Dark Matter Searches with the LUX and LZ Experiments

By

SALLY SHAW

Department of Physics and Astronomy
UNIVERSITY COLLEGE LONDON

A dissertation submitted to University College London in
accordance with the requirements of the degree of DOCTOR
OF PHILOSOPHY in the Faculty of Physics.

NOVEMBER 2016

ABSTRACT

The fundamental nature of our universe is still mostly unknown. Dark matter, known to make up 84% of the mass in the universe, is widely recognised as one of the biggest mysteries in modern science. A promising candidate is the Weakly Interacting Massive Particle, or WIMP; these enigmatic particles have so far evaded detection by anything other than their gravitational influence, but they are an elegant solution that may still be hiding in unexplored phase space.

WIMPs may be detected through elastic nuclear scattering in low-background experiments located in deep underground laboratories. The LUX dark matter experiment operates a time projection chamber with a xenon target, and has been world-leading for 3 years. LUX pioneered calibration techniques that have allowed great improvements in the sensitivity to WIMPs, developed novel signal identification algorithms towards rare-event selection, and demonstrated low energy nuclear recoil efficiency for WIMPs to unprecedented levels for noble gas targets. LUX has set the most-stringent constraints on WIMPs to-date with a a minimum sensitivity of 2.2×10^{-46} cm^2 at a WIMP mass of 50 GeV/c^2.

The LUX experiment has been completed, but the search will be taken up by its successor LUX-ZEPLIN (LZ), now under construction. Building on the technologies and techniques developed for LUX, LZ will probe theoretically well-motivated regions of unexplored electroweak parameter space to reach a sensitivity at 2×10^{-48} at 50 GeV/c^2, 100 times greater than LUX. Detailed Monte Carlo simulations have been conducted for LZ to inform and finalise design, and develop the background model, against which any potential signal will be evaluated for a first definitive discovery.

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

SIGNED: .... DATE: ....

The work contained in this thesis has been used in the following papers:

- D. S. Akerib *et al* (LUX), *Tritium calibration of the LUX dark matter experiment*, Phys. Rev. D. **93**, 072009 (2016)

- D. S. Akerib *et al* (LUX), *Improved Limits on Scattering of Weakly Interacting Massive Particles from Reanalysis of 2013 LUX data*, Phys. Rev. Lett. **116**, 161301 (2016)

- D. S. Akerib *et al* (LUX), *Low-energy (0.7-74 keV) nuclear recoil calibration of the LUX dark matter experiment using D-D neutron scattering kinematics,* arXiv:1608.05381, 161301 (2016)

- D. S. Akerib *et al* (LUX), *Results from a search for dark matter in the complete LUX exposure,* arXiv:1608.07648, accepted for publication in Phys. Rev. Lett. (2016)

- D. S. Akerib *et al* (LUX), *Signal yields, energy resolution, and recombination fluctuations in liquid xenon*, arXiv:1610.02076 (2016)

- D. S. Akerib *et al* (LZ), *LZ Technical Design Report*, in preparation

- S. Shaw *et al* (LZ), *Identification of Radiopure Titanium for the LZ Experiment,* in preparation

INTRODUCTION

We do not know or understand what makes up over 95% of our universe. This is a powerful and astonishing fact of modern physics: everything we know, all our physical matter and radiation, only accounts for a tiny fraction of the cosmos. Dark matter has been an enigma for over 30 years. It is one of the greatest mysteries of modern science; an unidentified substance only observable by its gravitational pull. It emits no light and yet its influence is profound; dark matter is responsible for the formation of galaxies and galaxy clusters. It is not an exaggeration to say our very existence depends on this unknown substance. Discovery of the nature of dark matter is recognised as one of the greatest contemporary challenges in science.

The mystery of dark matter dates back to the 1930s. Since then, we have found more and more evidence pointing to the existence of this massive, unreactive substance. This chapter will tell the story of dark matter, from its discovery and the leading theories explaining its nature to the current ongoing efforts to detect it.

1.1 Evidence for Dark Matter

1.1.1 Galaxy Rotation Curves

In the 1930s, radio-astronomers were taking measurements of galaxy rotation velocities, with Zwicky [1] and Oort [2] noticing the unusual property that the luminous rotational velocity of galaxies and nebulae was constant with increasing distance from the centre of the system, see figure 1.1. However, standard Newtonian dynamics and the virial theorem predict a fall off with

radius r:

$$v(r) = \sqrt{\frac{GM(r)}{r}} \tag{1.1}$$

where G is the gravitational constant, 6.674×10^{-11} nm^2/kg^2, and $M(r)$ is the mass contained within radius r. This observation, however, did not gain much interest until the 1970s, when astronomers, beginning with Rubin and Ford [3], were noticing that this effect was consistent between all galaxies and galaxy clusters. The proposed explanation was invisible, non-radiative matter; a dark, massive substance, existing in a halo around the galaxy, thus increasing its mass and allowing matter at high radii to rotate at the observed velocities, whilst preserving Newton's laws. There have been proposals that rather than the existence of a 'missing' matter, Newtonian
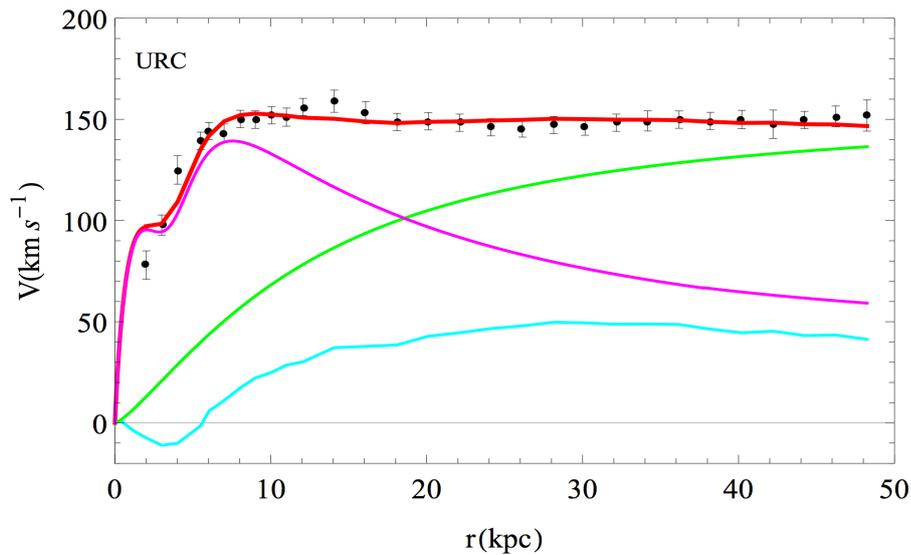
Figure 1.1: Velocity data from the spiral galaxy NGC 3198 (circles) modeled in red with several components: halo cored component (thick green line), the stellar disk (magenta line) and the H_I (hydrogen 21cm line) disk (azure line) [4].

physics requires a modification, specifically at low accelerations. These theories are known as Modified Newtonian Dynamics (MOND), and the original theory was proposed by Milgrom in 1983 [5]. MOND can satisfactorily explain some galactic rotation curves, however, it has been shown that there are issues with temperature profiles [6] and stability [7] of the galactic disc. Furthermore, it can not sufficiently explain all of the other observations, detailed in the next few sections that provide evidence for dark matter. There are theories and tests for MOND still in development [8], but the general scientific consensus is in strong favour of dark matter.

1.1.2 The Cosmic Microwave Background

The Cosmic Microwave Background (CMB) is often called 'relic' radiation. It was emitted from the surface of last scattering just 380,000 years after the Big Bang - the moment of recombination where atoms formed and the universe became transparent to photons—now stretched by inflation to a thermal radiation with a temperature of 2.7 K. The CMB is roughly homogeneous and isotropic, but contains tiny fluctuations (or anisotropies) at around 1 part in 100,000. These anisotropies can be mapped at different angular scales and the corresponding power spectrum is shown in figure 1.2. The angular scale is parametrised as the multipole moment of a spherical harmonic, ℓ, which determines the wavelength, $\lambda = 180°/\ell$ of the mode on the sphere of the CMB. Larger values of ℓ refer to smaller regions of the sky. Therefore, the power spectrum tells us how correlated the temperature at two points of the sky are on different angular scales. Information on the density of matter and dark matter in the universe can be extracted from the peaks; in particular the third peak is sensitive to the density ratio of dark matter to radiation. Notice that the peaks become damped as ℓ increases; photon diffusion washes out the initial fluctuations from inflation. The third peak, however. is boosted relative to the rest, demonstrating the domination of dark matter in the plasma before the time of recombination.
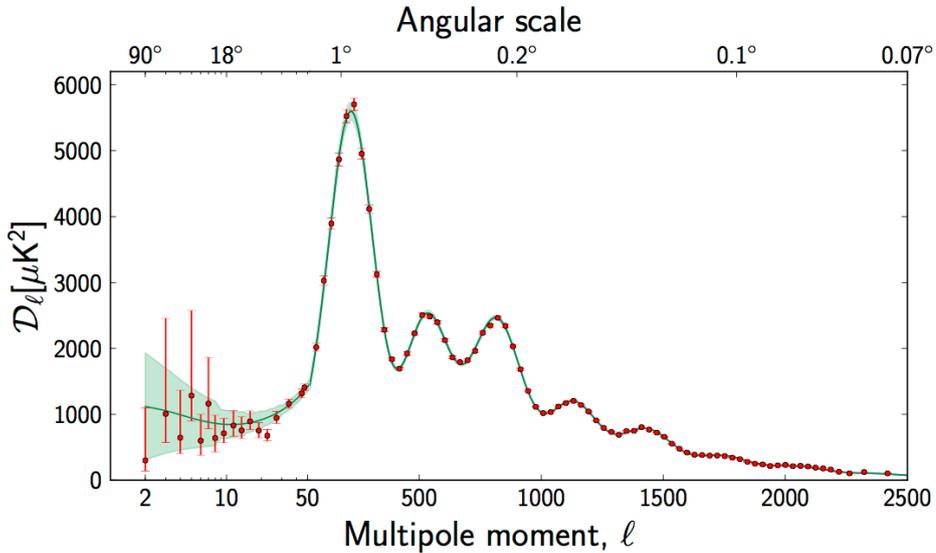
Figure 1.2: Power spectrum of the CMB temperature anisotropy in terms of the angular scale. The first peak provides information on the curvature of the universe, the second provides the baryon density, Ω_B and the third provides the dark matter density Ω_{DM} [9].

It is standard in cosmology to quantify densities using dimensionless ratios known as the present-day density parameters, Ω_x, where x can be several species:

$$\Omega_x = \frac{8\pi G \rho_x}{3H^2} = \frac{\rho_x}{\rho_c} \tag{1.2}$$

where $\rho_c = 3H^2/8\pi G$ is the critical density for which the spatial geometry of the universe is flat, and H is the Hubble parameter. Using this formalism, the Planck satellite has measured the density of matter in the universe ($x = m$) from the CMB as $\Omega_m = 0.3089 \pm 0.0062$ [10].

1.1.3 Baryonic Acoustic Oscillations

The anisotropies in the CMB, presented in figure 1.2, are caused by baryonic acoustic oscillations—density fluctuations in baryonic matter. In the very early universe regions of higher density attracted matter gravitationally. Any dark matter in these regions continued to collapse, whilst as the photon-baryon plasma density increased, it heated up and started to expand. The expansion caused a counteractive cooling, and the photons and baryons fell back inwards until thermal pressure overcame gravity and the process reversed. These oscillations continued until recombination, when photons streamed away leaving behind a shell of baryons and unperturbed dark matter at the centre of the over-density. Baryonic matter, now cool, fell back inwards, and these structures are thought to be the seeds of galaxy formation. The process led to the existence of hotter regions where the matter had contracted, and cooler where the photons had streamed away, measured as anisotropies in the CMB.
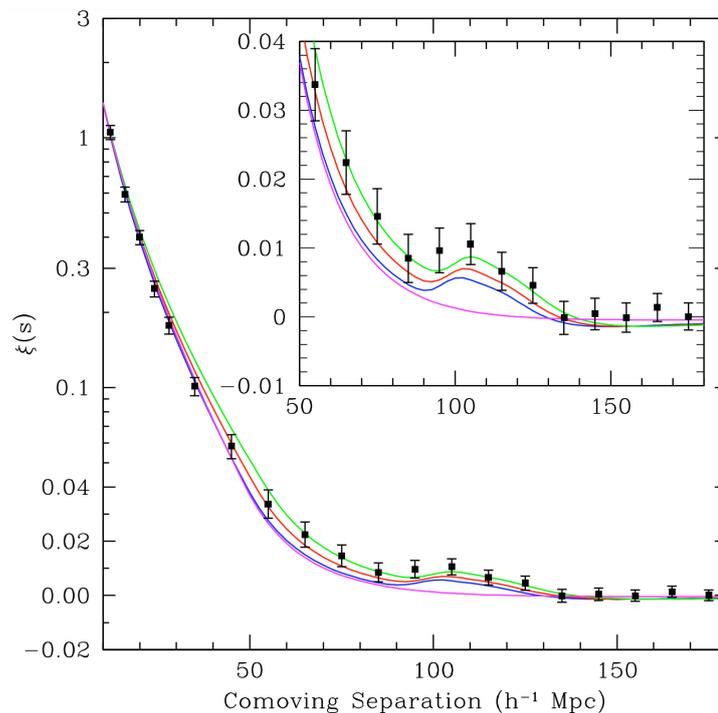
Figure 1.3: The correlation function ξ as a function of the comoving galaxy separation (s) as measured by the Sloan Digital Sky Survey using a sample of 46,748 luminous red galaxies [11]. ξ describes the probability of finding a pair of galaxies within the distance s of each other. A statically significant bump appears at ~100 Mpc. The curves show predictions from different models; green shows $\Omega_m h^2 = 0.12$, red 0.13 and blue 0.14, and magenta is a pure CDM model and shows no peak feature.

The CMB spectrum allows the determination of a length scale for baryons in the early universe of ~150 megaparsec (Mpc). The effect is detectable as a ~1% preference for pairs of galaxies to be separated by this length, called the BAO feature, see figure 1.3. This has been measured in two large galaxy redshift surveys, the Sloan Digital Sky Survey and the 2dF Galaxy Redshift Survey [12]. Combining the CMB observations with these BAO measurements allows a precise estimate of the Hubble constant and the baryonic matter density in the universe, $\Omega_B = 0.0486 \pm 0.0010$ [10]. Combining this with Ω_m leaves a dark matter component of $\Omega_{DM} = 0.2589 \pm 0.0057$, or about 26% of the energy density of the universe, and >83% of the matter density.

1.1.4 Big Bang Nucleosynthesis

Big Bang Nucleosynthesis (BBN) is the production of nuclei heavier than hydrogen in the early universe. Taking place 10 seconds to 20 minutes after the Big Bang, where temperatures were 10 MeV to 100 keV, the main processes of BBN are:

$$p + n \rightarrow\, ^2\mathrm{D} + \gamma \tag{1.3}$$

$$^2\mathrm{D} + p \rightarrow\, ^3\mathrm{He} + \gamma; \quad ^2\mathrm{D} +\, ^2\mathrm{D} \rightarrow\, ^3\mathrm{T} + p \tag{1.4}$$

$$^3T + D \rightarrow\, ^4\mathrm{He} + n \tag{1.5}$$

$$^3\mathrm{He} +\, ^2\mathrm{D} \rightarrow\, ^4\mathrm{He} + p \tag{1.6}$$

$$^4\mathrm{He} +\, ^3\mathrm{T} \rightarrow\, ^7\mathrm{Li} + \gamma; \quad ^4\mathrm{He} +\, ^3\mathrm{He} \rightarrow\, ^7\mathrm{Be} \tag{1.7}$$

$$^7\mathrm{Be} + n \rightarrow\, ^7\mathrm{Li} + p \tag{1.8}$$
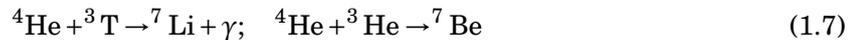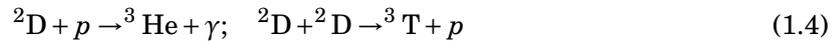
It can be calculated that BBN created almost all of the helium in the universe, and that heavier elements were made in decreasing quantities.

BBN provided further strengthening of the dark matter theory when in the 1970s, calculations of the deuterium concentration in the universe were found to be too high to be consistent with the usual BBN models, which presumed most of the universe consisted of baryons. Within this scenario, most deuterium would have been turned into ^{4}He, contradictory to measurements. However, the observed concentrations of deuterium can be matched to the model when assuming a smaller baryon density, and including dark matter as the remaining mass.

1.1.5 Gravitational Lensing

Gravitational lensing is the phenomenon by which light from a bright, extremely distant source is bent by the mass of an astrophysical object along its path. This can result in a replication of the source image or a distortion; an example is shown in figure 1.4. The amount of distortion is directly related to the mass of the lensing structure. This allows accurate measurements of the
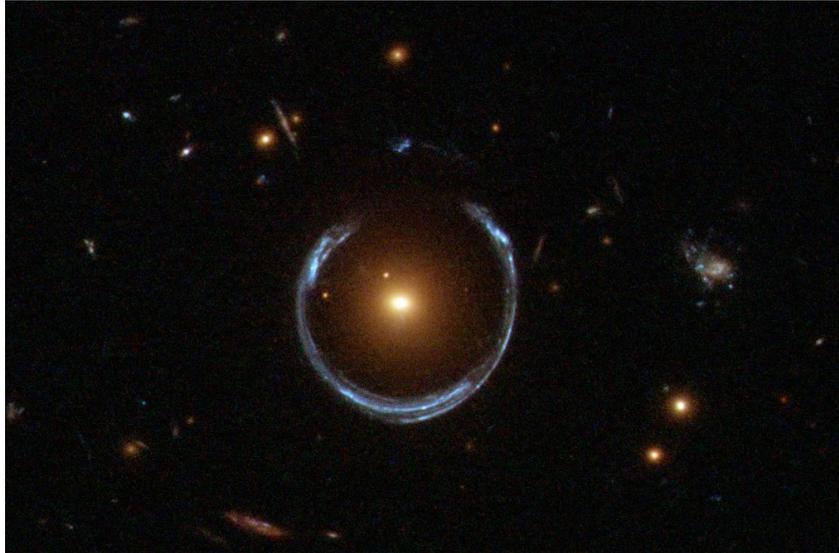
Figure 1.4: Example of an 'Einstein Ring' created by gravitational lensing of a distant blue galaxy by the red galaxy LRG 3-757. The image of the distant galaxy has been distorted into an almost perfect circle [ESA/Hubble.]

mass of distant galaxies and galaxy clusters, which confirms their mass obtained from rotation curves, providing support for the dark matter theory.

Mapping of the mass within a lensing structure is possible, and provides very convincing proof of a non-interacting dark matter component. A famous example of this is the Bullet Cluster, seen in figure 1.5. Here, two galaxy clusters have collided. False colour has been added to the image to map out the mass measured from x-ray emissions in pink and from gravitational lensing in blue. The gravitational lensing measurements show two clear 'cores' of mass, thought to be dark matter, that have passed through each other with very little interaction. The normal matter, in the form of x-ray emitting gas, has interacted and clustered together in the centre. These measurements can be used to place constraints on the self-interaction of dark matter particles [13].

