection effects, and can include PMT after-pulsing and the presence of electron trains. Simulations are performed with millions of electron bunches in sizes ranging from 1 to 80 electrons in size across the bulk. Similarly, simulations are performed specifically at the height of the cathode and gate in order to obtain sufficient statistics. These events contribute to extracted electrons in the range 0 to 20 electrons after electron lifetime (taken to be nominally 850 µs) and extraction efficiency effects are incorporated (see Figure 6.15). Using these events, the number of extracted electrons and the pulse width is reconstructed (using the same method as in the toy model) and is used within the PDFs. As before, the resulting PDFs are scaled according to the initial rate of background events. For the study discussed here, only simulated data is available and hence the true absolute rate of events from the grids is

unknown. Ultimately this will be measured with data. Hence, these grid spectra are normalised to 0.5 mHz between 5 and 20 extracted electrons based on approximations from LUX [141] in the same way as the toy model. Using raw liquid electrons in this

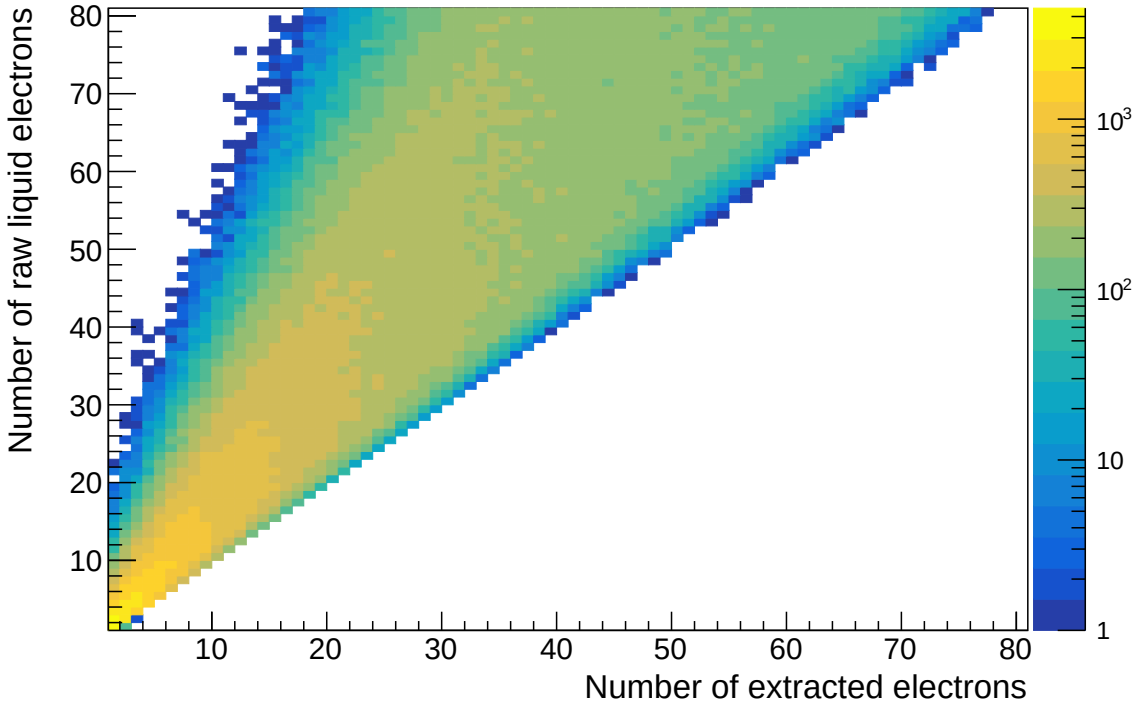

Figure 6.15: Extracted electrons from the number of raw liquid electrons. Various phenomena contribute to the loss of the electrons, including attenuation on impurities (electron lifetime effects where the electron lifetime is taken to be 850 µs) and non-unitary extraction efficiency. This however shows that a fraction S2-only signal is lost below the lower analysis threshold of 5 electrons.

way has the advantage of preventing a mixture of effects from the detector configuration differences between the full simulations and the parametrisation in the toy model. In all cases, a cut is placed on the number of raw extracted electrons at five times the average S2 response from a single electron. Again, this is due to the radial position reconstruction and is discussed in Section 6.2.5. The width distributions for each number of bulk liquid electrons are then weighted according to the number produced for each source. This process is illustrated in Figure 6.16.

Following the narrative of adding in physical features gradually, the simulation can be adjusted to add some effects which can broaden the pulse width further than diffusion.

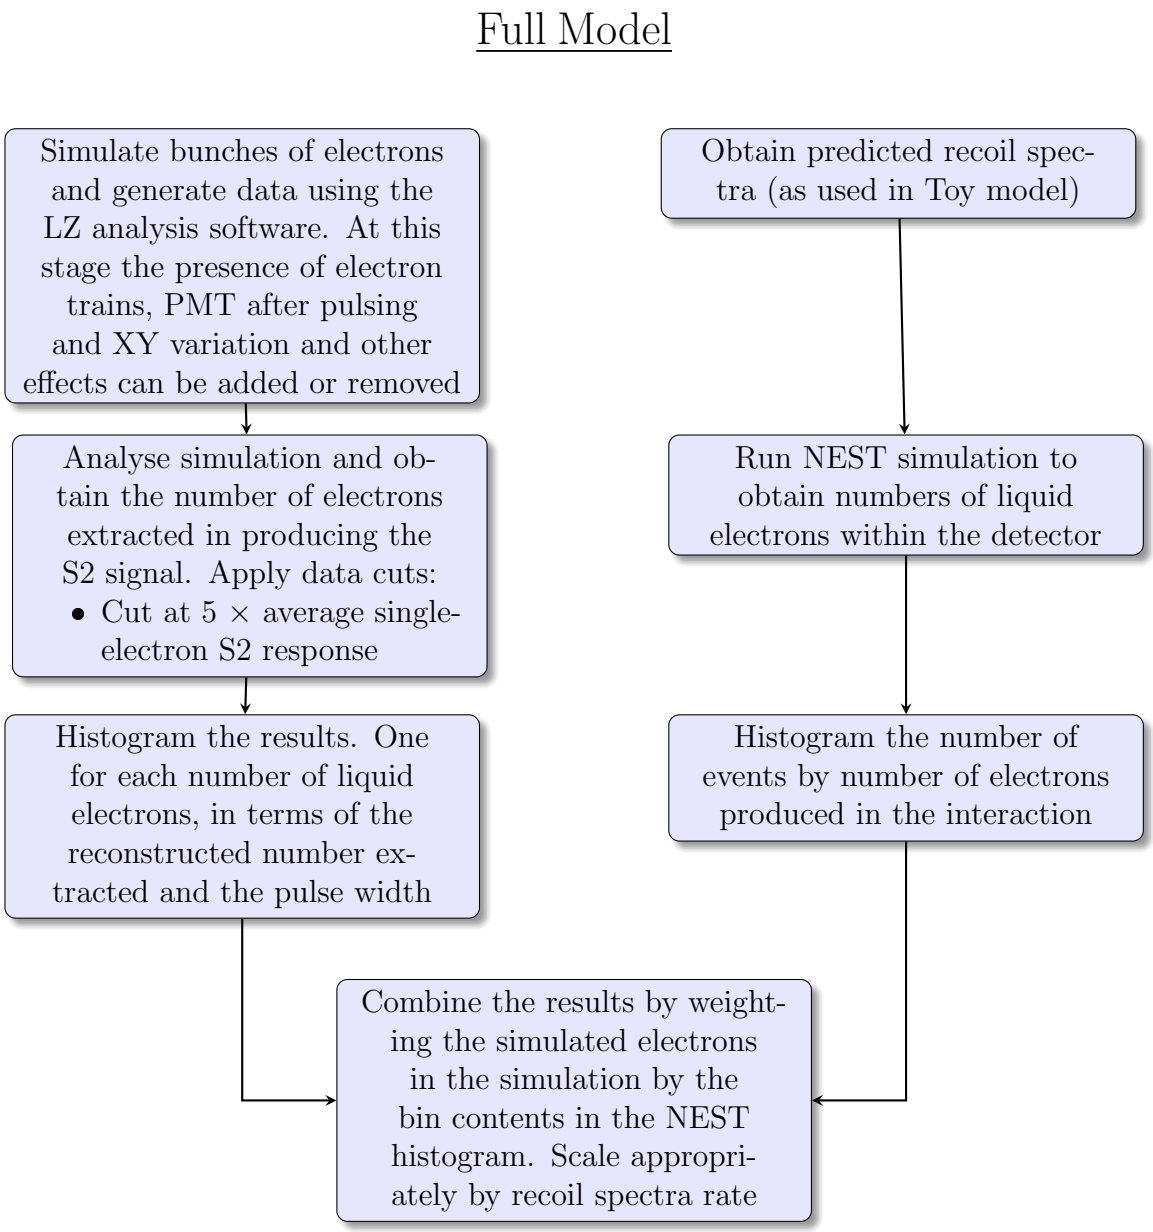

Figure 6.16: Flowchart explaining the steps in generating the full simulation model for the pulse width.

This scenario is one similar to the updated toy model study (Figure 6.13). With the resulting PDFs shown in Figure 6.17, all events are constrained to have a radius less than 5 cm from the centre of the TPC which removes the radial position dependence on the light collection efficiency (as discussed in Chapter 5.4). The delaying in time of single electrons in the electron train, and PMT after-pulsing have also not been simulated. Hence, in this situation, only the stochastic nature of photon arrival times and effects linked to pulse splitting and some reconstruction features remain. Even with many of the

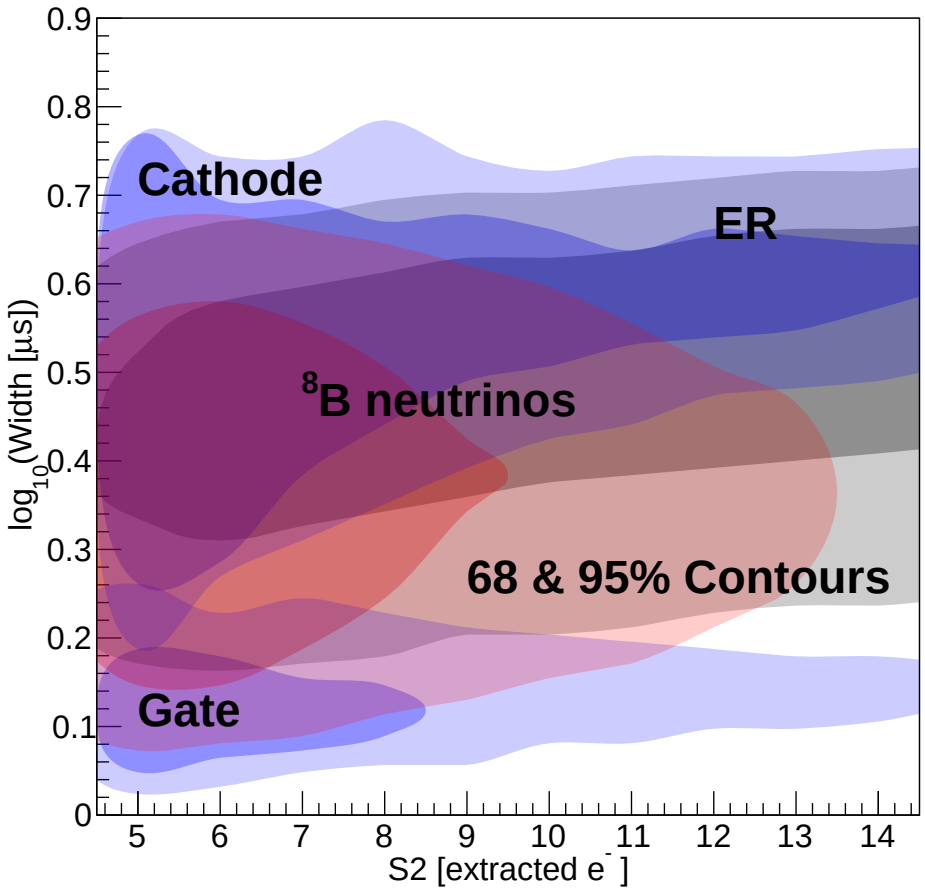

Figure 6.17: Two-dimensional probability distribution functions (PDFs) using the more accurate simulations. These simulations include photon propagation effects and reconstruction effects, but effects from electron trains and PMT after-pulsing are not included. The pulse width has been obtained using the Monte Carlo truth variables pertaining to the first to last photon time.

phenomena remaining switched off, Figure 6.17 shows there is greater variance in the width distribution of gate events using Monte Carlo truth, first to last photon times compared with the updated toy model study. This extra variance is due to reconstruction effects and the modelling of light reflections from the Teflon in the TPC, going beyond the boxcar considered before. Hence, it is evident that the modelling of the single electron waveform with a boxcar is not the best approximation. This is also seen when considering real pulses during the detector levelling with little diffusion. Apart from the gate population then there is good agreement between the more accurate simulations (Figure 6.17) and the updated toy model (Figure 6.13) despite changing from the toy model to the more

accurate simulations. This suggests that the diffusion modelling is good, however the modelling of the single electron S2 width is limited in the toy model case. .

Considering additional effects within the model and pulse width optimisation

From Figure 6.12, it is evident that the addition of PMT after-pulsing and electron trains, which occur at random, is going to increase the variance on measured pulse widths. To combat these effects, it is necessary to choose a definition of pulse width which minimises this variance. Indeed, this also must be a measurable quantity in the final experiment unlike the above methods which use Monte Carlo truth variables. Table 6.1 summarises the width metrics that are considered; the conversion factor for each metric to obtain σ_e (Equation 6.5, page 144) is included. To calculate this factor, metrics are calculated with a Gaussian function and compared with the standard deviation. Using the conversion to σ_e ensures that the metrics can be compared using a common standard.