1.1.6 Structure Formation

Observations of the large-scale structure of the universe today on the level of galaxy clusters show a universe made of filaments and voids. Dark matter is able to explain this structure and its formation from anisotropies in the CMB, if it is assumed to be 'cold'. The temperature of dark matter refers to its velocity (and correspondingly its mass) at the time of 'freeze-out', which was the moment in the early universe where annihilation of dark matter stopped and the relic density was set (see section 1.2.2 for more details). Cold dark matter (CDM) refers to a non-relativistic, heavy particle, hot dark matter to a ultra-relativistic, extremely light particle, and warm is somewhere in between. Assuming that the universe began as isotropic and homogeneous and that dark matter seeded the formation of over-dense regions, large scale computer simulations can be run to simulate the formation of structure in the early universe. These can be done assuming

Figure 1.5: The Bullet Cluster. False colour shows the mass distribution from gravitational lensing (blue) mapped from background galaxies, and from x-ray emitting gas (pink). This demonstrates the non-interacting nature of dark matter [14, 15].

different temperatures of dark matter and different initial conditions; it is found that cold dark matter produces universes very much like the one we observe today. Warmer dark matter washes out substructure because of its high velocity. This has led to the adoption of the ΛCDM cosmology. Λ is the cosmological constant, associated with dark energy, and ΛCDM is the leading model for the universe today. However, whilst a simulation run with a CDM component reproduces the filamentary structure of the universe well, there are still three main problems. The first of these is known as the missing satellites problem; simulations predict an over-abundance of dwarf galaxies that are not observed, with a difference on the order of thousands. There are potential solutions to this problem; the small DM halos do exist but most have not attracted enough enough baryonic matter to create a visible dwarf galaxy [16], that supernovae could have blown away much of the matter inside them [17], and also that some of the dwarf galaxies may have been merged into or tidally stripped apart by larger galaxies [18]. The second problem is known as 'too big to fail'; the dwarf galaxies predicted by simulations are much more massive than those that are observed [19]. Technically these should have been 'too big to fail' and we should be able to observe them today, but we do not. Finally, the third problem is known as the core-cusp problem; simulations predict a sharp rise in dark matter density in the centre of galaxies (a cusp), whilst observations suggest a flatter profile (a core) [20]. These problems suggest that either the ΛCDM model is not a satisfactory model of the universe, or that our computer simulations, which are less reliable on small scales, are not yet good enough to adequately predict small-scale structures. Inclusion of an additional warm dark matter component can solve some of these issues and this is an ongoing field of study in astrophysics.

1.2 WIMP Dark Matter

1.2.1 Properties of Dark Matter

From astrophysical observations, several key assumptions can be made about dark matter:

- It is a neutral or millicharged particle
- It is stable, or has a lifetime longer than the age of the universe
- It does not interact with photons by emission or absorption, but annihilation into photons is possible
- It is 'cold'—non-relativistic at the time of freeze-out
- It is non-baryonic

The Standard Model (SM) of particle physics contains no suitable candidates for dark matter. Whilst neutrinos share some of the above properties, their abundance is well known, and they constitute a small amount of hot dark matter: $\Omega_\nu \leq 0.003$. Therefore, the problem of dark matter requires physics beyond the Standard Model.

1.2.2 Motivation for WIMPs

The motivation for the dark matter candidate known as the Weakly Interacting Massive Particle, or WIMP, lies with the relic abundance of dark matter. In the early universe after inflation, all particles exist in a hot 'soup'. Dark matter particles will undergo self-annihilation within this soup; the rate of this process decreases as the universe expands because of the resulting drop in particle density. The corresponding Boltzmann equation for a particle species of number density n is [21]:

$$\frac{dn}{dt} + 3Hn = -\langle\sigma v\rangle(n^2 - n_{eq}^2) \tag{1.9}$$

where H is the Hubble parameter, $H = \dot{a}/a$, with a as the scale factor of the universe, σ is the particle interaction cross section, v is velocity and n_{eq} is the equilibrium number density. Particles 'freeze-out', i.e. stop annihilating and remain at what is known as the relic abundance, once $H \sim \langle\sigma v\rangle n$ - i.e. once the universe has expanded sufficiently enough that the particle density is too low for self-annihilation. Knowing the relic abundance today allows an estimate of σ, although without knowing the WIMP mass it cannot be derived directly. Defining the subscript f as referring to freeze-out, and denoting the WIMP as χ, the Boltzmann equation can be solved and is found to give [21]:

$$n_f \sim (m_\chi T_f)^{3/2} e^{-m_\chi/T_f} \sim \frac{T_f^2}{M_{Pl}\langle\sigma v\rangle} \tag{1.10}$$

where m_χ is the WIMP mass, T is temperature and M_{Pl} is the Planck mass, defined as $M_{Pl} = \sqrt{\bar{h}c/G}$. Notice the ratio $x_f = m_\chi/T_f$ in the exponential; this is found to be roughly constant at ~ 20 for all WIMP masses. The relic density then becomes [21]:

$$\Omega_\chi = \frac{m_\chi n_0}{\rho_c} = \frac{m_\chi T_0^3}{\rho_c}\frac{n_0}{T_0^3} \sim \frac{m_\chi T_0^3}{\rho_c}\frac{n_f}{T_f^3} \sim \frac{x_f T_0^3}{\rho_c M_{Pl}}\frac{1}{\langle\sigma v\rangle} \tag{1.11}$$

where $\rho_c = 3H^2/8\pi G$ is the critical density of the universe (for flat space-time) and the subscript 0 represents the universe at the current time. This is only valid if $m_\chi > T_f$—so the particle is non-relativistic at freeze-out. Figure 1.6 shows the evolution of the WIMP density with time, and how altering the cross section affects the relic density. The point of freeze-out is evident from where the density becomes constant.

We find that: $\langle \sigma v \rangle \sim G_F^2 m_\chi^2 \sim 10^{-10} \text{GeV}^{-4} m_\chi^2$ for $1 < m_\chi < 1000$ GeV (where G_F is the Fermi constant, 1.16637×10^{-5} GeV^{-2}) which is of the order of the scale of the weak interaction. This has been dubbed the 'WIMP miracle'; the fact that the measured relic density of dark matter requires a weakly interacting particle means that WIMPs are theoretically well-motivated. Furthermore, several WIMP candidates arise naturally in SM extensions such as supersymmetry (SUSY).
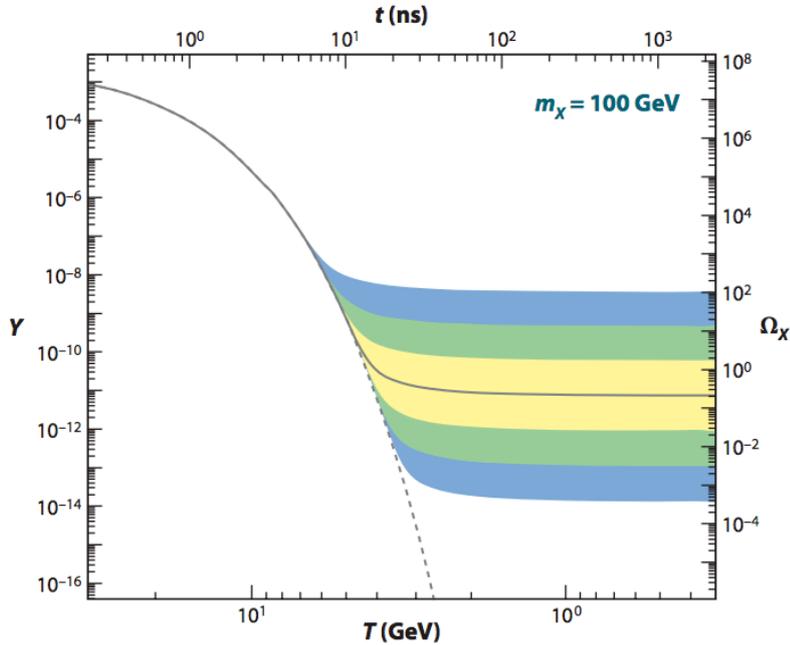
Figure 1.6: Here the WIMP yield $Y = n/s$ where s is the entropy density is shown as a function of temperature T and equivalently time, demonstrating how the density becomes constant at the relic density, known as freeze-out. Yellow, green and blue shaded regions represent the effect of altering the cross section by factors of 10, 10^2 and 10^3 respectively [22].

1.2.3 The WIMP Halo

Using the galactic rotation measurements as a basis, the WIMP distribution is usually assumed to be an isothermal, spherical halo around the galaxy. Several different models can be used for the dark matter density profile; a commonly used simple one is the pseudo-isothermal halo:

$$\rho(r) = \frac{\rho_0}{1 + \left(\dfrac{r}{r_c}\right)^2} \tag{1.12}$$

9

where ρ_0 is the finite central density and r_c is the core radius. Whilst this model provides a good fit to rotation curve data, the enclosed mass diverges as the radius tends to infinity, which is unrealistic.

A more advanced model, the NFW (Navarro-Frenk-White) halo [23], was predicted using numerical simulations. This model uses the critical density of the universe, and is described by:

$$\rho(r) = \frac{\rho_c \delta_c}{\left(\dfrac{r}{r_s}\right)\left(1 + \dfrac{r}{r_s}\right)^2} \tag{1.13}$$

where δ_c is a dimensionless density parameter characteristic for each galaxy, which is proportional to the density of the universe at the time the galaxy formed. The NFW halo works for halo masses spanning four orders of magnitude, from small galaxies to galaxy clusters. However, simulations with this model lead to the unsolved core-cusp problem mentioned in section 1.1.6.

The WIMPs within the halo have a velocity distribution described by a Gaussian distribution:

$$f(\bar{v}) = \frac{1}{(2\pi\sigma_v)^{3/2}} e^{\frac{-|v|^2}{2\sigma_v^2}} \tag{1.14}$$

where σ_v is the 1-dimensional velocity dispersion, with the subscript v present to separate it from the cross section.

1.2.4 WIMP-Nucleon Scattering

As WIMPs are by nature weakly interacting, they will undergo scattering with atomic nuclei. This is a potential detection route, see section 1.3.3 on direct detection. Here we will derive the quantities relevant to such a process.

The cross section will be dependent on the WIMP-quark interaction strength, and furthermore the distribution of quarks within a nucleus. There are three main steps to the cross section calculation; first, the interaction of WIMPs with quarks and gluons, secondly, the translation to nucleons, and finally to the whole nucleus.

There are two types of possible interaction depending on the mediator; axial-vector (spin-dependent) and scalar (spin-independent), which will be discussed separately below. The assumption is made that the WIMP is a supersymmetric neutralino, following Jungman, Kamionkowski and Griest [24].

1.2.4.1 Axial-Vector Interaction

The Lagrangian to describe the WIMP-nucleon axial vector interaction is given by [24]:

$$\mathscr{L}_A = d_q \bar{\chi} \gamma^\mu \gamma_5 \chi \bar{q} \gamma_\mu \gamma_5 q \tag{1.15}$$

where d_q is a coupling strength, which differs depending on the flavour of quark. $\bar{\chi}$, $\chi/\bar{q}, q$ are the spinors of the WIMP and quark respectively, and γ^μ/γ_5 have their standard form as Dirac

γ-matrices. The interaction can be mediated by either the Z boson or a supersymmetric quark, or squark. To extend this to a nucleon and evaluate the matrix element, the spin content of the nucleon n is required in the Lagrangian [24]:

$$\mathscr{L}_A = \bar{\chi}\gamma^\mu\gamma_5\chi\bar{n}s_\mu n \sum_{q=u,d,s} 2d_q\Delta q^{(n)} \tag{1.16}$$

where s_μ is the spin of the nucleon, and $\Delta q^{(n)}$ are values obtained using experimental data on lepton-proton scattering. Within the sum, advantage has been taken of an isospin rotation, $u \leftrightarrow d$, between protons and neutrons.

Next, we must consider the whole nucleus. At zero momentum transfer, the average spins of the neutrons and protons can be calculated to determine the spin of the nuclear state, but at non-zero momentum transfer of $|\mathbf{q}|$ there is a nuclear form factor to be considered. It can be shown that [24]:

$$\frac{d\sigma}{d|\mathbf{q}|} = \frac{8}{\pi v^2}\Lambda^2 G_F^2 J(J+1)\frac{S(|\mathbf{q}|)}{S(0)} \tag{1.17}$$

where v is the WIMP velocity relative to the target, J is the total angular momentum of the nucleus and $S(Q)$ (defining $Q = |\mathbf{q}|$ to avoid confusion with notation for quarks) is given by [24]:

$$S(Q) = a_0^2 S_{00}(Q) + a_1^2 S_{11}(Q) + a_0 a_1 S_{01}(Q) \tag{1.18}$$

where a_0 (isoscalar) and a_1 (isovector) are defined as:

$$a_p = \sum_{q=u,d,s}\frac{d_q}{\sqrt{2}G_F}\Delta q^{(p)}, \quad a_n = \sum_{q=u,d,s}\frac{d_q}{\sqrt{2}G_F}\Delta q^{(n)} \tag{1.19}$$

$$a_0 = a_p + a_n, \quad a_1 = a_p - a_n \tag{1.20}$$

and Λ in equation 1.17 is:

$$\Lambda = \frac{[a_p\langle S_p\rangle + a_n\langle S_n\rangle]}{J} \tag{1.21}$$

The form factors $S_{ij}(\mathrm{q})$ and the expectation values of the proton and neutron spin content $\langle S_p\rangle$ and $\langle S_n\rangle$ can be calculated by detailed nuclear calculations, which are beyond the scope of this work. As an example, for ^{131}Xe, it has been calculated that $\langle S_p\rangle = 0.0$ and $\langle S_n\rangle = -0.166$. $\langle S_p\rangle$ is null as expected as ^{131}Xe has only one unpaired neutron, which carries all of the spin.

Thus, the cross section becomes [24]:

$$\frac{d\sigma}{dQ^2} = G_F^2\frac{C}{v^2}F^2(Q) = \frac{\sigma_0}{4\mu_{\chi N}^2 v^2}F^2(Q) \tag{1.22}$$

where $\mu_{\chi N}$ is the reduced mass of the WIMP and nucleons, $\mu_{\chi N} = m_N m_\chi/(m_N + m_\chi)$, C is a model-dependent dimensionless number, and F is the normalised ($F(0) = 1$) form factor. For this interaction, also known as spin-dependent scattering, we find [24]:

$$F^2(Q) = \frac{32}{\pi}G_F^2\mu_{\chi N}\Lambda^2 J(J+1) \tag{1.23}$$

The cross section at zero momentum transfer is given by [24]:

$$\sigma_{0SD} = \int_0^{4\mu_{\chi N}^2 v^2} \frac{d\sigma}{dq^2} = 4G_F^2 \mu_{\chi N}^2 C \tag{1.24}$$

where SD will be used from now on to refer to spin-dependent scattering. Separating the zero momentum transfer cross section is useful when considering different targets, where the form factor for spin-dependent scattering is nucleus-dependent.

1.2.4.2 Scalar Interaction

The scalar interaction between WIMPs and quarks in the neutralino model arises from the exchange of a Higgs or a squark. In this model there is a further addition in the form of WIMPs interacting with gluons via heavy quark loops. Here we will just consider the scalar interaction of neutralinos with quarks, although Jungman, Kamionkowski and Griest explicitly calculate the quark-squark loops. The Lagrangian takes the form [24]:

$$\mathscr{L}_S = f_q \bar{\chi} \chi \bar{q} q \tag{1.25}$$

where f_q is the coupling to quarks. By summing over the quark couplings and taking into account the heavy quark loops, couplings to protons and neutrons f_p and f_n can be calculated It can then be shown that [24]:

$$\frac{d\sigma}{dQ^2} = G_F^2 \frac{C_{scalar}}{v^2} F^2(Q) = \frac{\sigma_{0SI}}{4\mu_{\chi N}} F^2(Q) = \frac{1}{\pi v^2}[Zf_p + (A-Z)f_n]^2 F^2(Q) \tag{1.26}$$

where SI will be used from now on to refer to spin-independent, Z and A have their usual meanings as number of protons and atomic mass and $F(Q)$ is a nuclear form factor; the most commonly used is the exponential form [24]:

$$F(Q) = e^{-\frac{Q}{2Q_0}} \tag{1.27}$$

where Q_0 is the coherence energy given by:

$$Q_0 = \frac{1.5}{m_N R_0^2} \tag{1.28}$$

and R_0 is the radius of the nucleus:

$$R_0 = 10^{-13} \text{cm}[0.3 + 0.91(m_N/\text{GeV})^{1/3}] \tag{1.29}$$

Note that this form factor will fall-off more rapidly for heavier nuclei as m_N increases, as well as with an increase in the energy transferred.

As before, we define a zero-momentum cross section [24]:

$$\sigma_{0SI} = \int_0^{4\mu_{\chi N}^2 v^2} \frac{d\sigma(Q=0)}{dQ^2} = \frac{4\mu_{\chi N}}{\pi}[Zf_p + (A-Z)f_n]^2 \tag{1.30}$$

Thus, SI scattering has an A^2 dependence that enhances the event rate for high mass targets.

Originally it was expected that the axial-vector coupling was the only interaction that occurred between neutralino WIMPs and nuclei. However, the scalar contribution becomes significant due to heavy quarks and is additionally enhanced if the Higgs boson is light. With the discovery that the Higgs mass is indeed relatively light (125 GeV) and the top quark is known to be heavy, it seems likely the scalar contribution may be dominant.

Whilst a vector interaction is also theoretically able to produce spin-independent scattering, the expected cross section is in excess of experimental limits for most couplings [25]. We can choose a coupling that allows agreement with current constraints, but for a WIMP with a mass in the range 10 to 1000 GeV annihilating via mainly vector interactions, the relic density is calculated as a factor up to $10^5\times$ higher than observed [25]. For this reason, we focus on the scalar interaction.