Pulse Width Metric	**Explanation**	**Multiplicative factor to σ_e**
aft90-aft10	Difference in 90% and 10% area fraction times (aftxy is time at which xy% of the total integrated charge from electroluminescence photons in an S2)	2.56
aft75-aft25	Difference in 75% and 25% area fraction times	1.34
RMS	Root Mean Squared value of the pulse. Calculated using information for each sample	1.00
FWHM	Full Width at Half Maximum of the pulse	2.36
FWHaft50	Full Width at half the evaluated pulse value at 50% of the integrated charge of the pulse	2.36
Pulse End - Pulse Start	Reconstructed start and end of the pulse using the LZ analysis package (LZap)	4.00‡

Table 6.1: Pulse width metrics used for this study. To obtain the value of σ_e the value from each metric discussed needs to be divided by that factor. Conversion factors have been obtained by evaluating the metrics on an ideal Gaussian distribution.

‡ For the final metric this factor has been approximated assuming 95% of the detected photons in the pulse are within the reconstructed pulse boundaries.

To perform a study as to determine the best metric, S2 signals were simulated originating from the centre of the drift region of the TPC volume. To ensure that these results are

robust against extra physical phenomena, simulations include effects from after-pulsing and XY variation of signals. Electron-trains are not included within these simulations, hence this makes the assumption that electron trains can be eliminated, which is considered separately below. To ensure that events can be separated at different drift-times, each metric can be used to estimate the value of σ_e and the distributions compared. The best pulse width metric is less susceptible to the additional physical phenomena and hence will have a distribution which has the smaller standard deviation. Figure 6.18 shows an example set of distributions for a set of about 5 to 6 electrons. Here, the area fraction time metrics perform better alongside the root mean squared of the pulse. In addition to the FWHM and FWHaft50 having larger variance, there are systematic offsets from the other metrics, suggesting that these are poor estimators of σ_e for small S2 signals. This result is not surprising as the small electron size S2 pulses are not smooth Gaussian functions, with split pulses and stochastic variations in light arrival times, leading to large variations within the interpolation algorithms to determine the FWHM and FWHaft50 metrics for each pulse. Within the next section, the RMS, aft90-aft10 and aft75-aft25 metrics are considered more closely with the electron trains added to fully optimise this choice.

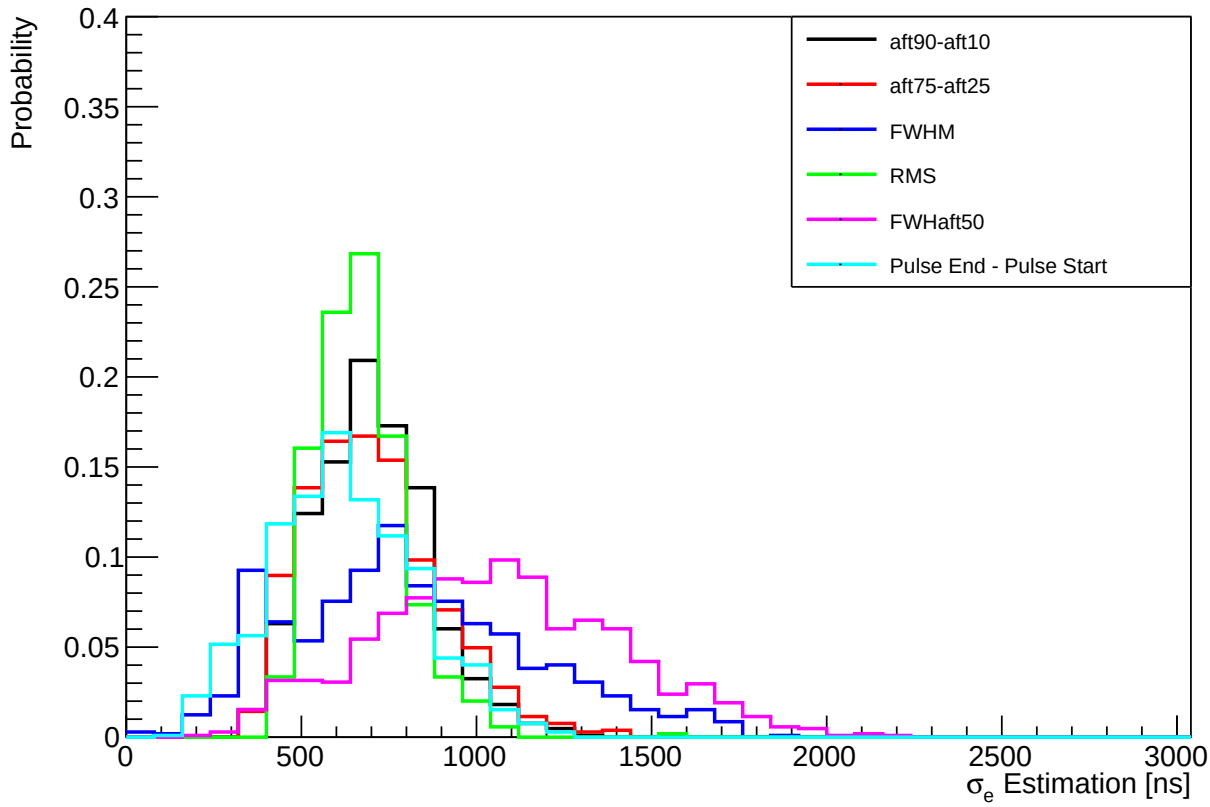

Figure 6.18: Width metrics from electron bunches simulated from the centre of the TPC in depth across the full radial range. This illustrates the distribution of pulse widths normalised to σ_e using values from Table 6.1. Metrics with a smaller variance are better as they lead to an improved estimate of σ_e and hence depth resolution.

Considering the impact of 'electron trains' on the model

Adding in the delayed electrons from the electron trains complicates the picture further. When defining the pulse width, it is necessary to exclude as many of these delayed electrons as possible, because they do not contain information encoded about depth by virtue of the temporal distribution not being due to diffusion and extraction region effects alone. Indeed, they can also mask many of the diffusion measurements for the 'main' part of the pulse. Predominantly, issues occur when electron trains are merged together with the largest S2 pulse and a very long reconstructed pulse, can be formed. Similarly, as Figure 6.19 illustrates, there can be a bias depending on the individual topology, where an electron could be included far from the main pulse if there are intervening electrons. Here, if the pulse boundaries are defined by the lines below the events, differing contributions

to the reconstructed pulse can be made. This then means that in the first two cases an electron following at a similar time can sometimes be included in some topologies. The first and last examples illustrate the masking of the main pulse width by topology effects of the delayed electron within the train. A reconstructed simulated event is illustrated in Figure 6.20, showing an event which would have this issue. Typically, this kind of reconstruction effect can occur for S2-only searches as the algorithms are typically tuned to much larger S2s that do not have these issues.

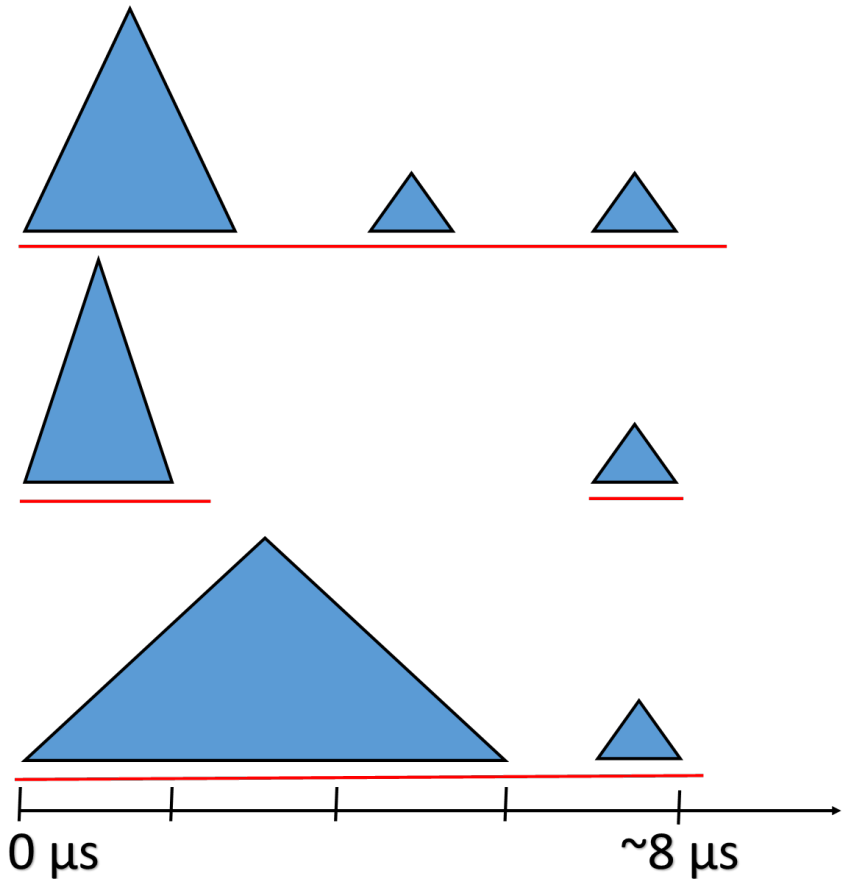

Figure 6.19: Diagram to explore three event topologies that can lead to the pulse classifier obtaining less useful results. The lines below the pulses signify a typical pulse identification, where often part of the electron train is included. Hence, this illustrates how the main pulse, represented by the first triangle in all three cases is adversely affected by this when estimating the pulse width. An example of such a waveform is illustrated in Figure 6.20.

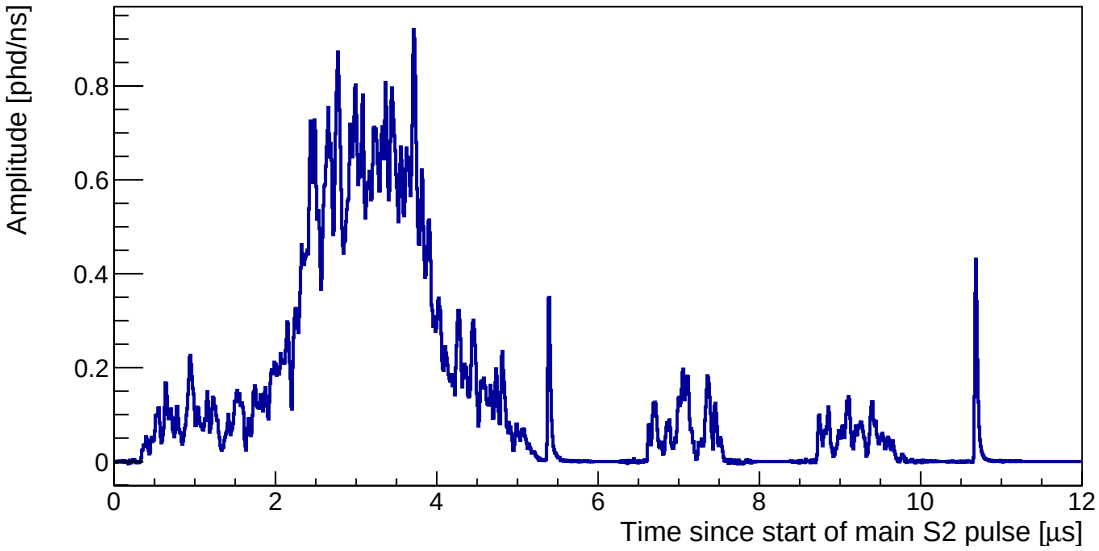

Figure 6.20: A simulated S2-only event with electron train. Everything after 6 µs can not be from diffusion effects alone and hence must be due to delayed transit. Such electrons should be ignored from width and S2 size calculations.

To help discriminate between electrons delayed in the train and those that are likely from the main pulse, a temporal window which is the length of an S2 pulse of an event originating at the cathode is defined. Anything after this time interval is not included in the pulse contribution. This cut is tuned such that it will not lead to the last single electron being bisected. It is done by checking the DAQ response has returned to near baseline. To obtain the maximum width of a pulse, a series of 15 electrons from the cathode can be approximated using a similar procedure to the toy model study. Instead of using a boxcar probability distribution function as before, a data-driven single-electron template can be formed. A Gaussian distribution for longitudinal diffusion is used to determine electron extraction times with corresponding photon arrival times sampled from this average single electron waveform. Therefore, a measurement of the first to last photon of the pulse can be estimated. Here, the individual photon width is negligible compared with the S2 width and so is neglected. Figure 6.21 suggests a 6 µs cut is appropriate because it ensures that the probability of an electron being missed that belongs with the pulse is less than 0.07%, bearing in mind that contributions are included after the 6 µs window. This method is preferred over running more simulations because it is a technique that can be used with real data. In forming the single electron template a series of 10,000 waveforms from

single electron pulses in the simulation are averaged (see Figure 6.22). This is done by temporally shifting the pulses so that aft50 on each pulse is aligned. The function itself is, by observation, not an ideal boxcar as is used in the toy model study. This is likely due to the photon propagation models and the variation in aligning pulses properly in the analysis with the aft50 metric. The presence of the tail, indicative of light delayed in its detection, is a key driver for the multiple electron pulses from Gate events being longer in Figure 6.17. It is important to note that this waveform does not have all the realistic effects included. There is likely to be amplification of the light yields as electrons approach the anode wires during the extraction, due to the potential changing from an approximately uniform field to the radial one of the wire. This is observed in waveforms when levelling the detector (Section 5.3.2).

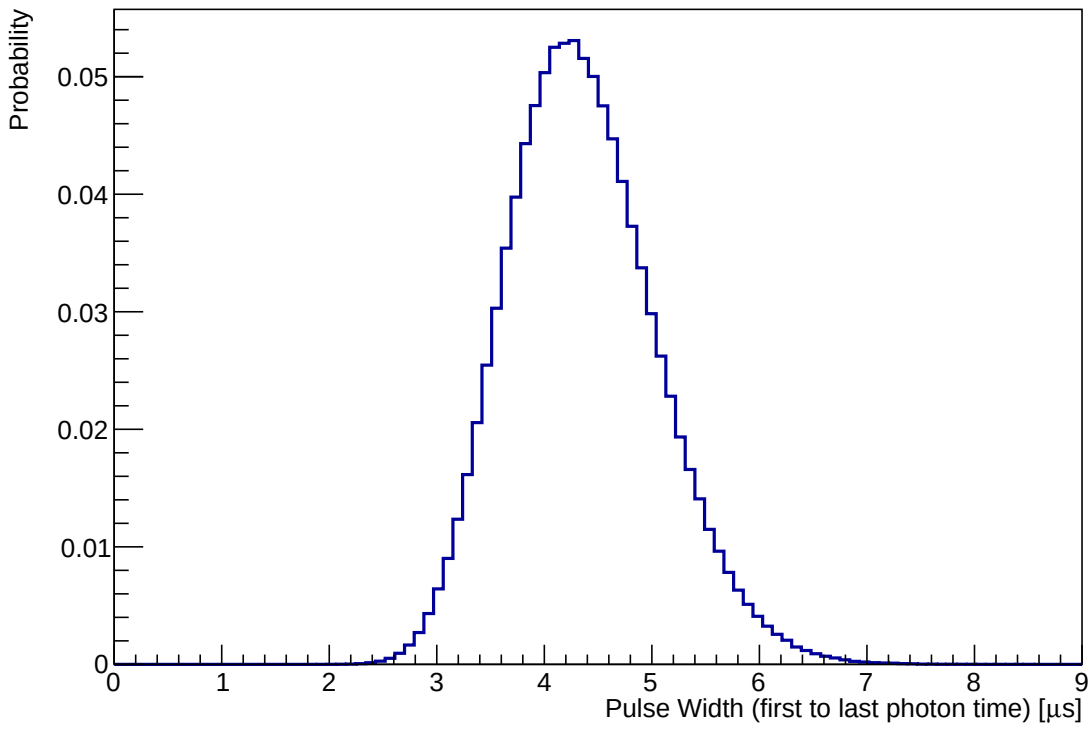

Figure 6.21: First to last photon time from events, with the diffusion constant, D_L = 17.1 cm^2/s, drift speed, v_d = 1.8 mm/µs, the maximum drift time as 850 µs, and the photon arrival times simulated using the single electron template function in Figure 6.22. This illustrates that applying a cut at 6us, whilst including single electron contributions that have begun is appropriate.

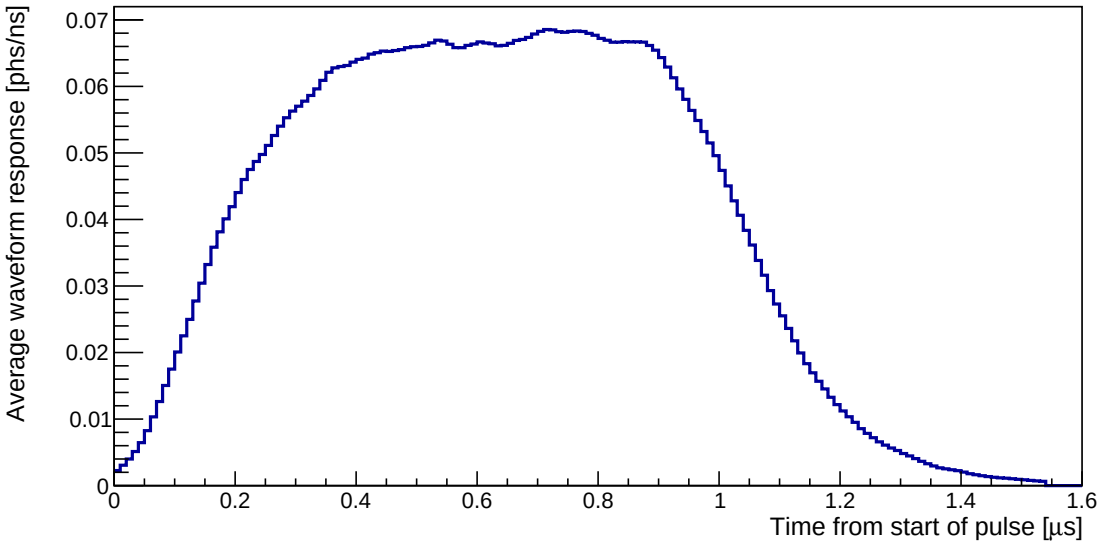

Figure 6.22: The average waveform for a single electron extracted from simulation. Single electron waveforms are aligned by the aft50 of each pulse to allow for temporal variation in the pulse classification boundaries.

Based on this variability of electron trains being included within the main pulse as defined by LZ software, using the standard reconstruction algorithms should be considered unreliable for an S2-only search. Hence, using the waveform directly and only including contributions within a temporal window from the start of the pulse is the preferred option. The pulse width metric discussed earlier needs to be re-evaluated based on this time window. Specifically, as events from the gate and cathode lead to the most extreme pulse widths, they are considered. Figure 6.23 shows that the metric defined by the aft75-aft25 is the best metric to use. It is well constrained for both the gate and cathode events, and it does not contain a long tail, extending upwards in the parameter values at the gate. Thus, meaning that the metric is less affected by electron trains that are invariably included. The RMS is obviously poor at discriminating events at the gate because it is highly influenced by the electron trains, and the aft90-aft10 contains an extended tail which precludes use of this in discriminating with bulk events. The systematic offset between the aft75-aft25 and aft90-aft10 metrics is due to the point where the metric looses sensitivity around the single electron width. When events originate from nearer the top of the TPC the aft75-aft25 metric reaches a limit before the aft90-aft10 metric. The aft75-aft25 metric constructed by the largest defined pulse in the reconstruction

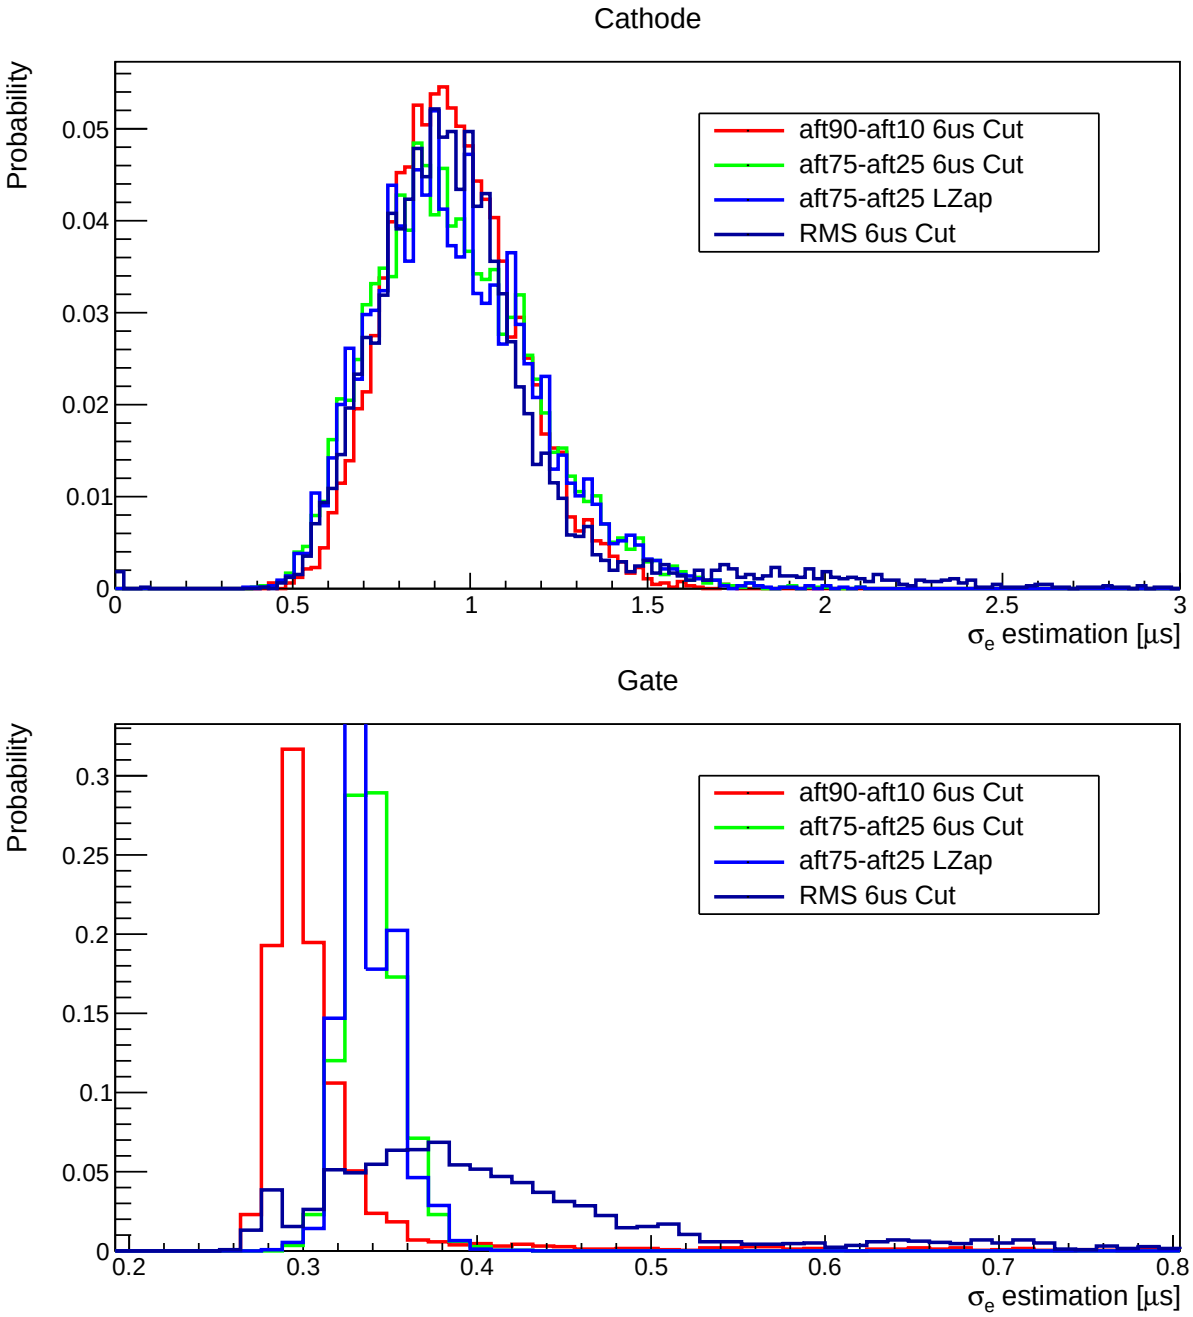

Figure 6.23: Plots illustrating the choice of width metric on σ_e for a series of the cathode and gate events of extracted size 800 to 1000 phd. The cathode events all have a very similar width distribution, hence the choice of metric is somewhat arbitrary, but the Gate events show systematic offsets and tails extending to much larger widths, making the aft90-aft10 and the RMS bad width metric choices. The reconstructed width distributions (LZap) are included for completeness.

software is included to validate that the waveform distribution is similar.

There is a small chance that electrons can randomly occur within the window and not be correlated with the diffusion properties of the main pulse. However, this is likely to be

small. From toy Monte-Carlo simulations it is estimated to occur for less than one event, assuming a 100 Hz rate of single electrons during the science runs from just random single electrons. Indeed, as the optimal pulse width metric for this analysis is less susceptible to these effects anyway, then even if a true event occurs within the same 6 µs window as random single electrons, then the width metric is somewhat robust at dealing with this, as is illustrated with gate events. The time window duration obtained here is dependent on the diffusion constant, the drift speed and the maximum drift time. These factors all have some dependence on the drift field in LZ and hence, this 6 µs cut will need to be tuned when the detector has sufficient data.

Some consideration of the impact of applying a 6 µs cut upon the size of the S2 signal itself is also required. There are situations where more than the 5 electrons, necessary to meet the lower analysis threshold cut, could be extracted. However, due to the addition of the 6 µs cut, some events could have a number of electrons removed and hence be located below this value. Figure 6.24 shows that the numbers extracted are often the same, however unfortunately there are cases such as for 10 electrons, where 50% of the signal is lost. This, fortunately is not the case for the majority of cases as the linear regression illustrates. However, in attempting to deal with edge cases, if the window is extended, then this has two fold impact. It increases the chance that single electrons from the electron train lagging the main pulse are included within the pulse boundary, and hence potentially bias the width. Additionally, it increases the chances of pileup with single electrons and other small S2 signals and hence lead to a poor estimation of extracted electrons.