1.2.4.3 Differential Event Rate in a Detector

For WIMP detection through WIMP-nucleon scattering, we are interested in the rate of events expected. Combining the cross sections calculated in sections 1.2.4.1 and 1.2.4.2 and the velocity distribution calculated in section 1.2.3, the differential event rate in a target can be derived to be [21]:

$$\frac{dN}{dE_R} = \epsilon \frac{\rho}{2m_\chi \mu_{\chi N}} \sigma_0 F^2(E_R) \int_{v_{min}} \frac{f(\bar{v})}{v} d^3v \tag{1.31}$$

where $\epsilon = \Delta t m_T$ is the experimental exposure (time of exposure multiplied by target mass, m_T), ρ is the local dark matter density and v_{min} is defined by [21]:

$$v_{min} = \sqrt{\frac{m_\chi E_{thr}}{2\mu_{\chi N}^2}} \tag{1.32}$$

This is the smallest velocity possible that can give a recoil of an energy at the detection threshold of the detector E_{thr}. The recoil energy the nucleus carries after an interaction is given by kinematics as [21]:

$$E_R = \frac{1}{2}m_\chi v^2 \frac{4m_\chi m_N}{(m_\chi + m_N)^2} \frac{1 + cos\theta}{2} \tag{1.33}$$

Note that in equation 1.31, there is a useful separation of physics. The particle and nuclear physics is contained within $\sigma = \sigma_0 F^2(E_R)$, the astrophysics within ρ and the velocity distribution integral, and the experimental factor within ϵ. From this differential rate, the number of expected events within an experimental exposure of ϵ and an energy range of E_{thr} to E_{max} is [21]:

$$N = \int_{E_{thr}}^{E_{max}} \frac{dN}{dE_R} dE_R \tag{1.34}$$

Therefore, it is possible to probe both the mass of the WIMP m_χ and its cross section σ through experiments. Figure 1.7 demonstrates the result of this calculation for four different elements commonly used in dark matter detection, for recoil energies of 0 - 100 keV. Notice the exponential fall-off with energy, which results from the form factor (equation 1.27).
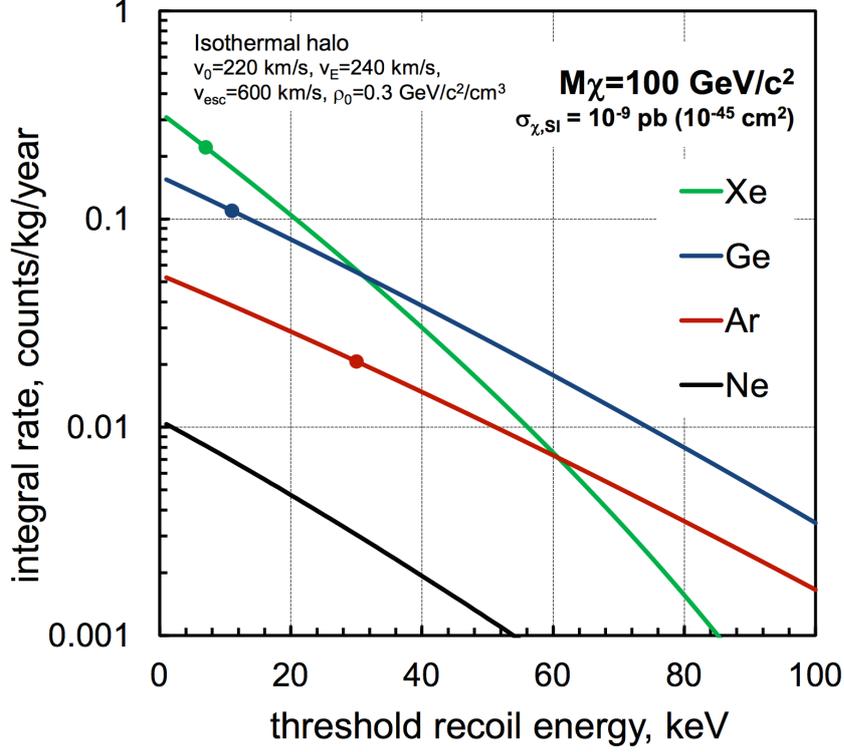
Figure 1.7: Integral rate of WIMP-nucleon scattering in counts per kg per year against recoil energy shown for a 100 GeV WIMP with $\sigma = 10^{-45}$ cm^{-2} and four different elements, Xe ($A = 131$), Ge ($A = 73$), Ar ($A = 40$) and Ne ($A = 20$), assuming perfect energy resolution and the isothermal halo model. The circles mark typical thresholds for each technology [26].

1.2.4.4 Annual Modulation

If we consider the isothermal spherical halo model (see section 1.2.3), it becomes clear that the motion of the Earth around the sun will produce an annual modulation in the event rate, due to an alteration in the relative WIMP velocity. The velocity of the halo relative to Earth is maximum at the start of June, and minimum at the beginning of December. The velocity of the Earth with respect to the halo can be calculated using an approximate orbital velocity of 30 km/s, an orbital velocity of the sun around the galactic centre of 220 km/s, and the angle between the Earth and the sun relative to the sun's motion, equal to 60° [21]:

$$v_E \approx v_0 \left[1 + \frac{30 \cdot \cos(60)}{220} \cos\left(\frac{2\pi(t - t_0)}{T} \right) \right] \approx v_0 \left[1 + 0.07 \cos\left(\frac{2\pi(t - t_0)}{T} \right) \right] \tag{1.35}$$

where T is the orbital period, i.e. 1 year. Therefore, the velocity modulation seen in the dark matter rest frame is about 7% of the total orbital velocity of the earth, v_0. On earth, this leads to more WIMPs above v_{min} in June, and so should cause an increase in event rate. The event rate can be separated into a time averaged rate A_0 and a modulation component of amplitude A_m [21]:

$$\frac{dN}{dE_R}(E_R, t) = A_0(E_R) + A_m(E_R) \cos\left(\frac{2\pi(t - t_0)}{T} \right) \tag{1.36}$$

The amplitude of the modulation is dependent on the energy window chosen and its proximity to the threshold energy of the detector. Seeing a modulation in event rate with a maximum in June and a minimum in December could be evidence of dark matter detection. There are some experiments that claim to have detected this signal, but this is disputed by others who failed to confirm the results. This will be further discussed in section 1.3.3.

1.3 Dark Matter Searches

There are three routes that dark matter detection experiments can take; the production of dark matter at colliders, the detection of its annihilation or decay products and the detection of a WIMP-nucleus interaction by the measurement of a nuclear recoil. The following sections give a description of each and current experiments using the technique.

1.3.1 Dark Matter Production

Dark matter production searches take place in particle colliders such as the Large Hadron Collider (LHC). Event reconstruction can be used to determine whether any energy (usually in the form of transverse momentum, p_T) went 'missing' in a collision, potentially signalling a dark matter particle leaving the detector. Neutrinos have the same signature in particle colliders, so often the searches focus on events with no leptons, and additionally a single energetic jet to tag the event. Figure 1.8 shows examples of dark matter production in proton-proton collisions.
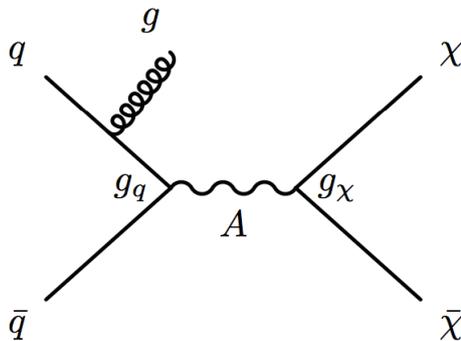
Figure 1.8: A dark matter production scenario at the LHC: pair production of WIMPs through an axial-vector mediator, A. The initial-state radiation of a gluon would be detected as a jet.

The results from these searches can be translated into the same WIMP mass vs WIMP-nucleon cross section phase space that direct detection uses. Their limits are complementary in the low WIMP mass region that direct detection is unable to probe, see figure 1.9, which shows limits placed by the CMS experiment at the LHC on dark matter properties. There is a caveat to the comparison of production and direct searches; the translation requires model assumptions, specifically on the nature of the mediator and its corresponding operator. The ATLAS

experiment has also placed bounds on the WIMP-neutron and WIMP-proton cross sections that are competitive with direct detection experiments [27].
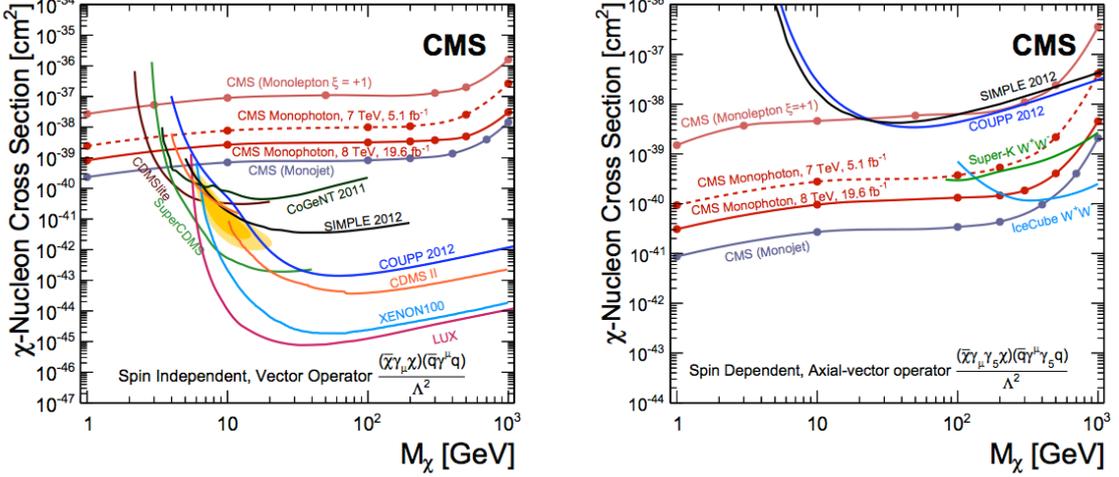
(a) Spin-independent cross section limit. (b) Spin-dependent cross section limit.

Figure 1.9: 90% confidence limits on the WIMP-nucleon spin-independent (left) and spin-dependent (right) cross sections from the CMS experiment. In both cases, the limits cover low mass regions currently unable to be reached with direct detection experiments, and for the spin-dependent limit, the limit exceeds other experiments for most WIMP masses.

1.3.2 Indirect Searches

Indirect detection experiments look for the annihilation or decay products of dark matter. They often focus on regions where the dark matter is expected to be densest, such as the galactic centre or dwarf galaxies. The basis of most indirect searches is to look for an excess of either γ-rays which arise from:

$$\chi + \chi \rightarrow \gamma + \gamma \tag{1.37}$$

where $E_\gamma = m_\chi$, or for a positron excess assuming:

$$\chi + \chi \rightarrow e^+ + e^- \tag{1.38}$$

Looking for an these excesses requires a good understanding of the γ-ray and positron spectra that that reach earth from known astrophysical sources.

Notable recent results are those of the Alpha Magnetic Spectrometer (AMS), a particle physics experiment designed to detect high energy cosmic rays. AMS-02 measured positron fraction, i.e. $e^+/(e^- + e^+)$, and found an increase in the positron fraction up to about 200 GeV, where it levels off, see figure 1.10(a). They also found an excess in the anti-proton to proton ratio that is currently unexplained, see figure 1.10(b). Many have jumped on these results in an attempt to tie them in with a dark matter model, however, for the positron excess, calculations of the

annihilation cross section assuming the positron excess is DM leads to a cross section much larger than is allowed by astrophysical constraints for all WIMP masses when using one annihilation channel. The only viable dark matter case that is allowed by γ-ray and cosmological constraints is a particle between 0.5 and 1 TeV, annihilating into four τ leptons with a branching ratio of 75% and 4 electrons with a branching ratio of 25%. More simply, a single pulsar can also satisfactorily explain the excess by itself [28]. For the anti-proton ratio, there are viable models that fit the data, for example 3 TeV supersymmetric 'wino' dark matter, but further data at high energies is needed to confirm or rule out the hypotheses [29].

(a) Positron fraction.   (b) Anti-proton proton ratio.
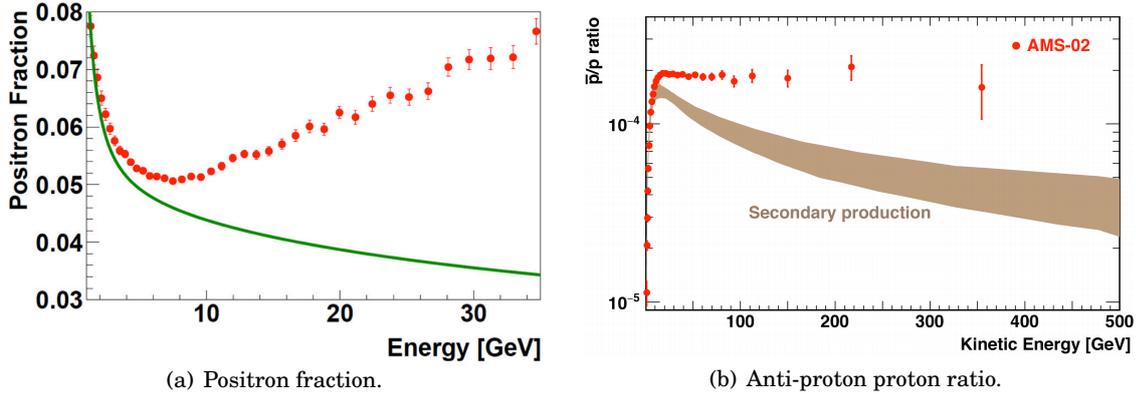
Figure 1.10: Left: AMS measurements of the positron fraction (red circles), showing an increase with energy that continues up to 200 GeV [30]. This rise is not seen in the prediction from the collision of ordinary cosmic rays (green line). Right: AMS-02 results on the anti-proton proton ratio, demonstrating that the data can not be explained by current models of secondary production [31].

Another experiment that looks for unexpected excesses is the Fermi Large Area Telescope (Fermi-LAT), which measures γ-rays from the centre of the Milky Way and has almost uniform full sky coverage from 20 MeV to 300 GeV. Fermi-LAT has placed constraints on dark matter mass and annihilation cross section using measurements of the isotropic gamma-ray background (IGRB), see figure 1.11 [32].

An additional channel is the annihilation to high energy neutrinos. Particular focus has been placed on the sun, where a flux of high energy neutrinos that are clearly not from solar processes could be an indication of dark matter annihilation. The IceCube and ANTARES experiments, both consisting of arrays of vertical strings of photomultiplier tubes viewing ice or water for muon Cherenkov light, have performed searches for dark matter annihilation neutrino products. They have placed limits on dark matter annihilation cross sections and spin-dependent WIMP-proton scattering that are competitive with direct detection searches [33, 34].

1.3.3 Direct Searches

Direct detection searches look for the interaction of our own Milky Way halo dark matter with atomic nuclei here on Earth. These searches rely on a non-zero weak interaction between dark matter and nuclei; section 1.2.4 describes the WIMP-nucleon interaction and and the differential
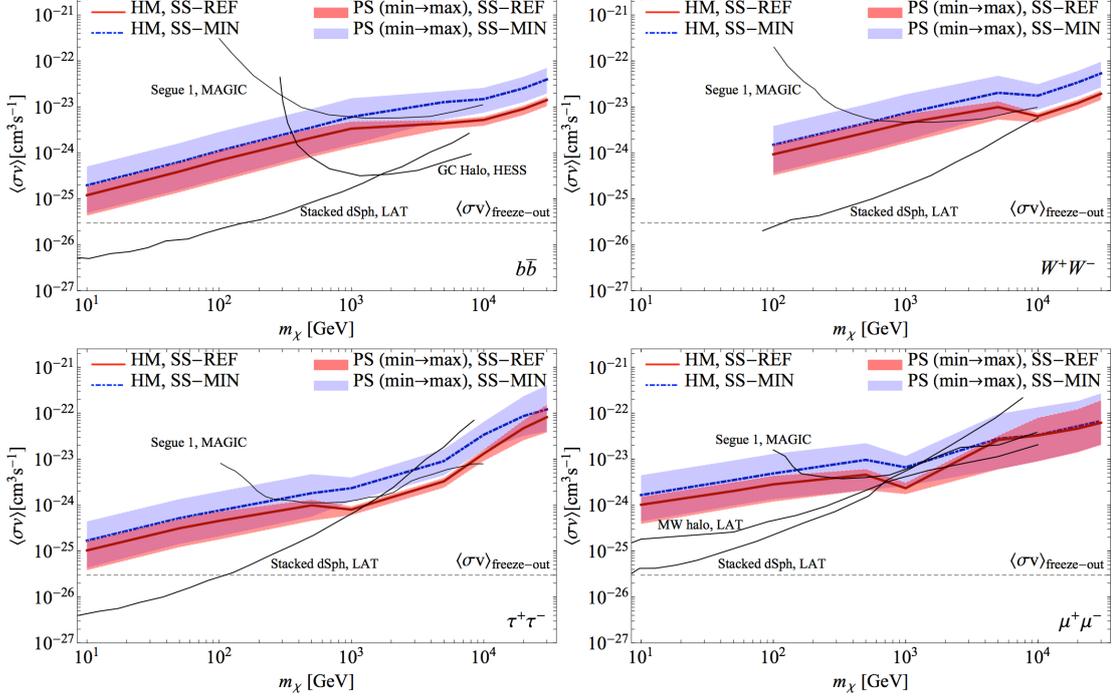
Figure 1.11: Upper limits (95% CL) on dark matter mass and annihilation cross sections through the (clockwise from top-left) $b\bar{b}$, W^+W^-, $\tau^+\tau^-$ and $\mu^+\mu^-$ channels from Fermi-LAT data [32]. Upper Red solid line shows limits obtained in a halo model scenario assuming the reference contribution from the galactic subhalo population, and the red band shows the theoretical uncertainty in the extragalactic signal. The blue dashed line with its corresponding uncertainty band refers instead to the limits obtained when the Milky Way substructure signal strength is taken to its lowest value.

rate within a detector. These experiments focus on achieving very low background event rates and so are placed deep underground in order to reduce the flux of particles from cosmic-rays.

Direct detection experiments require a form of signal production from a nuclear recoil caused by WIMP-nucleus scattering. There are three methods by which this can be done: scintillation, ionisation and heat, or alternatively the production of photons, charge and phonons. There are a plethora of direct detection experiments, some of which just use one of these channels, but many also use two, which can offer advantages in background rejection.

A consideration for direct detection targets is sensitivity to spin-dependent and spin-independent scattering. Spin-dependent scattering requires a nucleus with an unpaired nucleon, whilst for spin-independent any nucleon will do, but the heavier the better, as there is an A^2 dependence in the cross section, as explained in section 1.2.4.

1.3.3.1 Noble Gas Detectors

The noble gases xenon (Xe) and argon (Ar) are popular targets for direct dark matter searches. They are excellent scintillators, meaning they produce photons as particles pass through them. Both Xe and Ar scintillate in the vacuum ultraviolet (VUV) at 178 nm and 128 nm respectively. Furthermore, effective self-shielding, where background particles interact in the first few cm of the target, allows selection of an internal fiducial volume with a much lower background rate. The early noble gas detectors were single state—i.e. liquid only. ZEPLIN-I was the first single phase Xe dark matter detector, with the first results published in 2005 [35].

As seen in table 1.1, most progress has been made using dual-phase Xe detectors, which use both liquid and gaseous Xe to collect both light and charge. The LUX experiment, which is the focus of the next few chapters, is the world's first sub-zeptobarn ($< 1 \times 10^{-45}$ cm^2) detector [36].

Table 1.1: Past and current noble gas dark matter experiments. Active masses are given and lowest spin-independent cross sections where available are shown for comparison.

Gas	Single Phase	Lowest σ_{SI} (cm^2)	Double Phase	Lowest σ_{SI} (cm^2)
Neon	miniCLEAN (100 kg)	n/a		
	CLEAN (10-100t)	n/a		
Argon	DEAP-I (7 kg)	n/a	WArP (3.2 kg)	$\sim 1 \times 10^{-42}$ [37]
	DEAP-3600 (3600 kg)	n/a	WArP (140 kg)	n/a
	miniCLEAN (100 kg)	n/a	ArDM (1 ton)	n/a
	CLEAN (10-100t)	n/a	DarkSide-50 (46 kg)	2.0×10^{-44} [38]
Xenon	ZEPLIN-I (3.2 kg)	1.1×10^{-42} [35]	ZEPLIN-II (31 kg)	6.6×10^{-43} [39]
	XMASS (832 kg)	4.3×10^{-41} [40]	ZEPLIN-III (8 kg)	3.9×10^{-44} [41]
			XENON10 (10 kg)	4.5×10^{-44} [42]
			XENON100 (62 kg)	2×10^{-45} [43]
			XENON1T (3.5 t)	n/a
			PandaX-I (120 kg)	1.01×10^{-44} [44]
			PandaX-II (500 kg)	2.97×10^{-45} [45]
			LUX (250 kg)	0.6×10^{-45} [46]
			LZ (7 t)	n/a

1.3.3.2 Superheated Drop Detectors

Superheated Drop Detectors, or SDDs, are based on the same principles as classic bubble chamber experiments. Superheated drops of a liquid are suspended in a visco-elastic gel medium. When energetic particles pass through the medium, they cause the transition of droplets to the gas state, forming bubbles. SDDs can be operated under a low degree of superheating; a condition that make them sensitive only to recoiling nuclei, thus removing electromagnetic backgrounds. This is possible as the energy that must be deposited within a critical length for bubble nucleation is temperature dependent, so at the right degree of superheating only high linear energy transfer (LET) particles such as nuclei can deposit enough energy to trigger bubble formation. The bubbles are detected using acoustics—piezoelectric sensors surround the vessel. The main background for

this technique arises from α-particles, but as these produce a larger amplitude when compared to neutrons they can be effectively discriminated from WIMP candidate events [47].
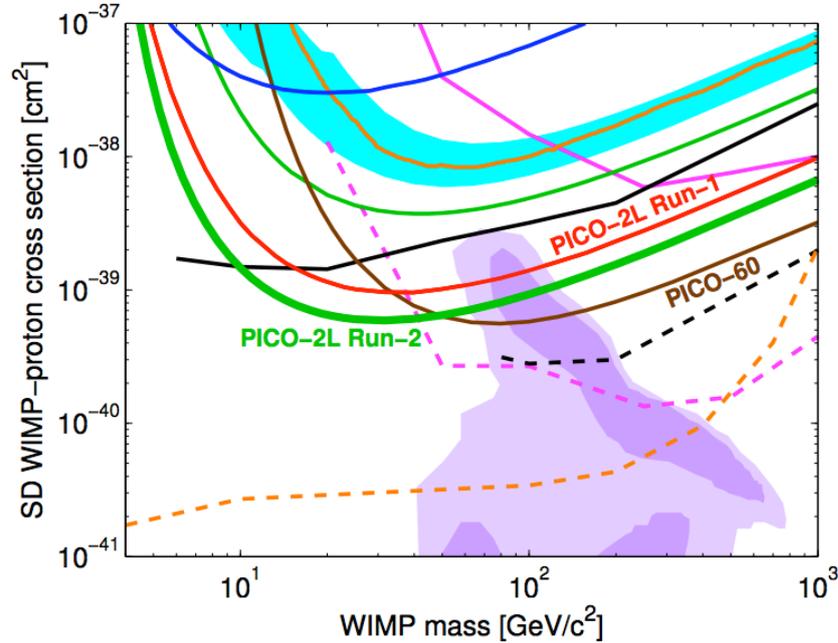
Figure 1.12: 90% C.L. limit on the SD WIMP-proton cross section PICO-2L is plotted in green [48], along with limits from PICO-60 [49] (brown), COUPP-4 (light blue region), PICASSO [50] (dark blue), SIMPLE [51] (thin green), XENON100 [52] (orange), IceCube [33] (dashed and solid pink), SuperK [53] (dashed and solid black) and CMS (dashed orange). The IceCube and SuperK (indirect detection experiments) results assume annihilation to W pairs (dashed) or b quarks (solid). The CMS limit is from a monojet search, assuming a a heavy mediator. Purple shading shows the parameter space of the constrained Minimal Supersymmetric Standard Model (cMSSM) [54].

Notable past SDDs are the Project In CAnada to Search for Supersymmetric Objects (PICASSO) at SNOLAB, the Superheated Instrument for Massive ParticLe Experiments (SIMPLE) at LSBB, and the Chicagoland Observatory for Underground Particle Physics (COUPP) at Fermilab. Currently, PICO (formed by the merging of PICASSO and COUPP), operate both a bubble chamber filled with 37 kg of CF_3I (PICO-60) and a chamber loaded with 3 kg of C_3F_8 (PICO-2L). The use of CF_3I makes PICO-60 sensitive to both SD and SI interactions; ^{19}F has an unpaired neutron, and the heavy ^{127}I is suited for SI interactions. PICO-60 is competitive in SD limits, reaching $\sim 5 \times 10^{-40}$ cm^2 for the WIMP-proton cross section for a 50 GeV WIMP [49]. PICO-2L has better sensitivity at low WIMP masses, being world leading below 50 GeV for the WIMP-proton cross section, see figure 1.12.

1.3.3.3 Cryogenic / Solid State Detectors

Cryogenic detectors operate below 100 mK and aim to detect tiny increases in temperature caused by nuclear recoils. All measure phonons, but depending on the target, either scintillation

or ionisation can also be measured. They have the highest sensitivity (among direct detection methods) for low mass (< 10 GeV) WIMPs.

Notable cryogenic detectors include the Cryogenic Dark Matter Search (CDMS) at Soudan, the Cryogenic Rare Event Search with Superconducting Thermometers (CRESST) at Gran Sasso, Expérience pour DEtecter Les WIMPs En Site Souterrain (EDELWEISS) at the LSM (Modane Underground Laboratory) and the future collaboration between CRESST and EDEWEISS, the European Underground Rare Event Calorimeter Array (EURECA), which will also be at the LSM. CDMS measures ionisation and phonons within germanium and silicon crystal substrates. CRESST uses scintillating calcium tungstate crystals, and EDELWEISS measures the ionisation produced in a semiconducting germanium crystal. They all operate with the same principle for detecting phonons; the resistance of semiconductors held below their critical temperature have a very strong temperature dependence. The resistance is constantly and accurately monitored and a sudden change indicates a particle interaction that has raised the temperature by a tiny amount.

SuperCDMS was a larger version of CDMS that began operation in 2009. It has been selected as one of the 'second generation' of direct dark matter experiments for funding by the US Department of Energy (DOE), alongside LZ (discussed in chapter 6) and ADMX-Gen2, an experiment searching for axions (an alternative dark matter model, see section 1.5). The next phase of SuperCDMS will be situated at SNOLAB and the design consists of modular detectors arranged in towers that provide electrical connections and cooling. Detectors will be operated in different ways; for example, some will have ultra-low energy thresholds but reduced background rejection. It is also designed to allow additional detector target mass to be added up to 400 kg. More than one type of detector can be used; SuperCDMS has tested HV and iZIP (interdigitated Z-sensitive Ionization and Phonon-mediated) detectors, with Ge and Si types of each. The HV detectors operate with a high gain for ionisation amplification and so can reach very low thresholds, but they lose background discrimination as the phonon signal is often buried by phonons from the charge propagation. iZIP detectors allow surface event rejection and NR/ER discrimination, but are less sensitive to low mass WIMPs. Commissioning of SuperCDMS phase 2 will begin in 2018.

1.3.3.4 Directional Detection

If we are able to resolve the direction of a nuclear recoil, we can determine the direction of the incoming particle, and due to the motion of the solar system around the galactic centre and therefore through the dark matter halo there is expected to be a directionality to WIMPs, see section 1.2.3. It has been calculated that WIMPs should appear to come from the constellation of Cygnus. Directionality is desirable as not only would it confirm a dark matter signal as galactic in origin, it may also provide a background rejection technique for neutrino-nucleus scattering, an irreducible background mimicking WIMP-nucleus interactions, as solar neutrinos could be identified as originating from the sun. Neutrino-nucleus scattering is discussed in section 1.4.

Directional detectors have so far proved to be far less sensitive than other direct detection methods. Current directional time projection chamber (TPC) technologies use a gaseous target in order to obtain particle tracks [55]. This limits sensitivity as it is much more difficult and expensive to get to large target masses for high exposures when using gas. The first gas TPC for directional dark matter detection was developed by the DRIFT collaboration [56]. DRIFT uses CS_2, CF_4 and O_2 in a mixture of 73%, 25% and 2% respectively. An ionisation signal is transported by CS_2^- ions, and the inclusion of O_2 causes ions to travel at slightly different velocities, resolving the position of ionising events in the z direction. The original motivation for using ions was that they do not experience the same spreading that an electron cloud does, which impacts position resolution. DRIFT has successfully operated a m^3-scale detector at Boulby, and intends to scale up to 10 m^3. Several other collaborations, mainly using CF_4, aim to develop electron-drift TPCs, including MIMAC (Modane) [57], NEWAGE (Kamioka) [58] and DMTPC (SNOLAB) [59].

1.3.3.5 Discovery Claims

The DAMA/NaI and DAMA/LIBRA experiments have claimed to see an annual modulation in event rate that they attribute to dark matter. DAMA/NaI consisted of nine 9.70 kg sodium iodide (NaI) crytals and published a series of results, ending with a claim of a conclusive confirmation of dark matter annual modulation at 6.3σ [60]. DAMA/LIBRA was a follow-up experiment with a 250 kg NaI target. It began operation in 2003 and published final results in 2013, claiming a 7.5σ discovery of dark matter annual modulation; combining this with the first result gives a 9.3σ result [61]. Figure 1.13 shows the low energy event rate modulation observed by DAMA/LIBRA. Although it may appear convincing, no other experiment has been able to confirm the results, and the WIMP masses and cross-sections DAMA/LIBRA favour have been repeatedly ruled out by limits set by other experiments. Many have attempted to explain the signal or reconcile it with other constraints, for example by assuming WIMP-electron scattering [62], or with non-dark matter explanations such as muon induced neutrons [63], but these were subsequently ruled out as a possible explanation [64].

The CoGeNT (Coherent Germanium Neutrino Technology) experiment was designed with high sensitivity to low mass WIMPs (m_χ < 10 GeV), with the aim to confirm the DAMA results. CoGeNT's first results found an annual modulation signal at a modest 2.2σ in its bulk events, compatible with the DAMA results [65]. The amplitude of their signal is $4 - 7\times$ larger than that expected for a standard galactic halo, suggesting a non-Maxwellian component in the velocity, by which they attempt to explain the lack of corroboration by other experiments. However, these claims are also disputed; for example, an independent reanalysis of the CoGeNT data suggested the preference for light dark matter nuclear recoils was $< 1\sigma$ [63], and a further analysis by CoGeNT using maximum likelihood methods found a significance of only 1.7σ [66]. The consensus within the dark matter community on the DAMA measurements is that it is not consistent with a dark matter signal, as the results from LUX rule out the light mass WIMP
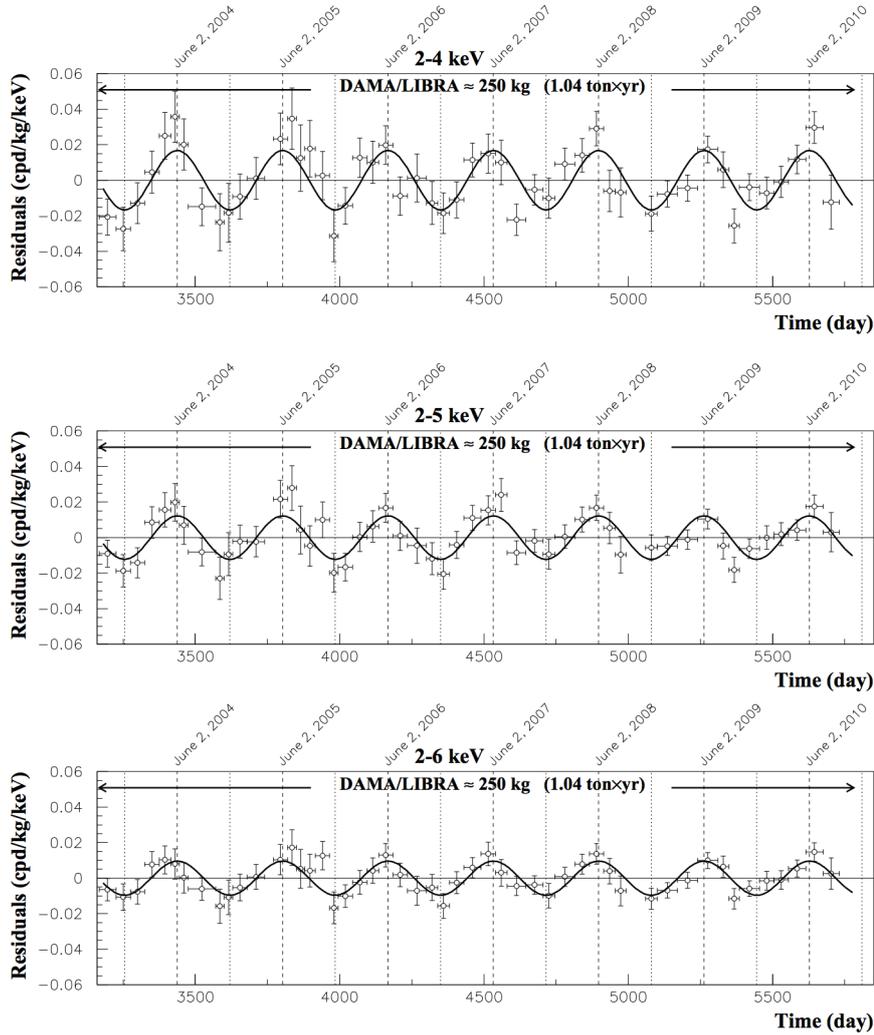
Figure 1.13: The DAMA/LIBRA annual modulation signal, shown as event count residuals for 2-4 (top), 2-5 (middle) and 2-6 (bottom) keV windows [61]. Superimposed curves are of the form $A\cos\omega(t-t_0)$ with a period $T = 2\pi/\omega = 1$ year, and $t_0 = 152.5$ days. This results in a maximum on June 2nd.

interpretation. DAMA and CoGeNT do not have electron and nuclear recoil discrimination, so the most likely interpretation is that there is some form of background with a modulation they have not quantified correctly, or that there is a fault in the apparatus.

There are experiments underway that aim to test the validity of these annual modulation signals. DM-Ice is a 17 kg NaI(Tl) experiment located in ice at the South Pole, that whilst not excluding the DAMA/LIBRA signal, has reported data consistent with no modulation in the energy range 4-20 keV [67]. Current plans are to extend the detector to 250 kg, also operating in both the northern hemispheres (Yangyang, South Korea). The SABRE experiment is similar, with plans to operate two high-purity NaI(Tl) crystals in both hemispheres [68]. Future results from the experiments will shed light on the true nature of the DAMA/LIBRA results.

1.3.3.6 Direct Detection Experimental Landscape

Each type of dark matter detector has its advantages and disadvantages, but together they can scan a large range of WIMP masses and cross sections. Figure 1.14 shows how different types of experiment have contributed to the continued increase in sensitivity to the spin-independent cross section, and how they will and will continue to do so.
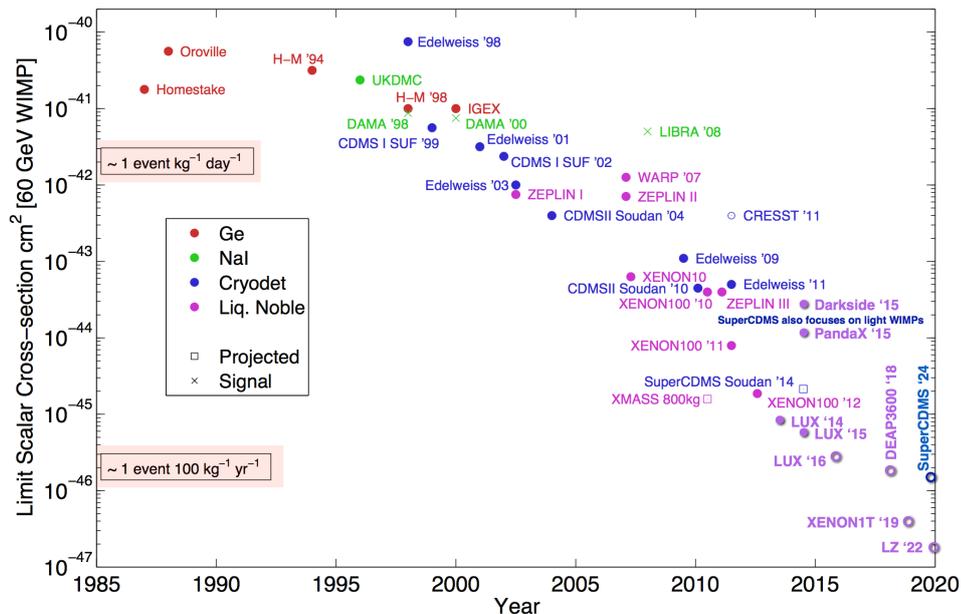
Figure 1.14: Evolution of the limit on the WIMP-nucleon spin-independent cross section for a 60 GeV WIMP, with colours corresponding to technology types. Early progress was made mainly with Ge, but more recently liquid noble detectors dominate. [69].

The landscape of the WIMP-nucleon scattering cross section is set to change over the coming years as the generation II experiments such as LZ, SuperCDMS and ADMX-Gen2 come online.