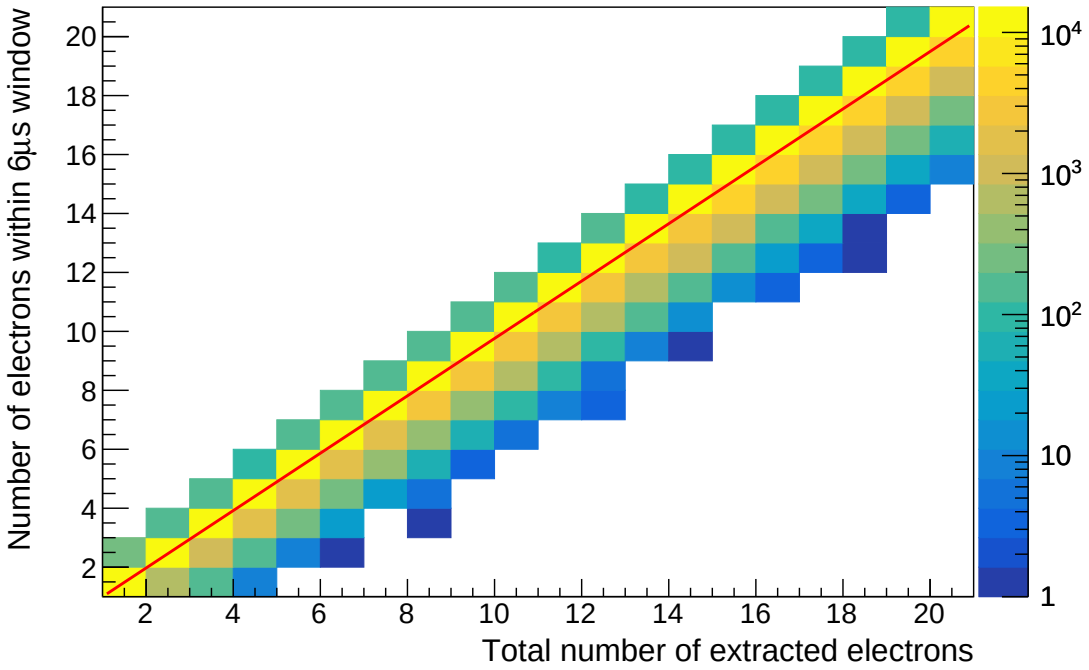

Figure 6.24: Plot to show how the S2 signal size is artificially reduced as a result of the 6 µs cut on waveforms. The trend line has gradient of 0.96 illustrating this is a measurable but small effect.

Based on the choice of the optimum width metric (aft75-aft25), Figures 6.25 and 6.26 show the PDFs in one and two dimensions. Comparing the Figures 6.26 and 6.17 it is possible to see that there is additional variance in the pulse width. This is despite the optimisation of the pulse width metric and is likely due to the additional phenomena discussed in this chapter. Optimistically though, there are regions remaining that are more separated from the grid backgrounds and hence this should show some improvement from the physics reach of a nominal S2-only search, where only S2 size would be considered. This can especially be the case if it is used in combination with a grid backgrounds reduction technique such as that used in LUX [146] to remove some of this background further.

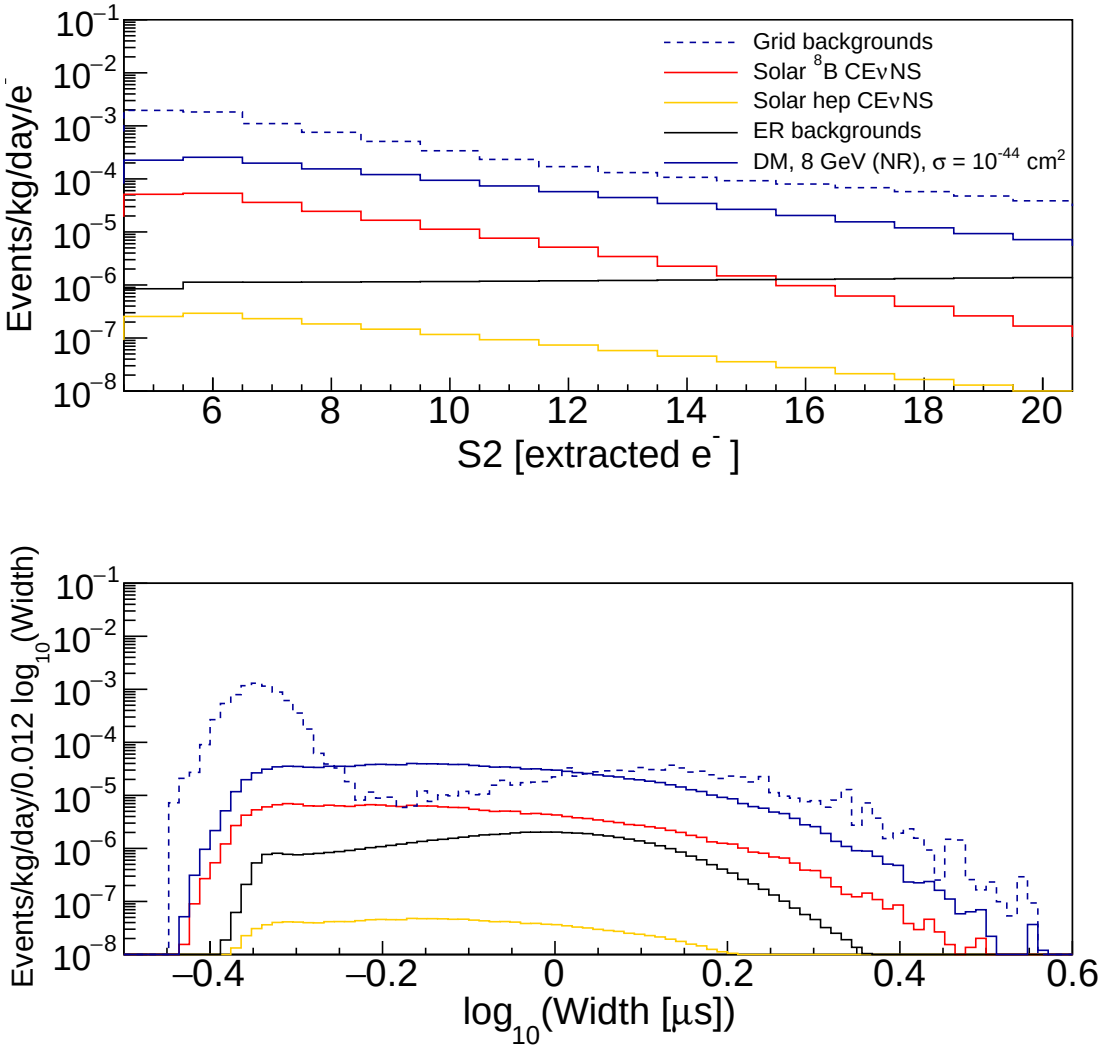

Figure 6.25: The principle backgrounds for an S2-only search are illustrated in both the S2 size (top) and S2 width (bottom) as measured by the aft75-aft25 metric. Despite the addition of the additional effects the pulse width metric optimisation ensures there are regions which are close to being solar neutrino dominated in width. Hence, there is still some improvement from doing this with S2 size alone. ER backgrounds includes contributions from Rn chains, ^{85}Kr and pp neutrinos.

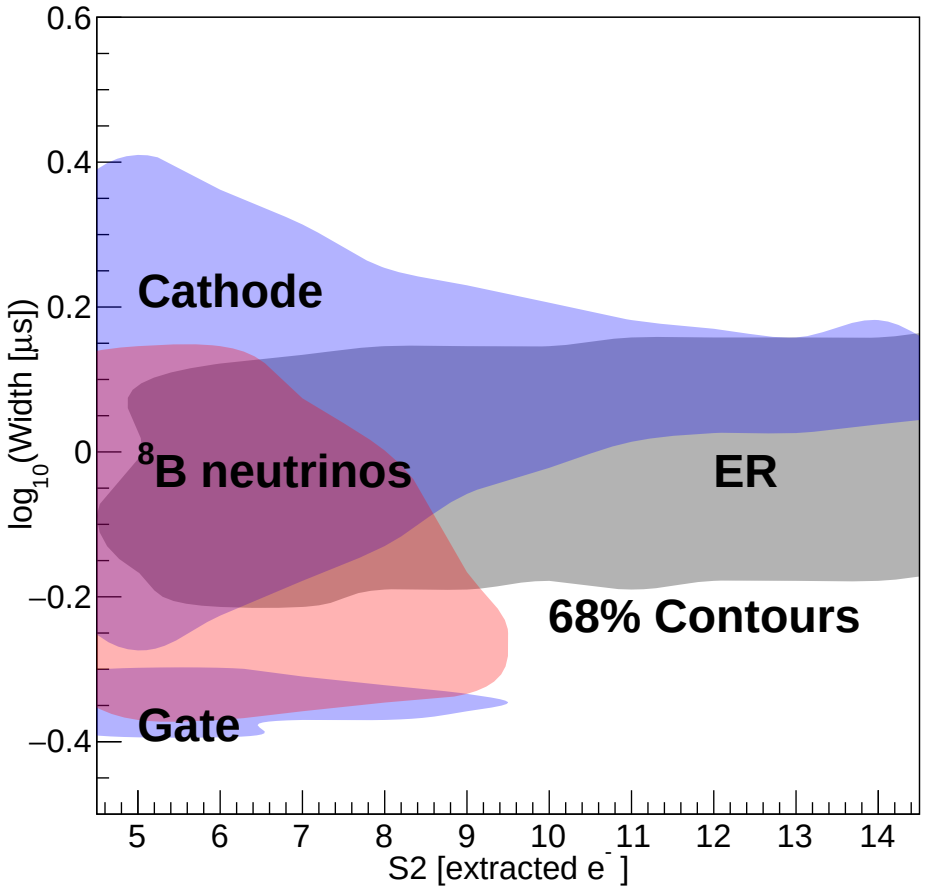

Figure 6.26: Two-dimensional PDFs from S2 principle backgrounds. In all cases these have been normalised. It is possible to see that there are some regions where the grid backgrounds are reduced and should improve sensitivity to dispersed source signals. Here, only the 68% contours are shown for clarity.

To illustrate the agreement between the toy models and the new method discussed here comparisons can be performed in one dimensional plots as shown in Figures 6.27 and 6.28. Also, included are reference LZ publication [141] plots and the output for the toy model. These act as validation that the results of this new model are consistent with the previous approach. There are a few differences which are evident within the spectra, and are likely due to the differences in the detector conditions between the parametrised toy model and the full chain simulations. Firstly, the rate in the first bin is greater in both cases. As radial variation of the light collection efficiency is now included, this leads to more smearing and so events of much higher numbers of electrons can now be found in nearby bins and hence the spectrum shows a softer cut-off at threshold. In addition, the values are slightly below the toy model in the case of ^{220}Rn and a few bins for ^{8}B. This level of variation is expected due to differences in reconstruction, radial effects of light detection efficiency, and also the loss of signal into the electron trains as illustrated in Figure 6.24. Furthermore, upon comparing the grids spectra the method also shows good agreement with the toy model, despite using a very different method. Although the agreement here is worse than for other sources, as the grids spectra are only an estimate then this level of reproduction is acceptable. The differences between the LZ publication and the updated toy model are solely due to the updates of the method to correct the spectrum as explained earlier.