1.4 Coherent Neutrino-Nucleus Scattering

WIMP limit plots often contain a shaded region at the bottom indicating a process known as coherent neutrino-nucleus scattering. This is a predicted Standard Model process, currently unmeasured, which poses a significant problem for future dark matter detectors. The effect of the scattering is a nuclear recoil, indistinguishable in most detectors from that caused by a WIMP. The coherent scattering cross section is given as [70]:

$$\frac{d\sigma_\nu}{dE_R} = \frac{G_F^2}{4\pi} Q_W^2 m_T \left(1 - \frac{m_T E_R}{2E_\nu^2}\right) F(Q^2)^2 \tag{1.39}$$

where Q_W is the weak charge and is dependent on Z and A:

$$Q_W = (Z - A) - (1 - 4\sin^2\theta_W)Z \tag{1.40}$$

24

where θ_W is the Weinberg (weak mixing) angle. All other symbols have their original meanings as in section 1.2.4. There is a 5% uncertainty on the cross section, which is dominated by the form factor $F(Q^2)$ [70]. Because of the dependence of the cross section on energy, neutrinos from different sources (solar, atmospheric, supernovae) will contribute by different amounts depending on their energy, see figure 1.15. In general, the solar neutrinos have the lowest energy; all are less than 20 MeV. Diffuse supernovae neutrinos reach to 100 MeV whilst atmospheric range from $\sim 10 - 1000$ MeV.
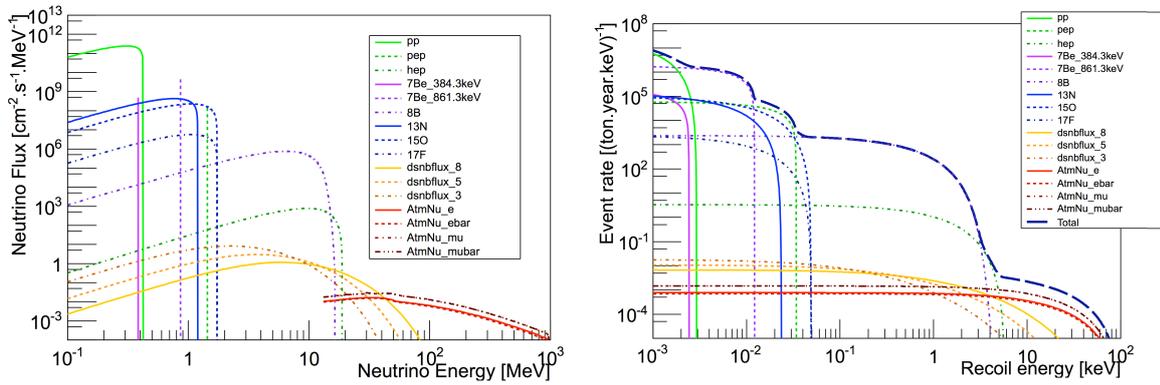
(a) Neutrino fluxes as a function of neutrino energy. (b) Event rates in xenon as a function of recoil energy.

Figure 1.15: Fluxes and the resulting nuclear recoil event rates in a Xe target for all neutrino sources relevant to dark matter [71]. Solar neutrinos are shown in green (fusion reactions) purple (^{7}Be and ^{8}B decay) and blue (CNO cycle decays). Diffuse supernovae neutrinos are shown in yellow. Atmospheric neutrinos and anti-neutrinos are shown in red/brown.

Notice in figure 1.15(b) that the recoil energy range near threshold is significant. For example, the LUX experiment (see chapter 2) used a 1.1 keV threshold for its 2015 results [46]. For high exposures, the neutrino rate at this energy is non-negligible ($> 10^3$ counts per ton year). Therefore, as experiments grow in size, these events will become a problematic aspect. However, once the background is large enough to measure and characterise, it can be treated as another well-understood background and could be subtracted.

There are lots of uncertainties in the neutrino flux, and so if a small signal was observed, it would be difficult to determine if it was a low mass WIMP or an excess of neutrinos. Directionality is a potential solution, as solar neutrinos come from the sun, so knowing the direction of the recoil could allow these events to be removed. A signal at higher WIMP mass, however, would be more problematic as supernovae and atmospheric neutrinos are uniform in direction.

As figure 1.16 demonstrates, the lower the cross section an experiment can reach to, the more neutrino events will be detected. It is a general hope of the dark matter community that the WIMP-nucleon cross section lies above the neutrino floor, as motivated by many WIMP models.
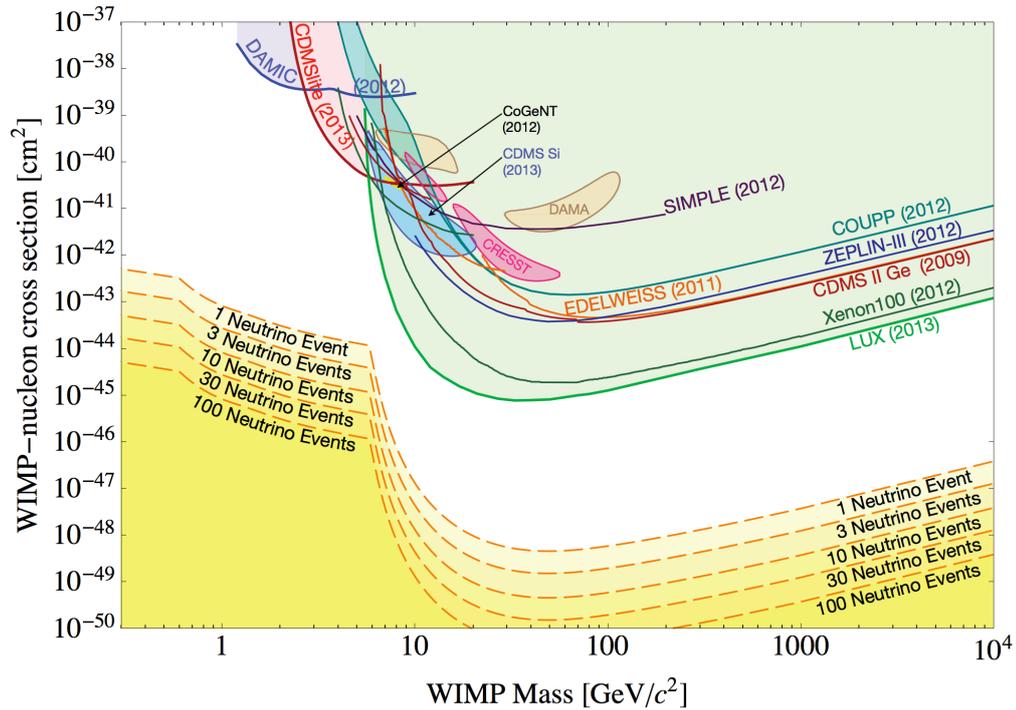
Figure 1.16: Neutrino isoevent contour lines are shown as dashed orange lines in the WIMP mass and cross section phase space. Experimental limits up to 2013 are shown from DAMIC, SIMPLE, COUPP (SDDs), EDELWEISS, CDMS, CDMSlite (cryogenics), XENON10, XENON100 and LUX (liquid xenon) [71]

1.5   Other Dark Matter Candidates

Whilst the focus of this work is on WIMPs, it is important to note there are many other potential candidates that may make up some or all of dark matter. Most are not only motivated by the evidence for dark matter, but by other unsolved problems in particle physics. Table 1.2 lists several well-motivated dark matter candidates and their properties, and they are discussed below.

Table 1.2: A summary of dark matter candidates, their properties and detection modes. Early universe refers to a signal or deviation from standard expectations in phenomena such as Big Bang Nucleosynthesis and the CMB. Collisonal refers to a high elastic scattering cross section between DM particles. GHP is gauge hierarchy problem and NPFP is new physics flavour problem [22].

	WIMP	SuperWIMP	Light $\tilde{G}$	Hidden DM	Sterile ν	Axion
Motivation	GHP	GHP	GHP/NPFP	GHP/NPFP	ν mass	Strong CP
Naturally Correct Ω	yes	yes	no	possible	no	no
Production Process	freeze-out	decay	thermal	various	various	various
Mass Range	GeV–TeV	GeV–TeV	eV–keV	GeV–TeV	keV	μeV–meV
Temperature	cold	cold/warm	cold/warm	cold/warm	warm	cold
Collisional	no	no	no	possible	no	no
Early Universe	no	yes	no	possible	no	no
Direct Detection	yes	no	no	possible	no	no
Indirect Detection	yes	possible	no	possible	yes	no
Particle Colliders	yes	yes	yes	possible	no	no

Axions are, after WIMPS, the second most popular dark matter candidate. They are motivated by the strong CP problem, which is the question of why quantum chromodynamics (QCD), the theory of the strong interaction, is not observed as breaking CP-symmetry experimentally, when the theory predicts it should. Axions are very light; astrophysical bounds place them at $m_A < 10$ meV. Axions interact with fermions, gluons and photons (at loop level), and could be detected directly by scattering with SM particles, or by the axio-electric effect, a photon-producing effect that requires a magnetic field. There are several experiments searching for axions; notably, the Axion Dark Matter Experiment (ADMX), utilities a resonant microwave cavity within a superconducting magnet to search for galactic dark matter. ADMX was able to exclude dark matter axions with masses between 1.9 and 3.53 μeV. [72]. A second generation of ADMX, ADMX-Gen2, will probe a second axion model in the range 2–20 μeV.

SuperWIMPs are Super Weakly Interacting Particles. These also exploit the WIMP miracle, as WIMPs are still produced and undergo freeze-out, but then decay later in the universe to SuperWIMPs, inheriting the relic density. There are several good candidates for SuperWIMPs in Supersymmetry, such as gravitinos and axinos (superpartners of the graviton and axion). Supersymmetry, or SUSY, is the theory that every particle has a heavier superpartner, with fermions having a boson counterpart and vice versa. The implication of SuperWIMPS for detection are interesting; direct detection becomes impossible, but indirect detection is possible through cosmic ray and neutrino experiments. If WIMPs and SuperWIMPs are charged, strong bounds on charged WIMPs suggest they may decay, meaning they could be produced by ultra-high energy cosmic rays. Collider experiments could also see long-live charged particles. SuperWIMPs could behave as warm dark matter and so may suppress small scale structures.

Light gravitinos, $\tilde{G}$, are motivated by the gauge hierarchy problem; this is the large discrepancy between the strength of gravity and all other forces. They arise in gauge-mediated supersymmetric breaking models and also solve the new physics flavour problem, which is the disagreement between the new physics scale needed to solve the hierarchy problem and the experimental bounds placed by flavour physics. The details of these models are beyond the scope of this work, but it is worth knowing that there are extensions to them that produce viable small components of dark matter. These can only be detected in particle colliders.

Hidden dark matter is that which has no Standard Model interactions whatsoever. If a whole host of hidden particles and interactions with the same couplings as the visible ones exist, the WIMP miracle could be replicated to give the relic density entirely in this 'hidden sector'. Such a sector could be produced in SUSY breaking. A true hidden particle could not be detected via direct detection, but there are suggestions of additional particles and non-gauge interactions that may couple the Standard Model to the hidden sector, so there could still be possibilities. Like the SuperWIMPS, there may also be astrophysical signals.

Sterile neutrinos help to solve the problem of neutrino masses not existing in the Standard Model, and a see-saw mechanism giving mass to both the sterile neutrino and the Majorana

neutrino explains why normal neutrinos are so light compared to other particles. Sterile neutrinos have no SM gauge interactions but can be produced in neutrino oscillations. Their abundance can be enhanced by a lepton number asymmetry in the early universe. Sterile neutrinos may be detected only through indirect detection of their decay products.

It is worth mentioning here that the scalar-tensor-vector gravity theory, often referred to as MOdified Gravity (MOG) , explains not only galactic rotation curves [73], but the mass of galaxy clusters [74], gravitational lensing in the Bullet Cluster [75], and cosmological observations without the need for dark matter [76]. However, because this includes a Lorentz-violating modification of Einstein's theory of general relativity, considered an extremely well-tested cornerstone of physics, it has not been considered particularity viable by most physicists.

DUAL-PHASE XENON TPCS AND THE LUX EXPERIMENT

The LUX (Large Underground Xenon) Experiment is a world-leading direct dark matter search and uses dual-phase (liquid and gas) xenon technology. It was the world's first sub-zeptobarn detector, placing stronger constraints on dark matter properties than all its predecessors. It is located in the Davis Cavern of the Sanford Underground Research Facility (Lead, South Dakota), 4,850 feet under the ground. This chapter will discuss the physical principles involved in producing a signal in a dual-phase TPC, give an overview of the components of the LUX detector itself and describe both the data taking systems that are in place and the simulation package developed for LUX.

2.1 Signal Production in a Dual-Phase Xe TPC

2.1.1 Energy Transfer

A particle that undergoes interactions with nuclei or electrons in liquid xenon will transfer an energy, E_0, through three channels: ionisation, excitation and heat:

$$E_0 = N_i E_i + N_{ex} E_{ex} + N_i \eta \tag{2.1}$$

where E_i and E_{ex} are the mean ionisation and excitation energies, N_i and N_{ex} are the number of ionised and excited atoms, and η is the mean kinetic energy of ionised electrons. η is thought to lie between 4.65 and 5.25 eV [26]. In a dual-phase xenon TPC, the energy loss to heat is undetectable. Other dark matter detectors utilise the heat channel with cryogenic crystals, and most direct dark matter searches attempt to use two of the three channels. In the case of a dual-phase xenon TPC those channels are excitation and ionisation; the principles of the detector rely on detecting these as two separate signals.

When considering nuclear recoils, an additional term is needed to account for recoil elastic scattering - collisions that do not result in excitations or ionisations. However, continuing with the electron recoil case and using the atomic ionisation potential I, the energies in equation 2.1 can be replaced, leading to:

$$\frac{E_0}{I} = N_i \frac{E_i}{I} + N_{ex} \frac{E_{ex}}{I} + N_i \frac{\eta}{I} \tag{2.2}$$

This is relevant only in the gas phase. In the liquid phase, I can be replaced with $E_g = 9.22$ eV, the band gap, due to the band structure of electronic states within liquefied xenon. Defining the W-value as $W = E_0/N_i$, i.e. the energy transfer per ionisation, we may rewrite:

$$\frac{W}{E_g} = \frac{E_i}{E_g} + \frac{E_{ex}}{E_g} \frac{N_{ex}}{N_i} + \frac{\eta}{E_g} \tag{2.3}$$

The ratio N_{ex}/N_i has been both calculated and measured in xenon for electrons; 0.06 is the calculated value but 0.2 is measured, so often their average is often used [26]. For nuclear recoils, a much higher ratio of $N_{ex}/N_i \sim 1$ has been found to better fit data [77]. This difference is the underlying principle of discrimination between the types of recoil in a two-phase detector, as we will see in chapter 5.

2.1.2 Primary Scintillation: S1

The scintillation process in liquid xenon (LXe) begins when a particle elastically scatters from either an electron or a xenon nucleus, causing a recoil, which in turn causes a cascade of secondary recoils. There are then two processes that lead to the emission of light known as the primary scintillation signal, S1. The first is impact of either a recoiling nucleus or electron with a Xe atom forming an excited state, which leads to eximer formation and the creation and decay of an excited diatomic molecule [26]:

$$X + \mathrm{Xe} \rightarrow \mathrm{Xe}^* + X \tag{2.4}$$

$$\mathrm{Xe}^* + \mathrm{Xe} \rightarrow \mathrm{Xe}_2^{*,v} \tag{2.5}$$

$$\mathrm{Xe}_2^{*,v} + \mathrm{Xe} \rightarrow \mathrm{Xe}_2^* + \mathrm{Xe} \tag{2.6}$$

$$\mathrm{Xe}_2^* \rightarrow \mathrm{Xe} + \mathrm{Xe} + \gamma \tag{2.7}$$

where X is either a recoiling nucleus or electron. The superscript v refers to states with vibrational excitation. Vibrational excitation is mostly non-radiative, but infrared photons are sometimes emitted. The VUV scintillation photon is emitted from one of two of the lowest excited states of the Xe_2 molecule, $^3\Sigma_u^+$, the triplet state, and $^1\Sigma_u^+$, the singlet state, as they transition to the ground state, $^1\Sigma_g^+$. This has a repulsive potential, causing dissociation of the molecule into two Xe atoms. The decay time depends on the state; 4.3 ns for the singlet state, and 22 ns for the triplet state, known as the fast and slow decay components respectively.

The second process that leads to scintillation light involves recombination of the ionisation electrons [26]:

$$X + \text{Xe} \rightarrow \text{Xe}^+ + e^- + X \tag{2.8}$$

$$\text{Xe}^+ + \text{Xe} + \text{Xe} \rightarrow \text{Xe}_2^+ + \text{Xe} \tag{2.9}$$

$$e^- + \text{Xe}_2^+ \rightarrow \text{Xe}^{**} + \text{Xe} \tag{2.10}$$

$$\text{Xe}^{**} + \text{Xe} \rightarrow \text{Xe}^* + \text{Xe} \tag{2.11}$$

$$\text{Xe}^* + \text{Xe}+ \rightarrow \text{Xe}_2^{*,v} \tag{2.12}$$

$$\text{Xe}_2^{*,v} + \text{Xe} \rightarrow \text{Xe}_2^* + \text{Xe} \tag{2.13}$$

$$\text{Xe}_2^* \rightarrow \text{Xe} + \text{Xe} + \gamma \tag{2.14}$$

As the final stage of this is similar to the first, the decay times and wavelengths are similar, but the population of the two excited states is different. Moreover, the recombination process introduces a significant time delay. In xenon, recombination is slow compared to the de-excitation times, and so this adds a non-exponential third component to the decay curve, on top of the fast and slow components. This can be removed by the application of an electric field, as this prevents recombination. The population of the two states is different depending on the species of the initial particle causing the recoil. This provides a method of discrimination (and in fact the only method in single phase detectors) by pulse shape. In argon, the decay times of the two states are different enough that the shape of the S1 pulse can be used to discriminate between electron and nuclear recoils extremely well. This is much more difficult in xenon, but it has been used, for example by the ZEPLIN-I experiment [35]. Note that in both the excitation and recombination routes, any additional heat generated is represented by the additional Xe atoms on the right hand side of the equations.

If the assumption is made that all excited atoms produce one VUV photon, and that all electron-ion pairs recombine, the number of emitted photons can be written as:

$$N_{ph} = N_{ex} + N_i \tag{2.15}$$

and so the scintillation efficiency in units of keV/photon is defined at zero electric field as:

$$W_s = \frac{E_0}{N_{ph}} \tag{2.16}$$

This varies strongly with electric field strength due to the fact the field will prevent recombination, thus partially removing one of the routes for VUV photon emission. The value of W_s depends on both the particle species and its energy. For example, for low energy (20-100 keV) electrons, $W_s = 18.3 \pm 1.5$ eV [78], whilst for 20 keV nuclear recoils, $W_s = 110 \pm 20$ eV [26]. A relationship between the W-value previously defined and W_s can be formulated using equation 2.15:

$$W_s^{min} = \frac{W}{1 + N_{ex}/N_i} = \frac{E_0}{N_i + N_{ex}} \tag{2.17}$$

where W_s^{min} is the minimum possible energy needed to produce a scintillation photon with no quenching processes in play. In this work, for liquid xenon, we will use an error-weighted average of 13.7 eV, using 13.8 ± 0.9 eV from Doke et al. [78] and 13.7 ± 0.2 eV from Dahl [77]; other values are not included as they are reported without uncertainty, with a calibration error or are simply old and have been superseded by technological improvements.

For particles with a high linear energy transfer (LET) such as α-particles, the particle track becomes dense enough with excited particles that these particles start to undergo collisions. This leads to the phenomenon of 'bi-excitonic quenching':

$$\text{Xe*} + \text{Xe*} \rightarrow \text{Xe} + \text{Xe}^+ + e^- \tag{2.18}$$

This process reduces the number of emitted photons from two (one from each Xe*) to one or none, depending on whether the Xe^+ recombines to reform an excited state. This is thought to explain why the amount of scintillation light falls significantly for high LET α-particles (see figure 2.2).

Recombination light is a significant contribution to the total light yield. For a low LET electron, the track is well-defined and can be modeled as a column of ions spaced with the average distance between them close to the Onsager radius:

$$r_c = \frac{e^2}{4\pi\epsilon_0\epsilon kT} \tag{2.19}$$

where r_c is 49 nm for liquid xenon. Here ϵ is the dielectric constant, k is the Boltzmann constant and T is temperature. The Onsager radius is determined by setting the electrostatic attraction of an electron to its parent ion equal to the kinetic energy of its thermal motion. Therefore, within r_c, electrostatic attraction dominates, whilst outside of it, thermal motion will draw the electron away from the ion. The thermalisation length for electrons in LXe is $r_{th} \sim 4.5$ μm and so many electrons will be too far from a positive ion to recombine after thermalisation, causing a reduction in scintillation light at low LET.

For low energy electrons and nuclear recoils below 100 keV, a different trend is seen, due to the different track structure of the particles. Recombination along low energy nuclear recoil tracks of interest for WIMP search is not well understood. The cylindrical model described above can not be used, due to the fact that the cascade of recoils tends to spread in all directions rather than along a track. The cascades are usually < 100 nm, much smaller than r_{th}, so most electrons should escape recombination. However, it is experimentally observed that nuclear recoils channel more energy into excitation, leading to the assumption $N_{ex}/N_i \sim 1$ as stated previously. Why exactly this ratio is much higher for nuclear recoils than electrons is still not clear, but one proposal is that some atomic energy levels may be lowered during collision of xenon atoms. This altering of energy levels may reduce the number of ionisation channels available, causing the ratio to increase [77]. We find that the scintillation signal, S1, is proportional to N_{ph}:

$$S1 \propto N_{ph} = N_{ex} + rN_i \tag{2.20}$$

where r is the charge recombination fraction.

2.1.3 Secondary Scintillation or Ionisation: S2

As previously discussed, for electron recoils $N_i >> N_{ex}$, whilst for nuclear recoils $N_i \sim N_{ex}$. Due to the difficulty in measuring the energy per electron-ion or electron-hole pair, $W = E_0/N_i$, the total ionised charge, Q_0, is defined:

$$Q_0 = N_i e \tag{2.21}$$

where e is the electric charge, and a parameterisation used to extrapolate the extracted charge as a function of electric field $\mathscr{E}$:

$$Q(\mathscr{E}) = \frac{Q_0}{1 + \dfrac{k}{\mathscr{E}}} \tag{2.22}$$

It should be noted that whilst $Q(\mathscr{E})$ works well, the physics used to derive it assumes recombination in a cylinder of positive and negative ions, which as mentioned before, is not a good model for low energy electrons and nuclear recoils. However, plotting $1/Q$ against $1/\mathscr{E}$ has an intercept of $1/Q_0$. This allows a measurement of N_i [26].

In order to produce the ionisation, or secondary scintillation signal, S2, the requirements are that the electrons and ions escape recombination, the charge carriers have high mobility and a low probability to form low mobility states along the drift path, and a high gain amplification mechanism. Xenon has a hole-type conductivity, meaning that charge moves through it via electron deficiencies or positive vacancies, and these have a higher mobility than other positive ions. It is essential that the liquid is as pure as possible as electronegative impurities have a large effect on the negative charge carrier mobility; for example, oxygen and water will capture electrons.

The electron drift velocity increases with the field; for example, it is 2.25 mm/μs at 1 kV/cm and 2.8 mm/μs at 10 kV/cm. After 10 kV/cm the velocity is saturated and no longer increases. Electrons will also diffuse along their drift paths. At zero field, they follow the diffusion equation isotropically:

$$\frac{eD}{\mu_0} = kT \tag{2.23}$$

where D is the diffusion coefficient and μ_0 is the zero-field mobility. In a field, this is modified to:

$$\frac{eD}{\mu} = F\langle\epsilon\rangle \tag{2.24}$$

where $\langle\epsilon\rangle$ is the mean electron energy and F is a constant depending on the electron energy distribution. Once the field is applied, diffusion is no longer isotropic, so D is split into D_T and D_L, transverse and longitudinal diffusion coefficients. For fields $\gtrsim 1$ kV/cm, the ratio D_L/D_T is found to be ~ 0.1; this should approach 1 as the field goes to zero. Whilst it may seem that the spread of an electron cloud as it drifts is detrimental to the performance of a dark matter detector as it will affect position resolution, it can be used for measuring the drift time, allowing ionisation-only searches. These are useful for detecting low energy events where no S1 was produced or measured, and this extends detector sensitivity to lower WIMP masses.

Extraction of electrons from the liquid to the gas is difficult, as it is energetically favourable for them to remain in the liquid. Many will not cross on the first attempt and will instead be reflected and scatter until they gain enough energy to cross. An external field is essential for reducing the height of the potential barrier.

Once extracted, the electrons accelerate in the gas and produce electroluminescence. Atoms are again excited to higher states and emit VUV photons as they decay into the ground state, producing an ionisation signal, S2. The observed signals for S2 is proportional to the number of and free electrons (N_e):

$$S2 \propto N_e = (1-r)N_i \qquad (2.25)$$
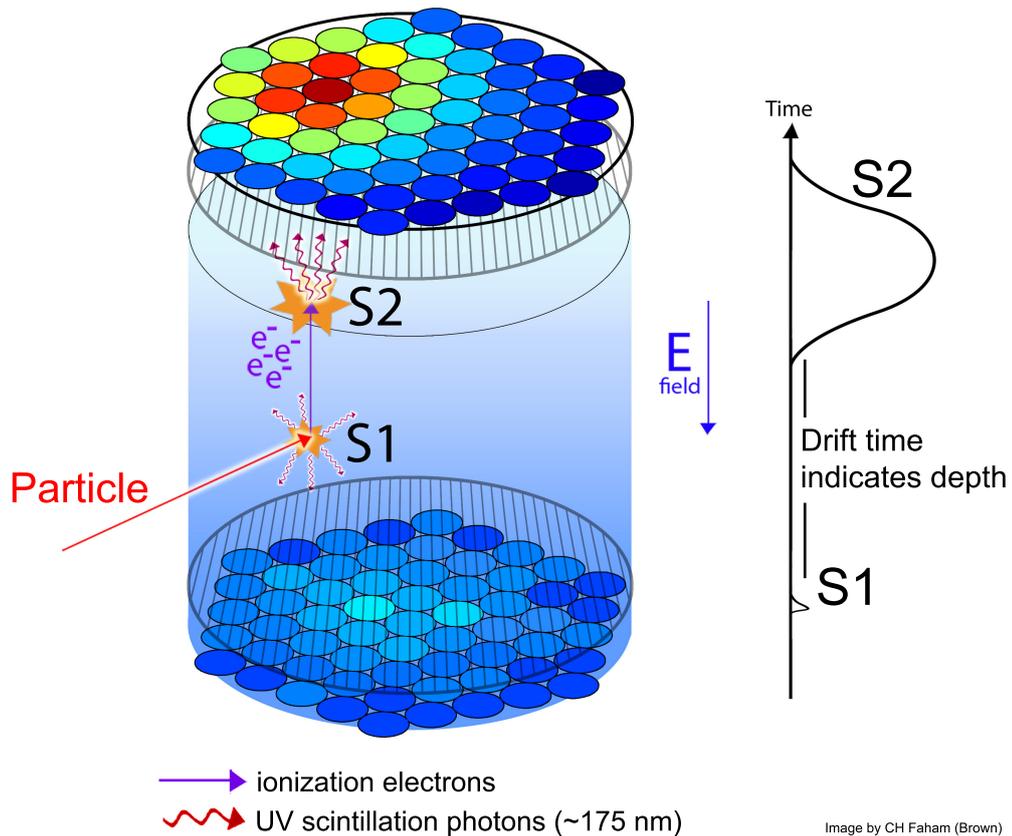
Figure 2.1: Schematic of a single scatter event in a dual-phase TPC, showing the production of an S1 signal in the liquid, the drifting of ionised electrons and the production of the S2 signal within the gas layer. The top and bottom are viewed by photo-multiplier tubes (PMTs). This setup allows 3D position reconstruction, as PMT arrays can locate the x and y co-ordinates, whilst the time between the two signals is a measure of the z co-ordinate.

2.2 Light Propagation in Liquid Xenon

Whilst liquid xenon is mostly transparent to VUV light, there will still be some attenuation of scintillation photons before they can be detected. Photons can be absorbed by impurities or may Rayleigh scatter elastically. Impurities can have a very large effect; 1 ppm of water vapour will absorb most scintillation light in less than 10 cm [79]. The attenuation length due to both processes can be written as:

$$\frac{1}{l} = \frac{1}{l_a} + \frac{1}{l_s} \tag{2.26}$$

Where l_a and l_s are the absorption and scattering lengths respectively. When reflective surfaces are utilised and surround the liquid, l_a can be > 100 cm [79]. It can be assumed that $l_a > l_s$, and so $l \approx l_s$, meaning Rayleigh scattering is the main contributor to light loss. For liquid xenon, $l_s \approx 30$ cm. The larger a liquid xenon detector is, the more Rayleigh scattering of photons will take place and this can impact position reconstruction. This has to be mitigated with accurate simulations of light propagation within the detector and uniformly distributed, mono-energetic calibration sources. Furthermore, light collection efficiency will impact energy resolution, and so it is important to choose a material with high reflectivity in VUV wavelengths to coat the inside of any liquid xenon detector. For LUX, Polytetrafluoroethylene (PTFE) has been used; PTFE reflects over 70% of LXe scintillation light in vacuum, and $\sim 100\%$ in LXe, although this may vary with the angle of incidence [80]. PTFE is also compatible with the low radioactivity requirements of dark matter experiments.

2.3 Light and Charge Yield in Liquid Xenon

2.3.1 Light Yield

Now that the processes behind signal production in LUX and other dual-phase xenon TPCs have been discussed, we move on to some important fundamental quantities that are important in determining detector efficiencies in analysis. The light yield, or $L_y(E_0, \epsilon)$, of liquid xenon is defined as the number of photons emitted in scintillation per keV, and is a function of the energy transfer E_0 and the electric field $\mathscr{E}$.

Figure 2.2 shows how the light yield increases with increasing LET, reaches a plateau corresponding to W_s^{min}, and falls again at higher LET, as discussed previously. This particle dependence can pose a problem, and so in many dark matter searches, the light yield for nuclear recoils is parametrised by an efficiency, $\mathscr{L}_{eff}$, using a W_s value, defined at zero field as [26]:

$$\mathscr{L}_{eff} = \frac{W_{s,e}(122 \text{ keV})}{W_{s,Xe}(E_0^{nr})} \tag{2.27}$$

where E_0^{nr} is the nuclear recoil energy (differentiated from E_0 as $\mathscr{L}_{eff}$ is only relevant for nuclear recoils) and the subscripts e and Xe relate to electron and xenon recoils respectively.

$W_{s,e}(122 \text{ keV})$ is the W_s value for for 122 keV γ-rays from a ^{57}Co source. ^{57}Co is convenient for calibration but the light yield for 122 keV γ-rays is not maximum, so it can introduce a complication into measurements, as quenching is usually calculated using the maximum light yield as a reference.

Using $\mathscr{L}_{eff}$, the light yield when a field is applied can be defined as [81]:

$$L_y = \frac{N_{ph}}{E_{nr}} = \mathscr{L}_{eff} \frac{S_n}{S_e} \frac{N_{ph}(^{57}\text{Co})}{122 \text{ keV}} \tag{2.28}$$

where S_n and S_e are scintillation reduction factors due to an applied electric field, and $N_{ph}(^{57}\text{Co})$ is the number of photons yielded at zero field from the ^{57}Co 122 keV γ-ray.
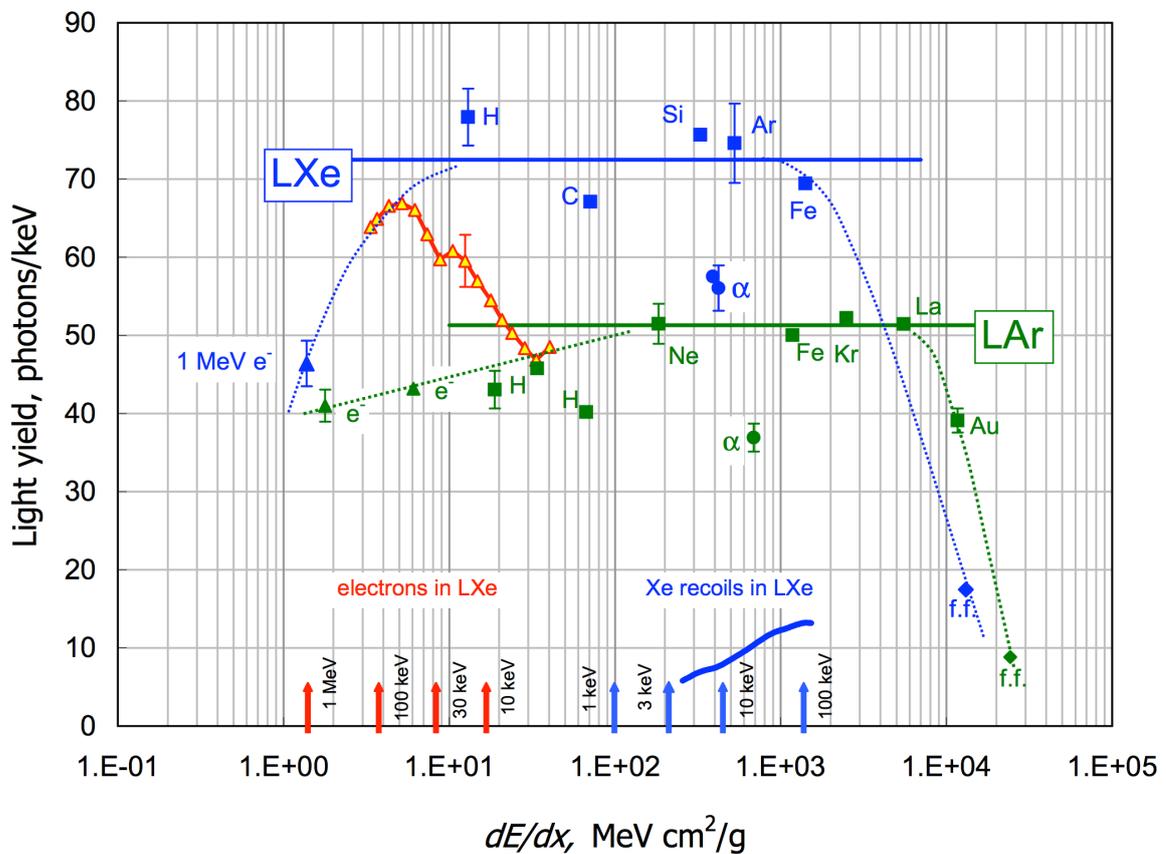
Figure 2.2: Scintillation light yield shown for both liquid xenon (blue) and liquid argon (green) as a function of linear energy transfer. Key in importance are the red triangles which show the light yield of electron recoils in xenon, and the blue line at the bottom of the plot, showing xenon nuclear recoils. The numbers beside the corresponding arrows show the recoil energy in keV. Note the significant quenching of the light yield of nuclear recoils compared to the electrons. The top plateau corresponds to $W_s^{min} = 13.7$ eV for liquid xenon [26].

Figure 2.2 demonstrates measurements of L_y for various particle species in both liquid xenon and liquid argon. Whilst measurements from relativistic electrons, ions and fission fragments are interesting and potentially useful for calibration, the important data for a dark matter search

is shown by the red triangles and the blue line at the bottom showing electrons and Xe nuclear recoils. Note the significant quenching of L_y for nuclear recoils, as previously discussed, and also the turn-over in L_y for electrons at low LET.

For the first LUX results paper, a conservative cut-off in L_y was assumed at 3 keV; the light yield was conservatively estimated to fall to zero below this energy. This was necessary because there were no measurements of L_y below this energy. The pioneering LUX D-D calibration has enabled measurements of the light yield down to 0.7 keV, see section 5.2.4.5.

2.3.2 Charge Yield

The charge, or ionisation yield is denoted $Q_y(E_0, \epsilon)$, and so like L_y is a function of the recoil energy and the applied electric field:

$$Q_y = \frac{N_e}{E_0} \tag{2.29}$$

Therefore, the charge yield describes the number of electrons ionised per keV of recoil energy. LUX measurements of Q_y can be found in section 5.2.4.4.

2.3.3 Combined Energy Scale

In this section we introduce the important quantities of detector-dependent light and charge detection efficiencies. For this purpose, we introduce slightly different notation than before, due to the introduction of the LUX in-situ D-D neutron calibration previously mentioned that eliminates the need to use $\mathscr{L}_{eff}$.

Both N_{ph} and N_e undergo statistical fluctuations dependently, and so using a linear combination of the two signals reduces fluctuations and improves energy resolution. Whilst it is difficult to determine the recombination fraction r, it is true for all values or r that:

$$N_i + N_{ex} = N_{ph} + N_e \tag{2.30}$$

therefore the total number of quanta produced is equal to the total number of ionisations and excitations. Therefore, equation 2.17 can be rewritten to calculate the recoil energy as:

$$E_0 = W_s^{min}(N_{ph} + N_e) \tag{2.31}$$

Then, using equations 2.28 and 2.29:

$$E_0 = W_s^{min}(E_0 L_y + E_0 Q_y) \tag{2.32}$$

and as N_{ph} and N_e are proportional to the observables S1 and S2, we can define two detector dependent gain values, g_1 and g_2:

$$g_1 = \frac{S1}{E_0 L_y(E_0, \mathscr{E})} \tag{2.33}$$

$$g_2 = \frac{S2}{E_0 Q_y(E_0, \mathscr{E})} \tag{2.34}$$

allowing equation 2.32 to be rewritten as:

$$E_0 = W_s^{min}\left(\frac{S1}{g_1} + \frac{S2}{g_2}\right) = 13.7 \text{ eV}\left(\frac{S1}{g_1} + \frac{S2}{g_2}\right) \tag{2.35}$$

For the case of nuclear recoils, we must apply a quenching factor, L, to account for the energy loss to atomic motion:

$$E_0^{nr} = \frac{W_s^{min}}{L}(N_{ph} + N_e) \tag{2.36}$$

and finally again rewrite equation 2.32 as:

$$E_0^{nr} = \frac{13.7 \text{ eV}}{L}\left(\frac{S1}{g_1} + \frac{S2}{g_2}\right) \tag{2.37}$$

g_1 and g_2 can be measured experimentally, see calibration sections 5.1.2, 5.1.4 and 5.2.4. The physical meaning of g_1 is the light collection efficiency referenced to the centre of the detector multiplied by the average quantum efficiency of the photo-multiplier tube (PMT) arrays, and g_2 is the product of the electron extraction efficiency by the average size of a single electron in detected photons. This allows the application of this equation to any event with an S1 and an S2 to reconstruct the total energy of the recoil.

Because of this significant quenching effect for nuclear recoils, we define two energy scales, known as the 'electron equivalent' energy, keV$_{ee}$, and the nuclear recoil energy, kev$_{nr}$ [1]. For energies calibrated with γ-ray sources, keV$_{ee}$ is always used to indicate that the energy corresponds to electron recoil events. As an example of the way these scales work, a γ-ray induced electron recoil of 6 keV$_{ee}$ would produce roughly the same total light and charge as a nuclear recoil of $\sim$30 keV$_{nr}$.