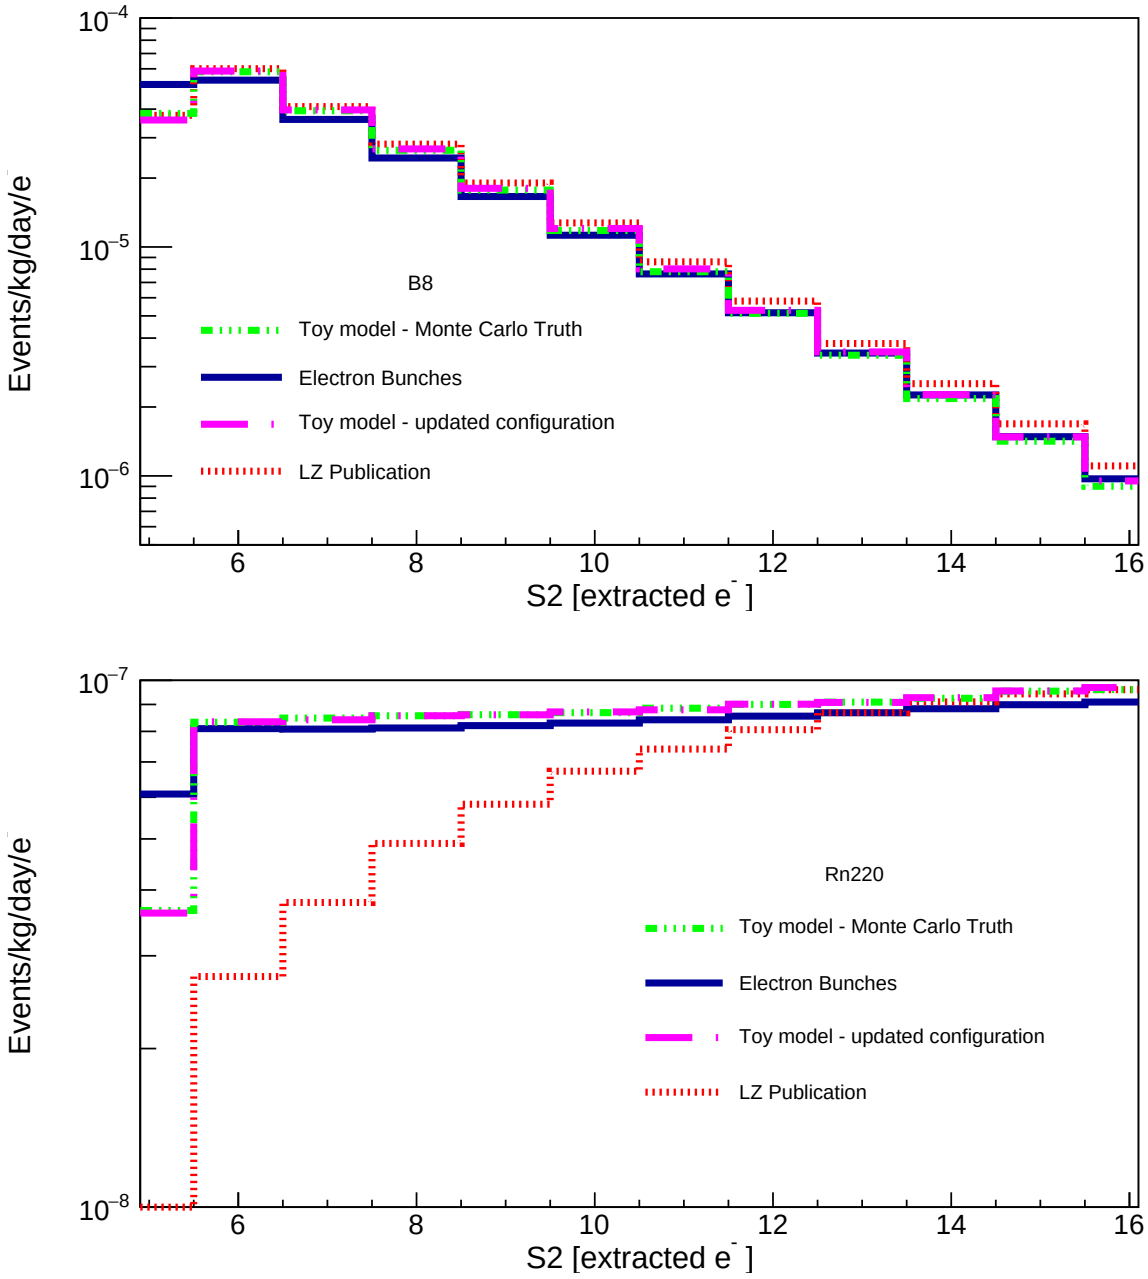

Figure 6.27: Comparison plots between the S2-only spectra for ^{8}B (^{220}Rn chain) events on the top (bottom). The toy model is the predicted spectrum using the fast parametrised model from NEST simulations with the updated configuration. The electron bunches are the output following the parametrised model, but using the more accurate simulations to perform much of the work. The results as used in the LZ Publication (Reference [141]) are included for comparison. The differences between these three spectra arise from different lower energy threshold cuts applied and detector effects in the simulations.

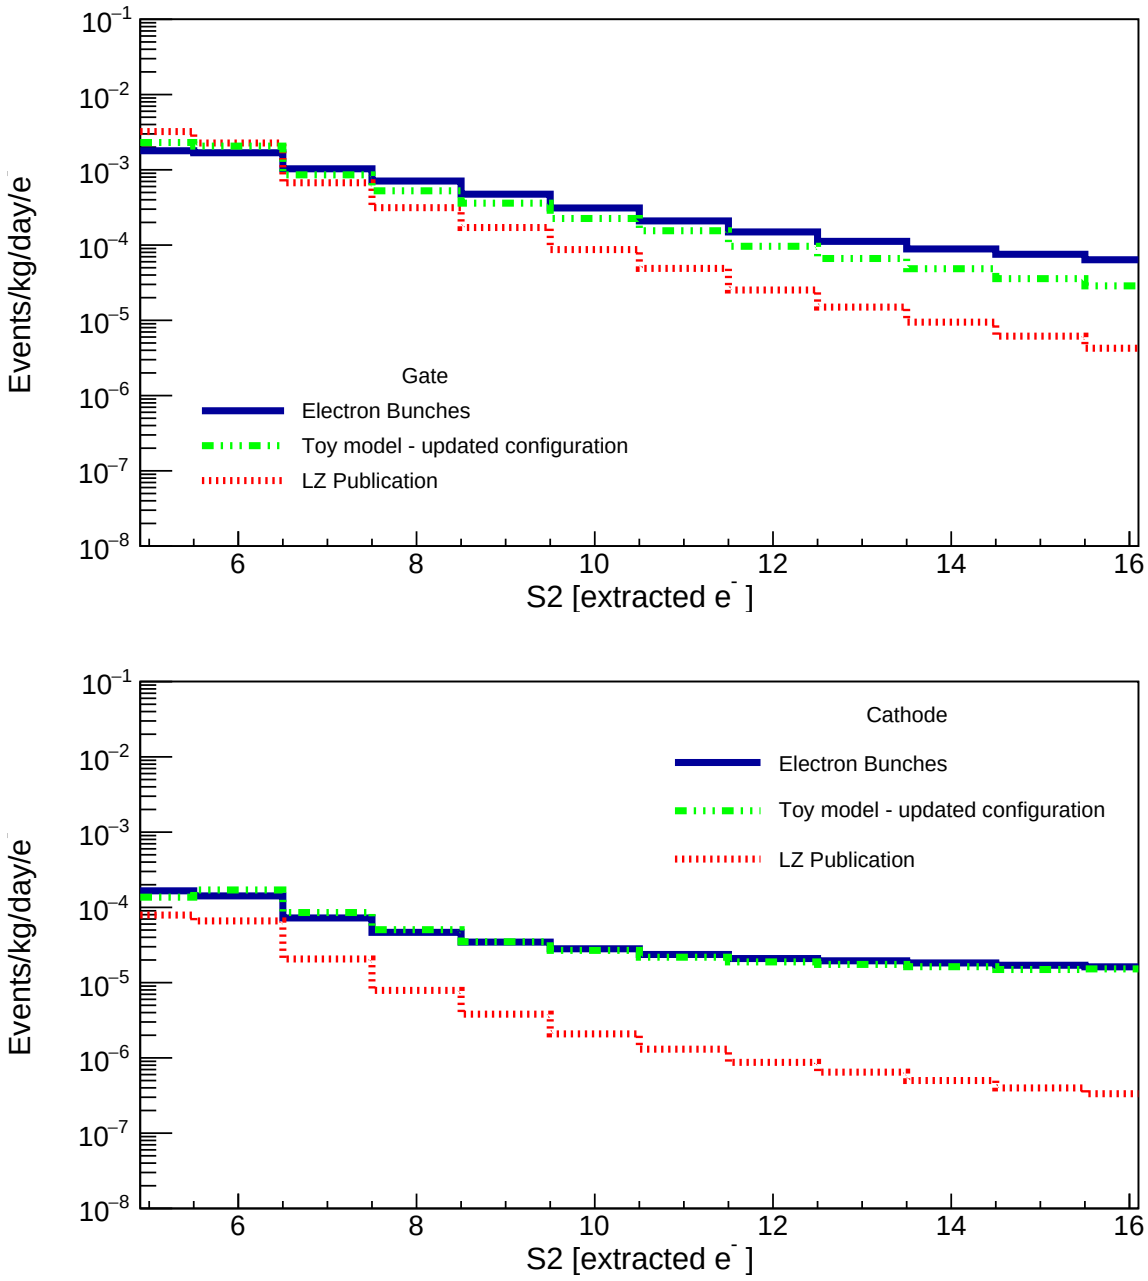

Figure 6.28: Comparison plots between the S2-only spectra for Gate and Cathode events. The toy model is the predicted spectrum using the fast parametrised model from NEST simulations with the updated configuration. The electron bunches are the output following the parametrised model, but using the more accurate simulations to perform much of the work. The results as used in the LZ Publication (Reference [141]) are included for comparison.

6.2.5 Position resolution and radial cut

The minimum S2 size considered within this study is driven by the position resolution. This is to ensure that the mass of the active fiducialised volume is well constrained and backgrounds from the walls (Chapter 3) can be removed. As explained in Section 3.1.2 the position reconstruction is determined by the PMT hit pattern on the top array with the Mercury algorithm [105]. The position resolution is expected to be better with larger S2 signals due to PMTs being better populated with photon hits and hence the relative statistical uncertainty is lower. Additional systematic factors occur due to the electric fields within the detector and hence need to be considered separately.

Figure 6.29 shows how the statistical variation is due primarily to the uncertainty from low statistics at smaller S2 sizes. This size dependence is expected from Poissonian errors and follows an inverse square-root dependence on S2 size as is fitted in the Figure. It also shows that the position uncertainty from the reconstruction is not radially dependent, which is expected from good experiment design. In both cases the uncertainty is calculated based on the covariance matrix elements, transformed to obtain the radial uncertainty.

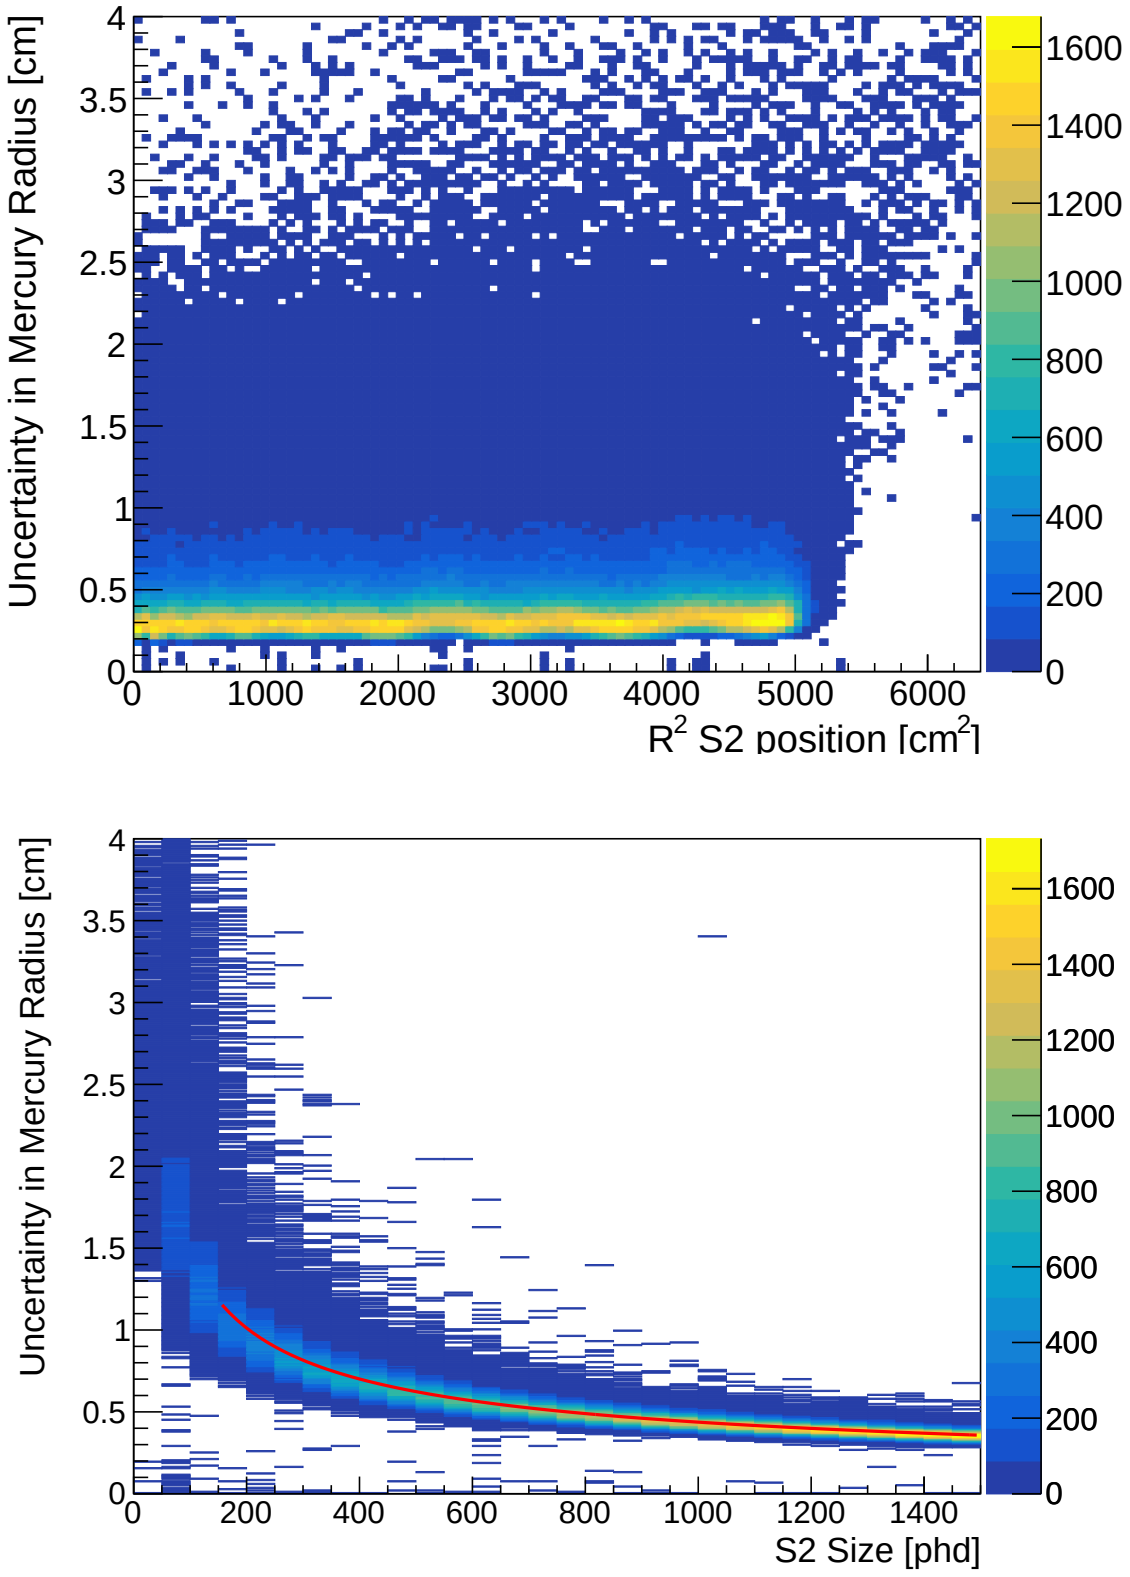

Figure 6.29: Variation in the position reconstruction uncertainty using the covariance from the PMT hit-pattern using the Mercury algorithm. The top plot shows no dependence of the uncertainty on radius, and the bottom plot shows the expected fall as the square root of the S2 size.

Systematic variation primarily occurs when the electron cloud drifts inwards due to the radial component of the electric field being non-zero. With the nominal S1-S2 analysis then information about the event depth enables this to be corrected. As Figure 6.30 shows, there is about a 1 cm deviation from the mean of the true and reconstructed positions using the LZ Monte Carlo (discussed in Section 3.5.3). There is the possibility that some of this offset can be corrected for, though this is not considered in this study. The wall background is not being modelled, as it is expected to be massively sub dominant within the fiducial region. Superimposed on this plot is also the statistical uncertainty previously considered.

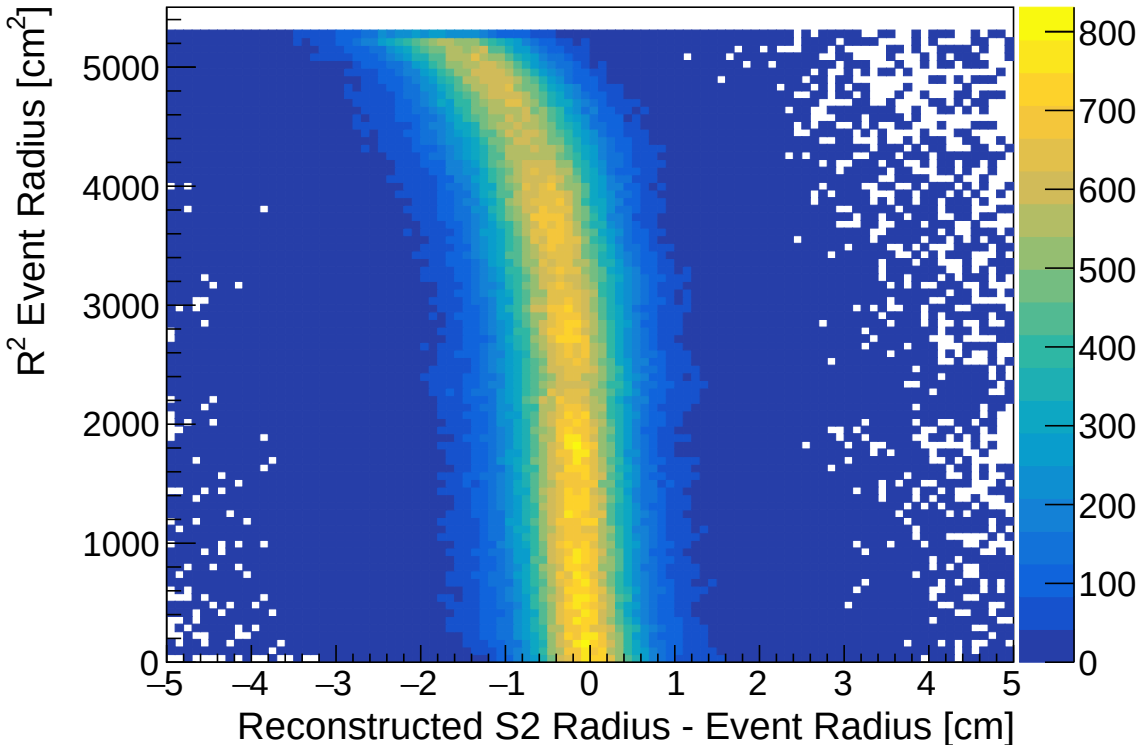

Figure 6.30: Comparison between the Monte Carlo truth position and the reconstructed position. This shows the systematic offset in the reconstruction due to the non-zero component of the electric field close to the walls. This leads to an approximate 1cm deviation on the mean.

The expected number of photons detected for 5 electrons in LZ is around 400. Below this value of photons detected the statistical uncertainty becomes much larger and hence restrictions of events using a radial cut becomes increasingly difficult. Hence, this is decided to be the lower analysis S2 cut for this study. During the running of the ex-

periment this threshold may be reduced further once the detector is better understood. For example the Xenon 1T analysis considered S2 signals down to 4.5 electrons, as is also done in LUX [147]. Hence, the study considered here should again be interpreted as being conservative. The exponentially falling signal models ensure that should this threshold be reduced then the sensitivity should improve further. The value of 1 cm is chosen to ensure that the statistical uncertainty does not dominate over the systematic offset.

As discussed in Chapter 3 the presence of event background topologies closer to the TPC walls necessitates the introduction of a radial cut. For this study the nominal radial cut considered in the WIMP analysis (4 cm from the wall) is used. With real data a data driven approach will be used to determine such a radial cut. However, there are reasons to suspect that the radial cut will not be significantly different from the one considered here and in the nominal S1-S2 search. Radiation leading to electronic recoil backgrounds will propagate less through the xenon than higher energy radiation. Hence, the self shielding does not need to be as large. Additionally, the grid background leakage into the xenon bulk is likely to mean that the effective radially independent grid background is comparable to the rate closer to the wall than in the nominal search. In combination these two effects mean that the theoretical cut could be closer to the walls. This effect however largely offsets the increase in the position uncertainty as previously discussed. Considering this together leads to the lower analysis cut in S2 size being driven by the position resolution. In addition, placing the radial cut at 4 cm from the wall ensures that the volume can be restricted. In total this therefore leads to a fiducial volume, considered in the analysis of 6.3 t with an uncertainty of about 4%.

6.2.6 Impact on sensitivity of more realistic models

The impact of this work on the ability of LZ to search for Dark Matter is a key metric to evaluate its successes. Using some modifications to the code structure that was developed for the statistical inference for Reference [112] for the nominal S1-S2 analysis, a sensitivity curve can be produced.

The general principle behind a statistical test is to evaluate a region of parameter space from which the presence of a signal can be rejected to some confidence level (here 90% is considered). To do this, a test-statistic for both the null (no signal) and alternative (signal present) hypotheses is constructed for different mass and signal strengths. A test statistic is a one dimensional representation of the data which can be used to evaluate the probabilities that the data meets each hypothesis [148][149]. This reduces the dimensionality down to a scalar function from which *p-values* can be evaluated. A profile likelihood ratio (PLR) is a common test statistic to use, especially within particle physics experiments such as LZ. It is a formulation of Neyman-Pearson lemma, which expresses that this is the best statistical discriminator [148]. By the PLR method, the test statistic is given by

$$\lambda = \frac{\mathcal{L}(\mu, \hat{\hat{\nu}})}{\mathcal{L}(\hat{\mu}, \hat{\nu})} \tag{6.6}$$

where $\mathcal{L}$ is the likelihood, μ donates the value of the parameter of interest (signal strength) and the ν relates to the vector of nuisance parameters, which contains the normalisations for the background sources. The normalisations and their uncertainties considered here are detailed in Table 6.2. Here, hats donate the maximum likelihood estimator, and the double hats relate to the conditional maximum likelihood estimator. The likelihood $\mathcal{L}$ is given by

$$\mathcal{L} = \text{Pois}(n_0, \mu + \mu_b) \times \prod_{i=1}^{n_0} \frac{1}{\mu + \mu_b} \left(\mu f_s(\mathbf{x_i}|m_{\text{wimp}}) + \sum_{j=1}^{N_b} \mu_b f_b(\mathbf{x_i}|\nu) \right) \times \prod_{p=1}^{N_p} f_p(\mathbf{g_p}|\nu_\mathbf{p}) \tag{6.7}$$

where the first term is the Poisson variation in the number of signal events, the central term examines the model of the expected number of events and the third term constrains the global parameters. Here n_0 is the number of events in the data set, over which the

product of each is taken with the index i. This central term contains the information based on the expected estimation from the probability distribution functions, f_s for signal and f_j for each background. For this study several sources of background are grouped together, as shown within Figures 6.26 and 6.25. This includes the electronic recoil backgrounds, the neutrino coherent scatters and the grid (cathode and gate) backgrounds.

It is often more common to use a derived test statistic t_μ which recasts the range of the parameter values allowed,

$$t_\mu = -2 \log \lambda(\mu). \tag{6.8}$$

Because both an over and under fluctuation of the signal can cause an incompatibility between the two hypotheses, this model is considered a two-tailed test. Estimating the test-statistic distributions can be done by running toy experiments involving sampling from the PDFs as obtained in the last section for each value of mass and for differing signal strengths. With each of these toy experiments the test statistic distribution is obtained by evaluating Equation 6.8 on the data. Using the median value for the alternative hypothesis as the best estimation for the 'observed data' evaluation of the test statistic enables a sensitivity study to be performed. The study uses the same Dark Matter parameters as discussed in Section 2.1.3 and used in Reference [112].

As Figure 6.31 shows, the sensitivity of LZ to low mass Dark Matter is improved compared with the nominal S1-S2 analysis. The bands show the one and two standard deviation confidence limits on the sensitivity. Significant improvement over an S2-only analysis, where only the signal size is considered, is also shown. In this case, the one dimensional PDFs are constructed by marginalising over the width parameter, and no further cuts have been applied. For completeness, the median sensitivity line for the LZ S2-only publication is included alongside the median sensitivity for the updated toy model which shows how optimistic this approach is. As discussed previously, the diffusion constant now considered is smaller than in Reference [141] as well as some different background assumptions and is reflected in Table 6.2 showing the predicted counts for 1000 days of exposure.

Background		Counts after cuts	
Type	Source	S2-only	σ/N
NR	ν ^{8}B	1481	4%
	ν hep	11.0	15%
	ν DSN	0.3	50%
	ν atm	0.5	25%
ER	^{222}Rn	166.4	10%
	^{220}Rn	27.4	10%
	ν pp	47.0	2%
	^{85}Kr	143.0	20%
	materials‡	52000	20%

Table 6.2: Expected number of events for the S2-only search as derived from the 6.3 tonnes nominal fiducial volume and 1000 days of exposure. A cut at 5 single electrons is applied to ensure position resolution in XY is better than 1 cm. The NR (ER) refers to the nuclear (electronic) recoil components. The dominant backgrounds are shown within Figures 6.26 and 6.25, where all the ER backgrounds (excluding grids) are grouped together.
‡ Note that this component includes the cathode, grid, accidental overlaps of single electrons as well as surface events and is obtained from scaling spectra from the LUX S2-only analysis.

6.3 Conclusion

The study shows that the sensitivity to Dark Matter interactions for the LUX-ZEPLIN experiment can be extended by considering an S2-only analysis. Considering the phenomena that affects such a study has been a focus of this Chapter. This required optimising the pulse width metric and the maximum pulse width to maximise the sensitivity to low mass Dark Matter. The work also shows that this sensitivity is not greatly different from the published study which made more basic assumptions.

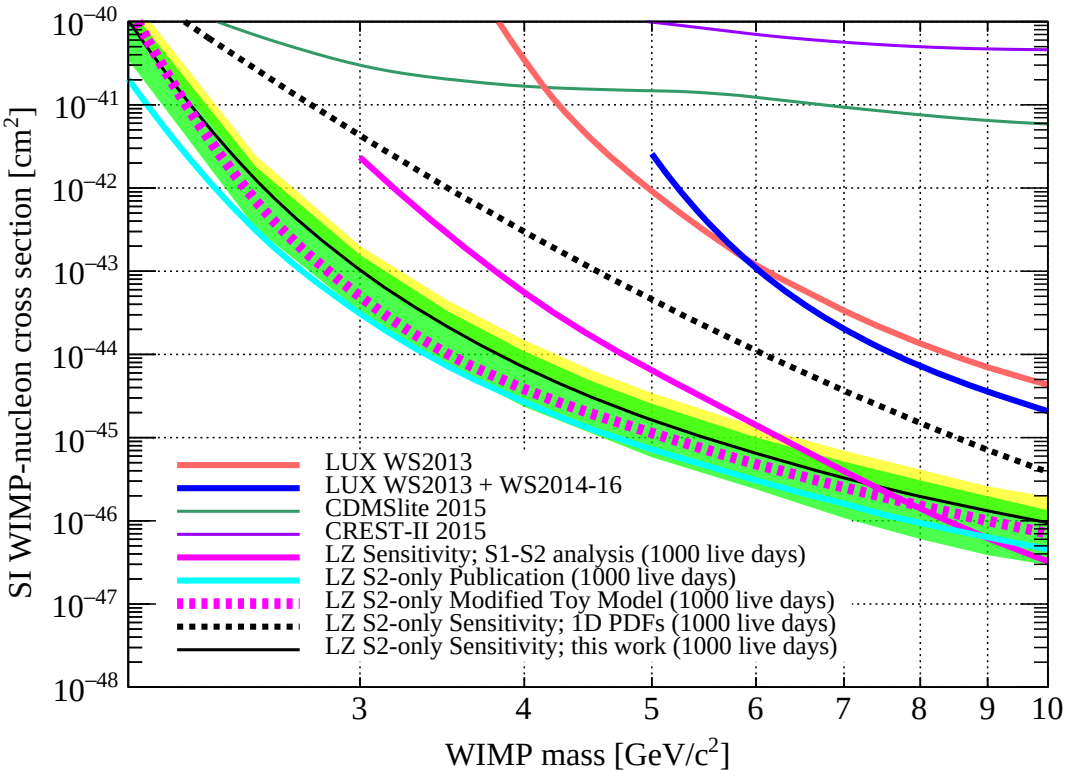

Figure 6.31: Expected sensitivity of LZ to spin-independent WIMP scattering from Xe nuclei. In the case of the most realistic model, curves are shown for the 90% confidence limit with the one and two standard deviation bands. Alongside former experimental limits, the nominal S1-S2 analysis sensitivity curve is also included, with three sensitivity curves for S2-only analyses. The publication curve from Reference [141] is shown to be optimistic when compared with the result from this study. However, this shows that a consideration of the optimised width metric still has an improvement over the one dimensional S2 size case. The comparison curves are included from LUX [150], CDMSlite [151] and CRESST-II [152].