2.3.4 Quenching Models for Nuclear Recoils

There are several different models that can be used for the quenching of nuclear recoils. Lindhard's theory [82] based on heavy ion quasi-elastic collisions gives the quenching factor as:

$$L = \frac{kg(\epsilon)}{1 + kg(\epsilon)} \tag{2.38}$$

where k is a constant of proportionality between electronic stopping power and the recoil velocity of the nucleus, $g(\epsilon)$ and ϵ are given by:

$$g(\epsilon) = 3\epsilon^{0.15} + 0.7\epsilon^{0.6} + \epsilon \tag{2.39}$$

$$\epsilon = 11.5(E_0^{nr}/\text{keV})Z^{-7/3} \tag{2.40}$$

where Z has its usual meaning as the atomic number; $Z = 54$ for xenon, meaning that L becomes important for $E_0^{nr} < 10$ MeV [81].

---

[1] In some works, keV$_r$ is used instead of keV$_{nr}$.

A second model, proposed by Bezrukov *et al* [83], suggests quenching of the form:

$$L = \frac{s_e}{s_n + s_e} \tag{2.41}$$

Where s_e and s_n are electronic and nuclear stopping power, several models of which can be assumed. s_e relates to the slowing down of a ion due to the inelastic collisions between bound electrons, and s_n is assumed to be proportional to the probability that a recoiling xenon atom scatters elastically from another xenon atom.

2.3.5 Requirements for Direct Dark Matter Detection

The signal production described above is valid for many particles across a range of energies, with the type of particle determining the type of recoil. The most relevant to WIMP detection are electrons and γ-rays, which induce electron recoils, and neutrons, which induce nuclear recoils. As the WIMP signal is expected to be a nuclear recoil, neutrons present a particularly dangerous background. However, there is also a finite discrimination factor for electron recoils, which means that a small fraction of these can also mimic a WIMP signal.

In order to partially mitigate the backgrounds from these particles in direct detection experiments, external sources of radioactivity and cosmic rays must be reduced as much as possible. For this reason, they are placed deep underground, and often have further external shielding such as water. Furthermore, they must be constructed from materials that are clean and radiopure, so careful construction and material screening is required.

Additionally, The previous sections have alluded to the ability to discriminate between an electron recoil and a nuclear recoil. The partitioning of energy between scintillation and ionisation in liquid xenon is significant, and is generally parametrised on an event by event basis by looking at S1 vs S2/S1. The two types of recoils from two bands in this phase space, with electron recoils higher in S2/S1 than nuclear recoils. As previously discussed in section 2.1.2, for electron recoils there is less recombination and therefore more electrons extracted ($N_i >> N_{ex}$). For nuclear recoils $N_{ex}/N_i \sim 1$ better fits data, although the mechanism leading to this is not fully understood. Despite this, very good (above 99.5%) discrimination is achieved using only the sizes of the S1 and S2 across energies relevant for WIMP search. Chapter 5 contains details of the LUX calibrations that map out the electron and nuclear recoil bands and measure discrimination.

Finally, as was seen in section 1.2.4 and figure 1.7, the event rate of WIMP-nucleus interactions in a target falls exponentially with energy, which is an effect of the nuclear form factor (equation 1.27). This means that achieving a low energy threshold for xenon interactions is important for maximising the possible number of events in a detector exposure. This ties in with measuring L_y and Q_y at very low recoil energies, which allows a reduction in the energy threshold assumed. Furthermore, a high target mass is also favorable, as it will also enhance the number of possible interactions. In reality, this has to be optimised based on the cost and feasibility of building a large detector.

2.4 The LUX Detector

2.4.1 Xenon

LUX holds 370 kg of liquid xenon, with 250 kg actively monitored. The xenon is continuously circulated in order to allow constant purification. Impurities reduce the electron drift length, which must be long enough for electrons from the bottom of the detector to reach the gas layer. Liquid xenon spills over into a weir, and then flows into a dual-phase heat exchanger and evaporates. The gas is then pumped with a diaphragm pump through a getter for purification. It is then recondensed and returned to the detector.

2.4.2 TPC and Electric Field

The LUX TPC is a dodecagonal structure of diameter 50 cm viewed from above and below by PMT arrays, see figure 2.3. The drift distance (height of the liquid xenon) is 49 cm. The TPC is lined with 12 highly reflective polytetrafluoroethylene (PTFE) panels for aiding light collection. The electric field is generated by a series of five electrodes; a bottom grid 2 cm above the bottom PMT array, the purpose of which is to shield the PMTs from the high voltage on the next grid, the cathode, which lies 4 cm above the bottom PMT array and 49.5 cm below the liquid xenon volume. 5 mm below the liquid surface is the gate grid, and 1cm above the gate is an anode grid. Finally, a top grid lies 4 cm above the anode and 2 cm below the top PMT array, and shields the PMTs from the field in the same way as the bottom grid. The field in the liquid has ranged from 2.9 kV/cm (for LUX's first WIMP search run) to 3.5 kV/cm (for the second and final WIMP search run). The anode and gate work in conjunction to generate a 5 kV/cm extraction field below the liquid, and a 10 kV/cm electro-luminescence field in the gas layer. The cathode high voltage is delivered by a warm feedthrough containing transformer oil and xenon gas. Up to 100 kV can be delivered through the xenon-gas filled umbilical. The drift field in the liquid is shaped uniformly using 48 copper rings of thickness 3.2 mm and spacing 1 cm, held within grooves in the PTFE panels.

2.4.3 PMTs

A total of 122 PMTs monitor the active xenon region. Two copper arrays, each holding 61 PMTs are located at the top and bottom of the TPC, facing towards the xenon volume. The PMTs are the low background Hammamatsu R8778 model with a diameter of 5.6 cm. They are ideal for use with the xenon scintillation light wavelengths, peaked at 178 nm, having a measured average 33% quantum efficiency and 90% collection efficiency at this wavelength, and a signal gain of up to 10^7 at 1,500 V [84]. The PMTs were designed for use at cryogenic temperatures in the liquid xenon range of 165-180 K.
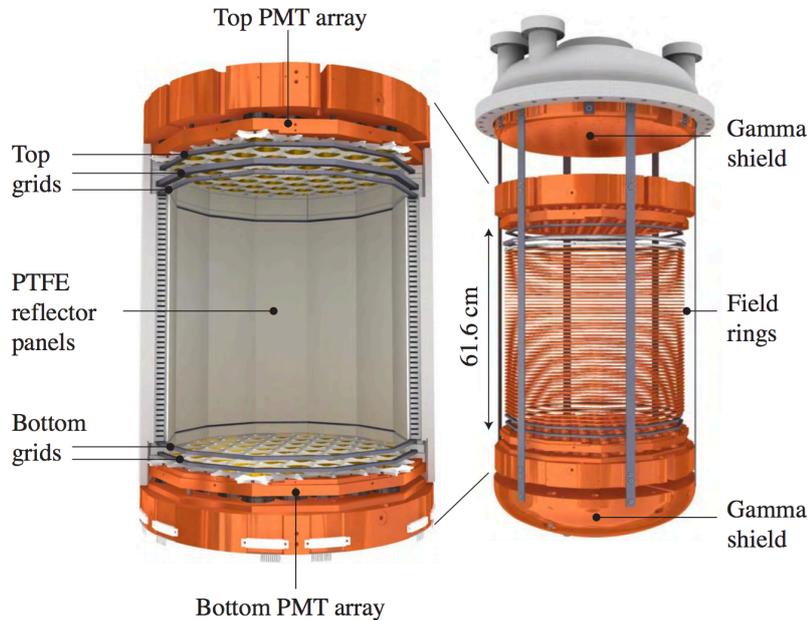
Figure 2.3: The LUX TPC, showing the locations of the grids, PMT arrays, field rings and shielding. On the right, the top of the cryostat is also visible.

2.4.4 Water Veto and Shielding

The location of the detector 4,850 feet (∼ 1.5 km) underground reduces the cosmic ray muon flux by 6 orders of magnitude with respect to the surface [85]. However, muons are still not negligible at an integral flux rate of 4.40×10^{-9} cm^{-2}s^{-1} (compared to 2.68×10^{-3} cm^{-2}s^{-1} at the surface) [86], and cavern rock is radioactive, and so the detector is further shielded with a water tank.

The LUX water tank is of diameter 7.6 m and height 6.1 m, see figure 2.4. The sides of LUX are shielded by 3.5 m of water, whilst the top and bottom are shielded by 2.75 m and 1.2 m respectively. Water lowers both γ-ray and neutron fluxes from cavern radioactivity significantly so that the detector components become the dominant background. The water can be circulated through a purifier to achieve suitable levels of radioactive contamination. The water tank is instrumented with 20 10-inch Hammamatsu R7081 large photocathode area PMTs designed to collect Cherenkov light from cosmic ray muons. These have a peak wavelength of 420 nm so are well suited to detect Cherenkov light, which is peaked in the UV. Any event within the TPC coinciding with light in the water tank can be vetoed as a background muon event. To increase light collection, the water tank is lined with Tyvek reflectors.
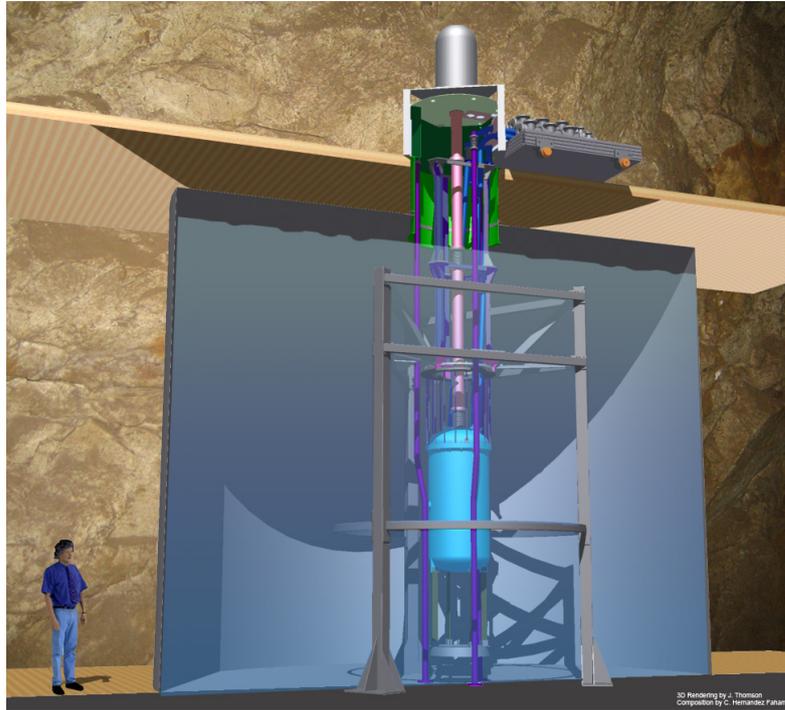
Figure 2.4: 3D CAD rendering of the water tank within the Davis Cavern. The support structure and LUX cryostat are shown, and cabling for the electronics and thermosyphon can be seen passing into the top of the water tank.

2.5 Data Acquisition

2.5.1 Introduction

In any experiment there must be in place a reliable, efficient and effective way to obtain data. The following sections give a brief overview of the LUX data acquisition (DAQ) systems, describing their use in selecting events of interest within the LUX TPC. The raw data in LUX is simply pulses of light measured by the PMTs as a change in voltage over time. However, because of background interactions of γ-rays, internal discharges of light from high voltage grids and radioactive decay on detector surfaces, there is an almost constant stream of light being produced within the xenon. This in principle sounds detrimental for a rare event search. However, the DAQ has the ability to only save events that contain pulses of light that resemble S1s and S2s, which helps remove a large amount of noise and useless data, narrowing the search for WIMP interactions even before any analysis has taken place. The technical details of the DAQ system are beyond the scope of this thesis, but they can be found in [87].

2.5.2 Electronics

Both digital and analog electronics are used in the LUX electronics chain. The PMT electronics must maximise signal to noise for dark matter searches at low energies but also be able to handle high energy calibrations. PMT signals are collected as negative voltages at the PMT base and are passed through a preamplifier and a postamplifer. Some attenuation of pulse height takes place but pulse areas are preserved. Signals from the postamplifier are passed to a digitiser, a triggering system and a discriminator.

2.5.3 Trigger

DAQ thresholds are set low enough to record more than 95% of single electron pulses in each PMT. The LUX trigger system uses pules shape information from the PMT signals and pattern recognition using time, energy and positioning to select events of interest. Overly large signals and invalid geometrical patterns in the PMTs can be removed at the point of data acquisition. Decisions are made on keeping an event within 1 μs of the interaction, and FPGA technology allows real-time processing.

The PMT signals are summed in groups of seven to eight and sent to the trigger to be digitised. There are two filters, S1 and S2, which are adjusted to detect S1 and S2 like pulses respectively. The filter width width range for S1 is 15.625 - 250 ns, sufficient for a typical S1-like pulse which has a FWHM of 80 ns. S2s are wider and more symmetrical than S1s and so have a filter of a longer width, 62.5 ns to 4 μs, and a different shape [88]. The trigger modes can be used individually or combined to look for S1 + S2 signal events. The usual trigger requires that at least two trigger channels (each containing a group of PMTs) have more than 8 detected photo-electrons (phe) within 2 μs.

The maximum possible data acquisition rate is 1.5 kHz - this can acquire the full raw detector event rate irrespective of the trigger, but is generally only used for debugging and examining the detector background offline, where triggered events can be sorted. Normal data acquisition is at a rate of 300 Hz, where only triggered events are downloaded. This is the maximum acquisition rate without deadtime [87].

2.5.4 Pulse Only Digitization

Pulse Only Digitization (POD) is a data acquisition mode optimized for signals that are dominated by long periods of baseline, interspersed with short pulses. POD mode will only acquire pulses whilst suppressing the baseline, by only saving to memory signals that rise above a threshold. POD mode operates on PMT channel pairs. Only one channel of a pair needs to go above threshold for data to be saved, and the partner channel will also be recorded.

Baseline suppression of the DAQ in POD mode is 99.99%, reducing storage size of a typical event by 50 times over the full drift length of LUX, compared to non-POD mode acquisition. The
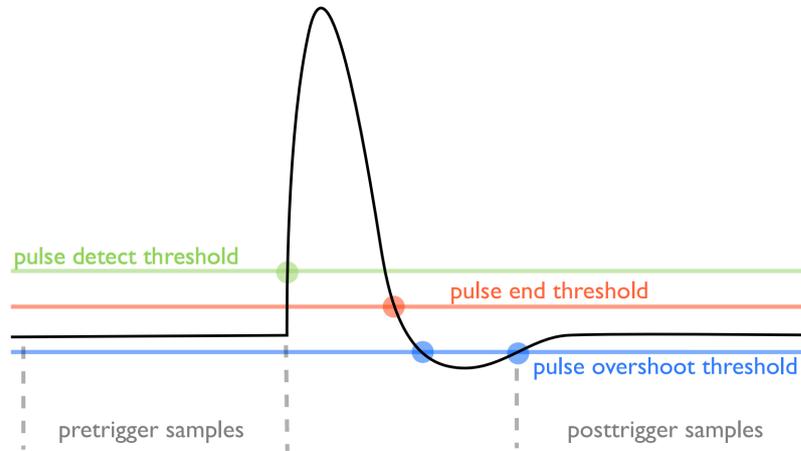
Figure 2.5: Illustration of relevant thresholds in Pulse Only Digitisation [87]. Here, a sample is 10 ns of real time, and the y-axis (not shown) would be in mV.

baseline of each channel is determined by a rolling average of 128 samples (sampling is done every 10 ns). The pulse detection thresholds are measured relative to this baseline. A certain number of samples, defined by pretrigger samples and posttrigger samples, are saved either side of a pulse, see figure 2.5.

2.5.5 Data Reduction and Format

The raw LUX data after digitisation is called a ".dat" file. These files are stored underground as well as transferred to the LUX Primary Mirror, located at Brown University. The following process also takes place underground, where dat files and trigger information are fed into an event builder. The event builder sorts the data into events, includes water veto PMT information and removes any pulses outside of event windows. It sorts each event by channel and then by pulse. The event building takes place underground, and one output file per raw data file is written out in a ".evt" format. These are not stored underground but are transferred to the Primary Mirror. These files can then go on to be processed in the LUX Data Processing Framework (DPF) to produce an ".RQ" file, where RQ stands for reduced quantity. The evt files are backed up onto several other mirrors across the world, and can be processed on any one. An RQ file contains all useful information on an event, such as pulse areas, classifications, timing, etc. See chapter 3.2 for more information on the DPF.

2.6 LUXSim

2.6.1 Introduction

For any physics experiment it is extremely useful to have a reliable simulation framework. An accurate Monte Carlo simulation is crucial for determining design feasibility, background models, detection efficiencies, and can be used to test all sorts of physics interactions. Simulation has been used in LUX extensively for design and then later for analysis and validation of results. In the particle physics community, GEANT4 [89, 90] is most commonly used. GEANT4 is a toolkit that simulates the passage of particles through matter and their interactions with it. It includes tracking, geometry, physics models of hadronic, electromagnetic and optical processes and energy deposits. It can be used to build full and accurate simulations of any detector, with physics ranging from eV to TeV scales. It is implemented in C++, and is widely used in particle physics, nuclear physics, accelerator design, space engineering and medical physics.

LUXSim is a simulation built in the GEANT4 framework with the primary use as simulation of backgrounds (see chapter 4 for more information), most of which involves radioactivity in various parts of the detector. Other purposes include checking optical properties such as light collection, simulating calibrations for real calibration data validation, and simulating signal for efficiency determination (see section 3.9). GEANT4 is is more commonly used o simulate the response of detectors at colliders and fixed target experiments where beams of particles enter the detector; in the case of LUX, a different approach is needed, so LUXSim has the capability of generating radioactive decays in any volume. Here, the basics of LUXSim needed to understand the background simulations are given, but further information can be found in [91].

2.6.2 Geometry and Tracking

Several geometries are available for use in LUXSim, including the full detector system (LUX1.0), a prototype (LUX0.1) and an empty cryostat for studying the water shield without needing to load the complex internal detector structure. The LUX model includes all major components that will affect optical response and backgrounds, including the 122 Hammamatsu R8778 PMTs, reflective PTFE and the wire grids and meshes, with each wire placed individually. The grids can be removed if optical physics is not required (e.g. for energy deposition only simulations).

LUXSim uses a component-centric approach; a custom class known as LUXSimDetectorComponent, which inherits from the GEANT4 class G4PVPlacement, is used to represent physical volumes of the detector. Each detector component can be used as a particle source, allowing easily modelling of radioactivity in a certain material. Furthermore, each component can be set as a sensitive detector with its own record level, which determines how much information is stored. The record levels for normal particles can be set as follows:

1. Record primary particle IDs, energy, direction and position, the record level and size, event ID, component ID and total energy deposition.

2. Record all the above, as well as particle name, track ID, energy, direction, step number, track number, parent ID, position, global time, energy deposition, for all steps with an energy deposition.