Conclusion

This thesis has outlined the experimental evidence for Dark Matter from several sources. It reveals that the current theory of the universe is incomplete, and a favoured solution is to add the presence of non-interacting, non-relativistic matter. Several candidates are outlined that span many orders of magnitude in mass.

This work makes significant contributions to the LUX-ZEPLIN detector, which improve the ability of the experiment to detect Dark Matter. The detector is a 10 tonne liquid xenon time projection chamber, situated at the Sanford Underground Research Facility in South Dakota and is searching for interactions of Dark Matter particles with xenon nuclei.

More specifically the work involves the construction and calibration of a thermometry system which enables the detector to operate with stability in detector conditions, and ensured the cool-down procedure could be enacted without damage. The thermometers are shown to better the experimental requirements which arise from these functions.

Additionally, substantial consideration is given to the origins of the position dependence on the S2 signal. The cool-down of the detector has been shown to be uniform despite the thermally sensitive Teflon TPC that could have contracted non-uniformly giving rise to additional S2 position dependence. The level sensors have also been characterised and are sensitive to departures of liquid level during operations. In combination with the S2 width position dependence from events within the extraction region, the level sensors were used to provide information about the tilt of the inner cryostat, enabling an adjustment to be performed. A study is performed, showing how the position dependence

can be removed using dispersed source calibration sources. This is an essential requirement for a successful Dark Matter search. Crucially, with early data, the attenuation due to xenon impurities is shown to be better than expected. This allows detector commissioning to continue and shows the successful implementation within the LZ analysis framework.

The thesis concludes by utilising the diffusion of drifting electrons to enable the sensitivity to be improved below the nominal threshold. Techniques introduced include a consideration of the optimal pulse width metric for S2 pulses and a restriction to only consider contributions within the first 6 µs of the S2 signal. This limits the effects of delayed electrons within an event affecting the pulse width measured. These techniques ensure that an LZ S2-only analysis is substantially advanced in development with realistic phenomena. The analysis will be performed in addition to the nominal search for Dark Matter and by removing the requirement of a scintillation signal (S1); this lowers the threshold. Should a Dark Matter candidate lie within this region of phase space the studies presented here show that a detection is possible.

Bibliography

[1] Aghanim N. et al. Planck 2018 results. VI. Cosmological parameters. arXiv:1807.06209 [astro-ph] (2018).

[2] Zwicky F. On the Masses of Nebulae and of Clusters of Nebulae. The Astrophysical Journal, 86:217 (1937).

[3] van Aleada T.S. et al. DISTRIBUTION OF DARK MATTER IN THE SPIRAL GALAXY NGC 3198. The Astrophysical Journal, 295:305–313 (1985).

[4] Rubin V.C. and Ford Jr. W.K. Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions. The Astrophysical Journal, 159:379 (1970).

[5] Zyla P. and et. al. Review of particle physics. PTEP, 2020(8):083C01 (2020).

[6] Milgrom M. A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis. The Astrophysical Journal, 270:365 (1983).

[7] Bekenstein J.D. Relativistic gravitation theory for the MOND paradigm. Physical Review D, 71(6):069901 (2005).

[8] Famaey B. and McGaugh S. Modified Newtonian Dynamics (MOND): Observational Phenomenology and Relativistic Extensions. Living Reviews in Relativity, 15(1):10 (2012).

[9] Vankov A. Einstein's Paper: "Explanation of the Perihelion Motion of Mercury from General Relativity Theory". page 34.

[10] Wambsganss J. Gravitational Lensing. J.P. Françoise, G.L. Naber, and T.S. Tsun, editors, Encyclopedia of Mathematical Physics, pages 567–575. Academic Press, Oxford (2006).

[11] Massey R., Kitching T., and Richard J. The dark matter of gravitational lensing. Reports on Progress in Physics, 73(8):086901 (2010).

[12] Madejski G. Recent and Future Observations in the X-ray and Gamma-ray Bands: Chandra, Suzaku, GLAST, and NuSTAR. https://digital.library.unt.edu/ark:/67531/metadc877627/m1/3/ (2005).

[13] Angus G.W., Famaey B., and Zhao H. Can MOND take a bullet? Analytical comparisons of three versions of MOND beyond spherical symmetry. Monthly Notices of the Royal Astronomical Society, 371(1):138–146 (2006).

[14] Chandra :: Photo Album :: 1E 0657-56 :: More Images of 1E 0657-56. https://chandra.harvard.edu/photo/2006/1e0657/more.html.

[15] Primack J.R. Dark Matter and Galaxy Formation. 1192:101–137 (2009).

[16] Bertone G. and Hooper D. A History of Dark Matter. Reviews of Modern Physics, 90(4):045002 (2018).

[17] Penzias A.A. and Wilson R.W. A Measurement of Excess Antenna Temperature at 4080 Mc/s. The Astrophysical Journal, 142:419–421 (1965).

[18] Liddle A.R. An Introduction to Modern Cosmology [Electronic Resource]. Ebook Central. Wiley, Chichester, England, third edition. edition (2015).

[19] Fixsen D.J. THE TEMPERATURE OF THE COSMIC MICROWAVE BACKGROUND. The Astrophysical Journal, 707(2):916–920 (2009).

[20] Peebles P.J.E. and Yu J.T. Primeval Adiabatic Perturbation in an Expanding Universe. The Astrophysical Journal, 162:815 (1970).

[21] Akrami Y. et al. Planck 2018 results. I. Overview and the cosmological legacy of Planck. arXiv:1807.06205 [astro-ph] (2018).

[22] ESA Science & Technology - Planck's power spectrum of temperature fluctuations in the Cosmic Microwave Background. https://sci.esa.int/web/planck/-/51555-planck-power-spectrum-of-temperature-fluctuations-in-the-cosmic-microwave-background.

[23] Freese K. Status of Dark Matter in the Universe. International Journal of Modern Physics D, 26(06):1730012 (2017).

[24] Simon J.D. and Geha M. The Kinematics of the Ultra-faint Milky Way Satellites: Solving the Missing Satellite Problem. The Astrophysical Journal, 670(1):313–331 (2007).

[25] Klypin A. et al. Where Are the Missing Galactic Satellites? The Astrophysical Journal, 522(1):82 (1999).

[26] Moore B. et al. Dark Matter Substructure within Galactic Halos. The Astrophysical Journal, 524(1):L19 (1999).

[27] Simon J.D. The Faintest Dwarf Galaxies. Annual Review of Astronomy and Astrophysics, 57(1):375–415 (2019).

[28] Kouvaris C. Limits on Self-Interacting Dark Matter. Physical Review Letters, 108(19):191301 (2012).

[29] Edward Kolb. The Early Universe. CRC Press (2018).

[30] Chepel V. and Araújo H. Liquid noble gas detectors for low energy particle physics. Journal of Instrumentation, 8(04):R04001–R04001 (2013).

[31] Feng J.L. and Kumar J. The WIMPless Miracle: Dark Matter Particles without Weak-scale Masses or Weak Interactions. Physical Review Letters, 101(23):231301 (2008).

[32] Peter A.H.G. Dark Matter: A Brief Review. arXiv:1201.3942 [astro-ph] (2012).

[33] Martin S.P. A Supersymmetry Primer. arXiv:hep-ph/9709356, 18:1–98 (1998).

[34] Peccei R.D. and Quinn H.R. CP Conservation in the Presence of Pseudoparticles. Physical Review Letters, 38(25):1440–1443 (1977).

[35] Perkins D.H. Particle Astrophysics. Oxford Master Series in Physics ; No. 10. University Press, Oxford, 2nd ed. edition (2008).

[36] Primakoff H. Photo-Production of Neutral Mesons in Nuclear Electric Fields and the Mean Life of the Neutral Meson. Physical Review, 81(5):899–899 (1951).

[37] Sikivie P. Experimental Tests of the "Invisible" Axion. Physical Review Letters, 51(16):1415–1417 (1983).

[38] Freese K., Fields B., and Graff D. Limits on Stellar Objects as the Dark Matter of Our Halo: Nonbaryonic Dark Matter Seems to be Required. arXiv:astro-ph/9904401 (1999).

[39] Tisserand P. et al. Limits on the Macho Content of the Galactic Halo from the EROS-2 Survey of the Magellanic Clouds. Astronomy & Astrophysics, 469(2):387–404 (2007).

[40] Abbott B.P. et al. Observation of Gravitational Waves from a Binary Black Hole Merger. Physical Review Letters, 116(6):061102 (2016).

[41] Mediavilla E. et al. Limits on the Mass and Abundance of Primordial Black Holes from Quasar Gravitational Microlensing. The Astrophysical Journal, 836(2):L18 (2017).

[42] Kopp J. Sterile Neutrinos as Dark Matter Candidates. arXiv:2109.00767 [hep-ph] (2021).

[43] Dodelson S. and Widrow L.M. Sterile neutrinos as dark matter. Physical Review Letters, 72(1):17–20 (1994).

[44] Ng K.C.Y. et al. New Constraints on Sterile Neutrino Dark Matter from $NuSTAR$ M31 Observations. Physical Review D, 99(8):083005 (2019).

[45] Abratenko P. et al. Search for Neutrino-Induced Neutral Current Δ Radiative Decay in MicroBooNE and a First Test of the MiniBooNE Low Energy Excess Under a Single-Photon Hypothesis. page 7.

[46] An H. et al. Direct Detection Constraints on Dark Photon Dark Matter. Physics Letters B, 747:331–338 (2015).

[47] Carr J., Lamanna G., and Lavalle J. Indirect detection of dark matter. Reports on Progress in Physics, 69(8):2475–2512 (2006).

[48] Aad G. et al. The ATLAS Experiment at the CERN Large Hadron Collider. Journal of Instrumentation, 3(08):S08003–S08003 (2008).

[49] Chatrchyan S. et al. The CMS experiment at the CERN LHC. Journal of Instrumentation, 3(08):S08004–S08004 (2008).

[50] Cerdeno D.G. and Green A.M. Direct detection of WIMPs. Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms, 269(4):498–503 (2011).

[51] McCabe C. Astrophysical uncertainties of dark matter direct detection experiments. Physical Review D, 82(2):023530 (2010).

[52] Binney J. Galactic Dynamics. Princeton Series in Astrophysics. Princeton University Press, Princeton, N.J. ; Oxford, 2nd ed. edition (2008).

[53] Moore B. Evidence against dissipation-less dark matter from observations of galaxy haloes. Nature, 370(6491):629–631 (1994).

[54] Peter A.H.G. Particle Dark Matter in the Solar System. Ph.D. thesis, Princeton University, United States – New Jersey (2008).

[55] Engel J. Nuclear form factors for the scattering of weakly interacting massive particles. Physics Letters B, 264(1):114–119 (1991).

[56] Lewin J. and Smith P. Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil. Astroparticle Physics, 6(1):87–112 (1996).

[57] Duda G., Kemper A., and Gondolo P. Model Independent Form Factors for Spin Independent Neutralino-Nucleon Scattering from Elastic Electron Scattering Data. Journal of Cosmology and Astroparticle Physics, 2007(04):012–012 (2007).

[58] Peter A.H.G. et al. WIMP physics with ensembles of direct-detection experiments. Physics of the Dark Universe, 5–6:45–74 (2014).

[59] Olcina Samblas, Ibles. Private communication (2021).

[60] Boehm C. et al. How high is the neutrino floor? arXiv:1809.06385 [astro-ph, physics:hep-ph] (2018).

[61] Mount B.J. et al. LUX-ZEPLIN (LZ) Technical Design Report. arXiv:1703.09144 [astro-ph, physics:hep-ex, physics:physics] (2017).

[62] Akerib D.S. et al. Calibration, event reconstruction, data analysis, and limit calculation for the LUX dark matter experiment. Physical Review D, 97(10) (2018).

[63] Cui X. et al. Dark Matter Results from 54-Ton-Day Exposure of PandaX-II Experiment. Physical Review Letters, 119(18):181302 (2017).

[64] Aprile E. et al. Dark Matter Search Results from a One Tonne x Year Exposure of XENON1T. arXiv:1805.12562 [astro-ph, physics:hep-ex, physics:hep-ph, physics:physics] (2018).

[65] Agnes P. et al. Low-Mass Dark Matter Search with the DarkSide-50 Experiment. Physical Review Letters, 121(8):081307 (2018).

[66] Ajaj R. et al. Search for dark matter with a 231-day exposure of liquid argon using DEAP-3600 at SNOLAB. arXiv:1902.04048 [astro-ph] (2019).

[67] Agnese R. et al. Results from the Super Cryogenic Dark Matter Search Experiment at Soudan. Phys. Rev. Lett., 120(6):061802 (2018).

[68] Abdelhameed A.H. et al. First results from the CRESST-III low-mass dark matter program. arXiv:1904.00498 [astro-ph] (2019).

[69] Bernabei R. et al. New limits on particle dark matter search with a liquid Xenon target-scintillator. Physics Letters B, 436(3):379–388 (1998).

[70] Alner G.J. et al. First limits on nuclear recoil events from the ZEPLIN I galactic dark matter detector. Astroparticle Physics, 23(5):444–462 (2005).

[71] Alner G.J. et al. First limits on WIMP nuclear recoil signals in ZEPLIN-II: A two phase xenon detector for dark matter detection. Astroparticle Physics, 28(3):287–302 (2007).

[72] Akimov D.Y. et al. WIMP-nucleon cross-section results from the second science run of ZEPLIN-III. Physics Letters B, 709(1):14–20 (2012).

[73] Akimov D.Y. et al. The ZEPLIN-III dark matter detector: Instrument design, manufacture and commissioning. Astroparticle Physics, 27(1):46–60 (2007).

[74] Abe K. et al. A direct dark matter search in XMASS-I. Physics Letters B, 789:45–53 (2019).

[75] Tan A. et al. Dark Matter Results from First 98.7 Days of Data from the PandaX-II Experiment. Physical Review Letters, 117(12):121303 (2016).

[76] Meyers P.D. et al. DarkSide-50: A WIMP Search with a Two-phase Argon TPC. Physics Procedia, 61:124–129 (2015).

[77] Aalseth C.E. et al. DarkSide-20k: A 20 Tonne Two-Phase LAr TPC for Direct Dark Matter Detection at LNGS. The European Physical Journal Plus, 133(3):131 (2018).

[78] Baudis L. WIMP dark matter direct-detection searches in noble gases. Physics of the Dark Universe, 4:50–59 (2014).

[79] Amaudruz P.A. et al. Design and construction of the DEAP-3600 dark matter detector. Astroparticle Physics, 108:1–23 (2019).

[80] CDMS Story. https://web.archive.org/web/20020601184806/http://cdms.berkeley.edu/experiment.html (2002).

[81] Agnese R. et al. Demonstration of surface electron rejection with interleaved germanium detectors for dark matter searches. Applied Physics Letters, 103(16):164105 (2013).

[82] Agnese R. et al. CDMSlite: A Search for Low-Mass WIMPs using Voltage-Assisted Calorimetric Ionization Detection in the SuperCDMS Experiment. Physical Review Letters, 112(4):041302 (2014).

[83] Armengaud E. et al. Final results of the EDELWEISS-II WIMP search using a 4-kg array of cryogenic germanium detectors with interleaved electrodes. Physics Letters B, 702(5):329–335 (2011).

[84] Juillard A. Status of the EDELWEISS-II Experiment. Journal of Low Temperature Physics, 151(3):806–811 (2008).

[85] Jochum J. et al. The CRESST dark matter search. Progress in Particle and Nuclear Physics, 66(2):202–207 (2011).

[86] Bernabei R. et al. First model independent results from DAMA/LIBRA-phase2. Nuclear Physics and Atomic Energy, 19(4):307–325 (2018).

[87] Amaré J. et al. Annual modulation results from three-year exposure of ANAIS-112. Physical Review D, 103(10):102005 (2021).

[88] Aprile E. et al. Projected WIMP Sensitivity of the XENONnT Dark Matter Experiment. arXiv:2007.08796 [astro-ph, physics:hep-ex, physics:physics] (2020).

[89] Schumann M. DARWIN. Science, 172 (2011).

[90] Szydagis M. et al. A review of basic energy reconstruction techniques in liquid xenon and argon detectors for dark matter and neutrino physics using nest. Instruments, 5(1):13– (2021).

[91] Sorensen P. and Dahl C.E. Nuclear recoil energy scale in liquid xenon with application to the direct detection of dark matter. Physical Review D, 83(6):063501 (2011).

[92] Dahl C.E. The Physics of Background Discrimination in Liquid Xenon, and First Results from Xenon10 in the Hunt for WIMP Dark Matter. Ph.D. thesis, Princeton University, United States – New Jersey (2009).

[93] López Paredes B. et al. Response of photomultiplier tubes to xenon scintillation light. Astroparticle Physics, 102:56–66 (2018).

[94] Lenardo B. et al. A Global Analysis of Light and Charge Yields in Liquid Xenon. IEEE Transactions on Nuclear Science, 62(6):3387–3396 (2015).

[95] Miyajima M. et al. Average energy expended per ion pair in liquid argon. Physical Review A, 9(3):1438–1443 (1974).

[96] Akerib D.S. et al. Discrimination of electronic recoils from nuclear recoils in two-phase xenon time projection chambers. Physical Review D, 102(11):112002 (2020).

[97] Doke T. et al. Absolute Scintillation Yields in Liquid Argon and Xenon for Various Particles. Japanese Journal of Applied Physics, 41(Part 1, No. 3A):1538–1545 (2002).

[98] Shutt T. et al. Performance and fundamental processes at low energy in a two-phase liquid xenon dark matter detector. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 579(1):451–453 (2007).

[99] Szydagis M. et al. NEST: A comprehensive model for scintillation yield in liquid xenon. Journal of Instrumentation, 6(10):P10002 (2011).

[100] Aiham, Al Musalhi. Private communication (2021).

[101] Akerib D.S. et al. The LUX-ZEPLIN (LZ) experiment. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 953:163047 (2020).

[102] Davis R. A review of the homestake solar neutrino experiment. Progress in particle and nuclear physics, 32(C):13–32 (1994).

[103] Linehan R. et al. Design and production of the high voltage electrode grids and electron extraction region for the LZ dual-phase xenon time projection chamber. arXiv:2106.06622 [astro-ph, physics:hep-ex, physics:physics] (2021).

[104] Neves F. et al. Measurement of the absolute reflectance of polytetrafluoroethylene (PTFE) immersed in liquid xenon. Journal of Instrumentation, 12(01):P01017–P01017 (2017).

[105] Akerib D.S. et al. Position Reconstruction in LUX. arXiv:1710.02752 [hep-ex, physics:physics] (2017).

[106] Akerib D.S. et al. Identification of radiopure titanium for the LZ dark matter experiment and future rare event searches. Astroparticle Physics, 96:1–10 (2017).

[107] Dobi A. et al. Study of a zirconium getter for purification of xenon gas. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 620(2-3):594–598 (2010).

[108] Abe K. et al. Radon removal from gaseous xenon with activated charcoal. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 661(1):50–57 (2012).

[109] Table of Isotopes decay data. http://nucleardata.nuclear.lu.se/toi/nuclide.asp?iZA=270057.

[110] Akerib D.S. et al. Low-energy (0.7-74 keV) nuclear recoil calibration of the LUX dark matter experiment using D-D neutron scattering kinematics. arXiv:1608.05381 [astro-ph, physics:hep-ex, physics:physics] (2016).

[111] Akerib D.S. et al. Kr83m calibration of the 2013 LUX dark matter search. Physical Review D, 96(11) (2017).

[112] Akerib D.S. et al. Projected WIMP sensitivity of the LUX-ZEPLIN (LZ) dark matter experiment. arXiv:1802.06039 [astro-ph, physics:hep-ex, physics:physics] (2018).

[113] Kudryavtsev V.A. Muon simulation codes MUSIC and MUSUN for underground physics. Computer Physics Communications, 180(3):339–346 (2009).

[114] Hasert F.J. et al. Observation of neutrino-like interactions without muon or electron in the gargamelle neutrino experiment. Physics Letters B, 46(1):138–140 (1973).

[115] Billard J., Strigari L., and Figueroa-Feliciano E. Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments. Physical Review D, 89(2):023524 (2014).

[116] Akerib D.S. et al. Measurement of the gamma ray background in the Davis cavern at the Sanford Underground Research Facility. Astroparticle Physics, 116:102391 (2020).

[117] Meng Y., Busenitz J., and Piepke A. A new method for evaluating the effectiveness of plastic packaging against radon penetration. Applied Radiation and Isotopes, 156:108963 (2020).

[118] Akerib D.S. et al. The LUX-ZEPLIN (LZ) radioactivity and cleanliness control programs. The European Physical Journal C, 80(11):1044 (2020).

[119] Akerib D.S. et al. Simulations of events for the LUX-ZEPLIN (LZ) dark matter experiment. Astroparticle Physics, 125:102480 (2021).

[120] Mathematica, Version 13.0.0. Wolfram Research, Inc., Champaign, Illinois (2021).

[121] Lemmon E.W. et al. NIST Standard Reference Database 23: Reference Fluid Thermodynamic and Transport Properties-REFPROP, Version 10.0, National Institute of Standards and Technology (2018).

[122] Monteiro C. et al. Secondary Scintillation Yield in Pure Xenon. Journal of Instrumentation, 2 (2007).

[123] Akerib D.S. et al. LUX-ZEPLIN (LZ) Conceptual Design Report. arXiv:1509.02910 [astro-ph, physics:hep-ex, physics:physics] (2015).

[124] Michels A. and Wassenaar T. Vapour pressure of liquid xenon. Physica, 16(3):253–256 (1950).

[125] Akerib D.S. et al. The LUX-ZEPLIN (LZ) radioactivity and cleanliness control programs. The European Physical Journal C, 80(11):1044 (2020).

[126] M 222-100-A — Heraeus PT100 Sensor, RS Components. https://uk.rs-online.com/web/p/platinum-resistance-temperature-sensors/4538098/.

[127] Thermal properties of Copper. http://www-ferp.ucsd.edu/LIB/PROPS/PANOS/cu.html.

[128] Jones W. Theoretical Solid State Physics. London, (1973).

[129] Scovell P.R. Internal Note - Rn Assays in Denys Wilkinson Building (2016).

[130] Scovell P.R. Internal Note - Dust Assays in Denys Wilkinson Building (2016).

[131] Lake Shore Cryotronics. PT 100 Reference Sensor. https://www.lakeshore.com/products/categories/overview/temperature-products/cryogenic-temperature-sensors/platinum.

[132] Brechna H. et al. Brookhaven National Laboratory selected cryogenic data notebook (1980).

[133] Oxford Physics: Atmospheric, Oceanic and Planetary Physics: Weather Data. http://eodg.atm.ox.ac.uk/eodg/weather/weather_nocol.html.

[134] Akerib D.S. et al. Investigation of background electron emission in the LUX detector. arXiv:2004.07791 [physics] (2020).

[135] Boast K. A Position-Sensing System for the Top Pmt Array in the Lux-Zeplin Dark Matter Experiment. Ph.D. thesis, ProQuest Dissertations Publishing (2017).

[136] Kirby R.K. Thermal expansion of polytetrafluoroethylene (Teflon) from -190 to +300 C. Journal of Research of the National Bureau of Standards, 57(2):91 (1956).

[137] Hidnert P. Thermal expansion of titanium. page 5.

[138] Farrow G.J. et al. Thermal expansion of PEEK between 80 and 470K. Journal of Materials Science Letters, 9(6):743–744 (1990).

[139] Hahn T.A. Thermal Expansion of Copper from 20 to 800 K—Standard Reference Material 736. page 7.

[140] Essig R., Sholapurkar M., and Yu T.T. Solar Neutrinos as a Signal and Background in Direct-Detection Experiments Searching for Sub-GeV Dark Matter With Electron Recoils. Physical Review D, 97(9):095029 (2018).

[141] Akerib D.S. et al. Enhancing the sensitivity of the LUX-ZEPLIN (LZ) dark matter experiment to low energy signals. arXiv:2101.08753 [astro-ph, physics:physics] (2021).

[142] Gillespie D.T. and Seitaridou E. Simple Brownian Diffusion: An Introduction to the Standard Theoretical Models. Oxford University Press (2012).

[143] Sorensen P. Anisotropic diffusion of electrons in liquid xenon with application to improving the sensitivity of direct dark matter searches. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 635(1):41–43 (2011).

[144] Njoya O. et al. Measurements of electron transport in liquid and gas Xenon using a laser-driven photocathode. arXiv:1911.11580 [nucl-ex, physics:physics] (2019).

[145] Szydagis M. et al. Noble Element Simulation Technique v2.0. Zenodo (2018).

[146] Akerib D.S. et al. Improving sensitivity to low-mass dark matter in LUX using a novel electrode background mitigation technique. arXiv:2011.09602 [astro-ph, physics:hep-ex, physics:physics] (2020).

[147] Aprile E. et al. Light Dark Matter Search with Ionization Signals in XENON1T. arXiv:1907.11485 [astro-ph, physics:hep-ex] (2019).

[148] Cowan G. Statistical Data Analysis. Clarendon Press, Oxford (1998).

[149] Cowan G. et al. Asymptotic formulae for likelihood-based tests of new physics. The European Physical Journal C, 71(2):1–19 (2011).

[150] Akerib D.S. et al. Extending light WIMP searches to single scintillation photons in LUX. Physical Review D, 101(4):042001 (2020).

[151] Agnese R. et al. New Results from the Search for Low-Mass Weakly Interacting Massive Particles with the CDMS Low Ionization Threshold Experiment. Physical Review Letters, 116(7):071301 (2016).

[152] Angloher G. et al. Results on light dark matter particles with a low-threshold CRESST-II detector. The European Physical Journal C, 76(1):25 (2016).