3. Record all the above, for every step, even when there is no energy deposition.

4. Record all information about the particle as above, then kill the track.

There is also the option to set record levels for optical photons, which are as follows:

1. Record primary particle IDs, energy, direction and position, the record level and size, event ID, component ID and total photon number, then kill the track.

2. Record the same as above, but do not kill the track.

3. Record all the above, as well as particle name, track ID, energy, direction, step number, track number, parent ID, position, global time and photon wavelength, then kill the track.

In both cases, there is also a record level of 0, which means nothing is recorded in that volume. The lower record levels allow for the reduction of file sizes when running simulations, saving only the most necessary information. The option to kill the tracks is useful for debugging and for when only the flux of particles into a component is needed.

2.6.3   Optical Properties

Several critical optical properties of xenon must be included in LUXSim to ensure reliable and realistic simulations. Data on the refractive index of liquid xenon for wavelengths from 361.2 to 634.9 nm (ranging from 1.41 to 1.37) [92] was combined with a measurement of 1.69 ± 0.02 [93] at the scintillation wavelength of 178 nm using a polynomial fit. Furthermore, a Rayleigh scattering length of 30 cm is used [94], and a absorption length of 100 m at 178 nm. The absorption length is a conservative minimum estimate made using a fit to the convolution of the expected absorption coefficients of impurities expected in the xenon with the scintillation spectrum.

   Reflectivity coefficients of the PTFE are extremely important to light collection. Tests showed the total reflectance at normal incidence to be 70% at $\lambda = 175$ nm in vacuum (expected to be 100% in LXe), and that it has both diffuse and specular components [80]. The reflectivity in LUXSim is modelled as diffuse only as it was found that changing the type of reflection had only a minor effect on light collection efficiency when compared to changing the total reflectivity [91]. Comparison of simulated data to calibration data allowed best fit values for the reflectivity to be calculated as $1.00^{+0}_{-0.0049}$ in the liquid and $0.80^{+0.06}_{-0.04}$ in the gas.

   For all grids except the anode, transparency at 0° is 98%–99%, other than the anode where it is 88%. Because of the little data available for the optical properties of steel at 178 nm, LUXSim is conservative, using a 10% reflectivity for baseline estimates of the light collection efficiency.

The R8778 PMTs are implemented with the fused silica window they feature in real life, so that photon reflection from the quartz is included. The refractive index has been shown to be 1.59 at 178 nm [91].

Finally, also included are optical properties relevant to Cherenkov light present in the water tank veto, so this can be correctly modeled in muon simulations.

2.6.4 Event Generators

A variety of generators are available for LUXSim, including single radio-nuclide decays, full ^{238}U and ^{232}Th chains, an AmBe neutron source, a D-D neutron source, ^{252}Cf fission neutrons and gammas and cosmic muons (including spallation neutrons). These are used to model and map the events expected from backgrounds and detector calibrations. Generators can be placed in any volume, and either a total activity or an activity per unit mass specified. It is also possible to place single particle sources with either a set energy or an energy spectrum if necessary.

2.6.5 NEST

LUXSim uses the NEST (Nobel Element Simulation Technique) library to generate S1 and S2 siganls from energy deposits in liquid xenon. The following sections describe the implementation of NEST within LUXSim.

2.6.5.1 Modelling Recombination

NEST uses a simplified Platzman equation to describe the energy partition between excitation and ionisation:

$$E_{dep} = n_{ex}W_{ex} + N_iW_i = N_i(\alpha W_{ex} + W_i) = W(1 + \alpha)N_i \qquad (2.42)$$

where $\alpha = N_{ex}/N_i$, E_{dep} is the energy deposited in a single interaction, W_{ex} is the work function for exciting atoms, and W_i is the ionization work function. A mean work function W is used for simplicity. N_i and N_{ex} are related to the number of photons and electrons by equations 2.15 and 2.25. NEST calculates the recombination probability r using Birks' Law in the Doke-Birks approach [95].

$$r = \frac{A\dfrac{dE}{dx}}{1 + B\dfrac{dE}{dx}} + C, \quad C = 1 - A/B \qquad (2.43)$$

in which the first term represents the recombination of a wandering ionisation electron with an ion other than its parent, or 'volume' recombination, and the second term C is 'geminate' or Onsager recombination, where the electron recombines with its parent ion. Note how this is dependent on the LET, dE/dx.

47
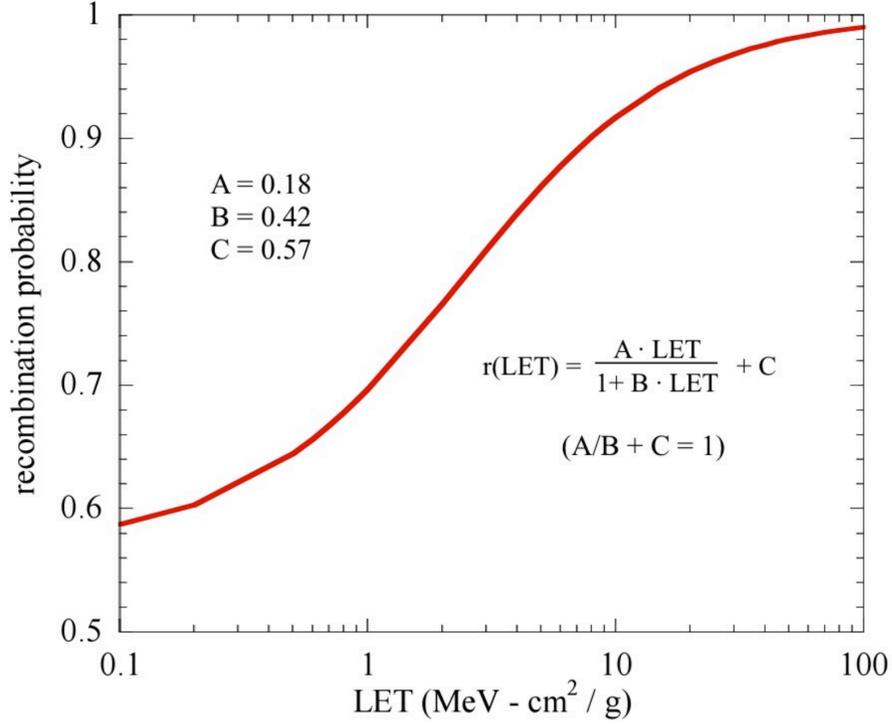
Figure 2.6: The electron recombination probability r implemented in NEST as a function of Linear Energy Transfer (LET). A, B and C are treated as free parameters and constrained using experimental light yield data for γ-rays incident on liquid xenon, and the condition that $A/B + C = 1$, which imposes complete recombination at infinite LET [95].

For tracks shorter than the mean ionisation electron-ion thermalisation distance, 4.6 μm in liquid xenon, a different model is used; the Thomas-Imel box model:

$$r = 1 - \frac{ln(1+\xi)}{\xi}, \quad \xi = \frac{N_i \alpha'}{4a^2 v} \tag{2.44}$$

where α' is a constant dependent on ionisation electron and hole mobilities and the dielectric constant, v is the mean electron velocity and a is a length scale that defines the ionisation density volume. $\alpha'/4a^2 v$ can be estimated at zero electric field using the dielectric constant $\epsilon = 1.96$, a temperature T of 165 K and $a = 4.6$ μm:

$$\frac{\alpha'}{4a^2 v} = \frac{\xi}{N_i} = \frac{e^2}{a\epsilon kT} = 0.14 \tag{2.45}$$

For events where a particle multiple scatters, each track is analysed separately to decide whether to apply the Birks' or Thomas-Imel model. The use of the Thomas-Imel box model at small track lengths successfully reproduces the turn-over in scintillation yield observed at low energies [96], shown in figure 2.2. This is due to low energy particles always being in the Thomas-Imel regime, where recombination probability depends on energy (via N_i) instead of dE/dx. A, B, C and ξ/N_i were all treated as free parameters in order to match the NEST model to all known experimental data. There was some inherent difficulty in using all data available, as sometimes error bars

were not available, there was more than one data point for certain γ energies and yields are also quoted as relative to either 122, 511 or 662 keV depending on the calibration source. It was decided to translate all data to 122 and 511 keV, except for low energy measurements in [96], which used a 59.5 keV ^{241}Am line. Data were treated as equally as possible, and in general, agreement was good between datasets. The fit was done using minimisation of the mean-squared of the residuals, and the results were $A = 0.18, B = 0.42, C = 0.57$ and $\xi/N_i = 0.19$, which was close to the estimate of 0.14. ξ/N_i was obtained using low energy (< 15 keV) data only, as this is where the Thomas-Imel model dominates. Parameters held fixed during the fit were $W = 13.7$ eV, $\alpha = 0.06$ and $a = 4.6$ μm.

However, new results on low energy electron recoils from Compton scattering experiments contradicted these results [97, 98], showing an overestimation of the yield by $\sim 20\%$. This was fixed by the insertion of a 2 V/cm field into the power law fit for the field dependence of r; this is the approximate strength of the screened electric field of a xenon nucleus at 4.6 μm. This reduced ξ/N_i to 0.05 [99], which agreed with new results, and appeared to be a more physically accurate representation as it includes the field from the xenon nucleus, which had previously been ignored.

Individual energies and ranges for secondary electrons are tracked in GEANT4 to determine recombination, and so variation in the tracks due to Monte Carlo fluctuation causes a variation in the light yield.

For electrons, it proves more difficult to obtain a match between the model and data. The model as it stands as described produces a yield of 57.4 ph/keV for a 1 MeV electron, whilst the best experimental data gives 46.4 ph/keV [78]. Manipulation of the recombination processes can reconcile the model with this data point; if non-geminate recombination is assumed to cease for a minimally ionising particle, the yield moves down to 43.1 ph/keV. This is controversial, as the thermalisation distance of 4.6 μm is much larger than the Onsager radius, meaning that geminate recombination should not be dominant. However, it is possible that an ionised electron could either be re-attracted to its parent ion if no other ion is available for recombination, or that it may become trapped within its parent ion's Coulomb field if it is repelled by nearby electron clouds. Including some initial excitation by increasing α brings the yield to 46.7 ph/keV, even more in agreement with experimental data.

2.6.5.2 Electric Field Quenching

As the strength of the applied electric field increases, electrons are less likely to recombine with ions as they drift out of the liquid. This results in a quenching of scintillation light, and so will decrease S1 light whilst increasing S2 light. In order to parameterise recombination at different electric fields, various datasets were used at five differing field strengths to collect values of ξ/N_i. A power law was then used:

$$\frac{\xi}{N_i} \propto 0.057 E^{-0.12} \tag{2.46}$$
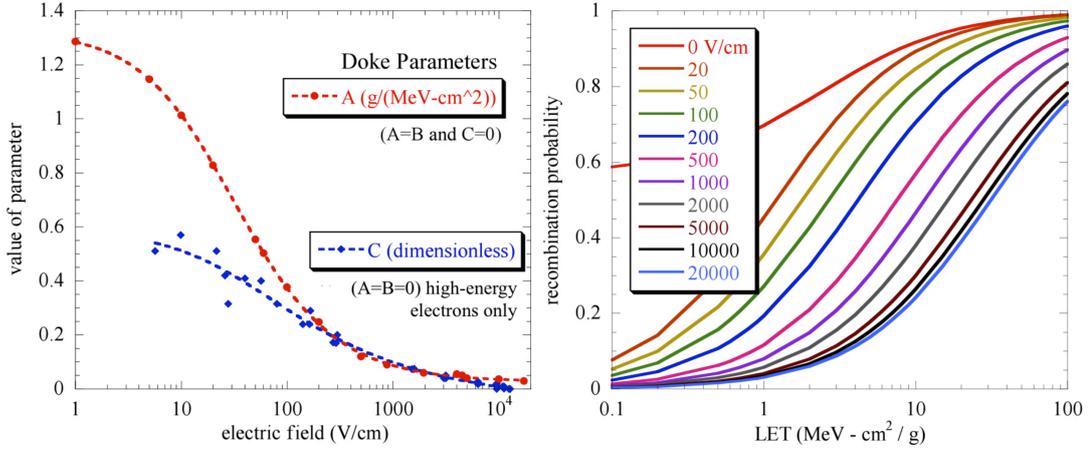
Figure 2.7: Left: The Doke parameters A, B and C as functions of electric field. Red shows the result of matching γ data, with most weight given to 122 keV. Blue is for relativistic electrons using 1 MeV electron data. A power law fit is used to interpolate between discrete fields. Right: Recombination probability r as a function of LET at different field strengths, based on the red curve on the left (i.e. $C = 0$) [95].

where E is in V/cm. To find the Doke parameters (A, B and C in equation 2.43), results for 122 keV from ^{57}Co sources were used as the largest dataset was available for this energy. For data taken in the presence of an electric field, the best fit gave negligible C, suggesting that geminate recombination is heavily suppressed. This is expected as the external electric field is far stronger than the Coulomb field of a Xe atom. It was found that there are two distinct cases; the first being for γ data, where $C = 0$ and $A = B$, the second being relativistic electrons, where $A = B = 0$, see figure 2.7.

The cross over between the Doke-Birks and Thomas-Imel regime also needs to be given an electric field dependence to avoid a discontinuity. A power law was also used, giving a cross over distance in μm as $69/\sqrt{E}$.

2.6.5.3 Nuclear Recoil Quenching

The only modification made to the above framework in NEST for nuclear recoils is the addition of the Lindhard factor to equation 2.42:

$$E_{dep}L = (1 + \alpha)N_i W \tag{2.47}$$

Quenching was previously described in section 2.3.4. NEST uses the Lindhard model with a modification from Hitachi [100]. The methodology followed is similar to that in [101], but starts with an ER-motivated Thomas Imel parameter at zero electric field instead of one from fits to NR data. Predictions were found to be consistent with some recent measurements such as Plante et al [102], but was in conflict with others such as Aprile et al [103]. A later modification to NEST changed α from 0.06 to 1.00 for nuclear recoils. This resulted in excellent agreement with experimental data, again from Plante et al [102], but also Horn et al [104]. The motivation for

using this higher ratio comes from several places. NR yields are observed to change slower with electric field than ER, implying that recombination is less dominant, and fluctuations due to NR recombination are smaller. Furthermore, when fitting to data if α is left as a free parameter the best fit is close to 1. Finally, almost the same Thomas-Imel recombination probability parameter described by a power law as used for ER can be used if α is 1.

NEST is continually updated with calibration data from LUX and tested. Its predictions have been shown to consistently agree with data and it has been adopted for use by other dark matter experiments. Further information on NEST can be found in references [95, 99, 105].

LUX Signal Identification

In rare event searches such as the LUX WIMP search, it is essential to have a system of signal identification with high efficiency, accuracy and background rejection. This chapter describes the LUX data processing system and algorithms that are used to select candidate WIMP search events. Much of my work for LUX has been on the development of and improving its pulse finding and classification algorithms with an aim to increase their efficiency over all energies, but with a focus on low energies where difficulties with small signals can arise. In this chapter I give an overview of the LUX Data Processing Framework, discuss the improvements and validation I performed on the pulse finding and classification system, and finally present the nuclear recoil efficiency curves I calculated and provided for the run 3 reanalysis, which were used for limit setting.

3.1 Some Definitions

To avoid repetition in the coming chapter, some important quantities/definitions in LUX analysis are given in table 3.1.

Table 3.1: Important definitions for use within LUX analysis.

Name	Definition
Run 3	data taken between April and August 2013, published in [46, 106].
Run 4	data taken between September 2014 and May 2016, published in [36].
NR	nuclear recoil
ER	electron recoil
Golden event	event containing a single scatter, 1 S1 and 1 S2
phe	area of a pulse in photoelectrons (emitted from PMT photocathode)
phd	area of a pulse in detected photons (originally striking photocathode)
Bad area	total area of the waveform not contained in the S1 and S2

The basic selection of a golden event within LUX data has a couple of caveats. The event needs to have 1 S1 followed by 1 S2, but any S2s before the first S1 or any S1s after the S2 are also allowed, as they are coincidental with the main scatter.
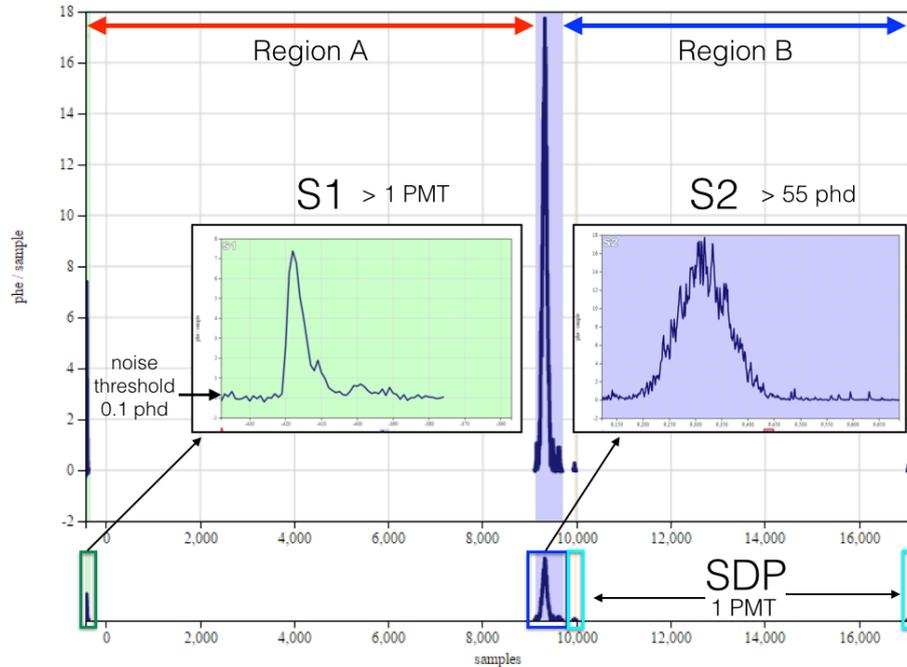
Figure 3.1: A single scatter event from LUX data. The S1 and S2 are shown with their minimum requirements of at least 2 PMTs for the S1 and the minimum S2 size of 55 phd. This event also contains two single detected photons, denoted SDP, after the S2. The bad area of this event would be just the area of those two photons as the rest of the non-signal waveform is empty. Quantities important to the pulse finder (Region A/B, noise threshold—see section 3.3.1) are also shown.

The differentiation between phd and phe in table 3.1 is necessary due to the non-zero probability of one photon ionising two photoelectrons from a PMT photocathode. For PMT gains—the conversion from voltage to pulse area in phe—these were initially assumed to be equivalent, leading to an overestimation of pulse sizes. Once this was discovered, it was measured in Hammamatsu R8778, R8520, and R11410 PMTs and accounted for within the LUX PMT gains. The probability of 2 phe emission was found to be between 18% and 24% depending on the tube and measurement method [107]. This is a vacuum ultra-violet (VUV) effect.

3.2   LUX Data Processing

The Data Processing Framework is a version controlled collection of C++, Matlab, Python and XML code, with a modular format, used to identify, parameterise and characterise triggered events. Figure 3.2 shows a simple version of the DPF. There can be several versions of each
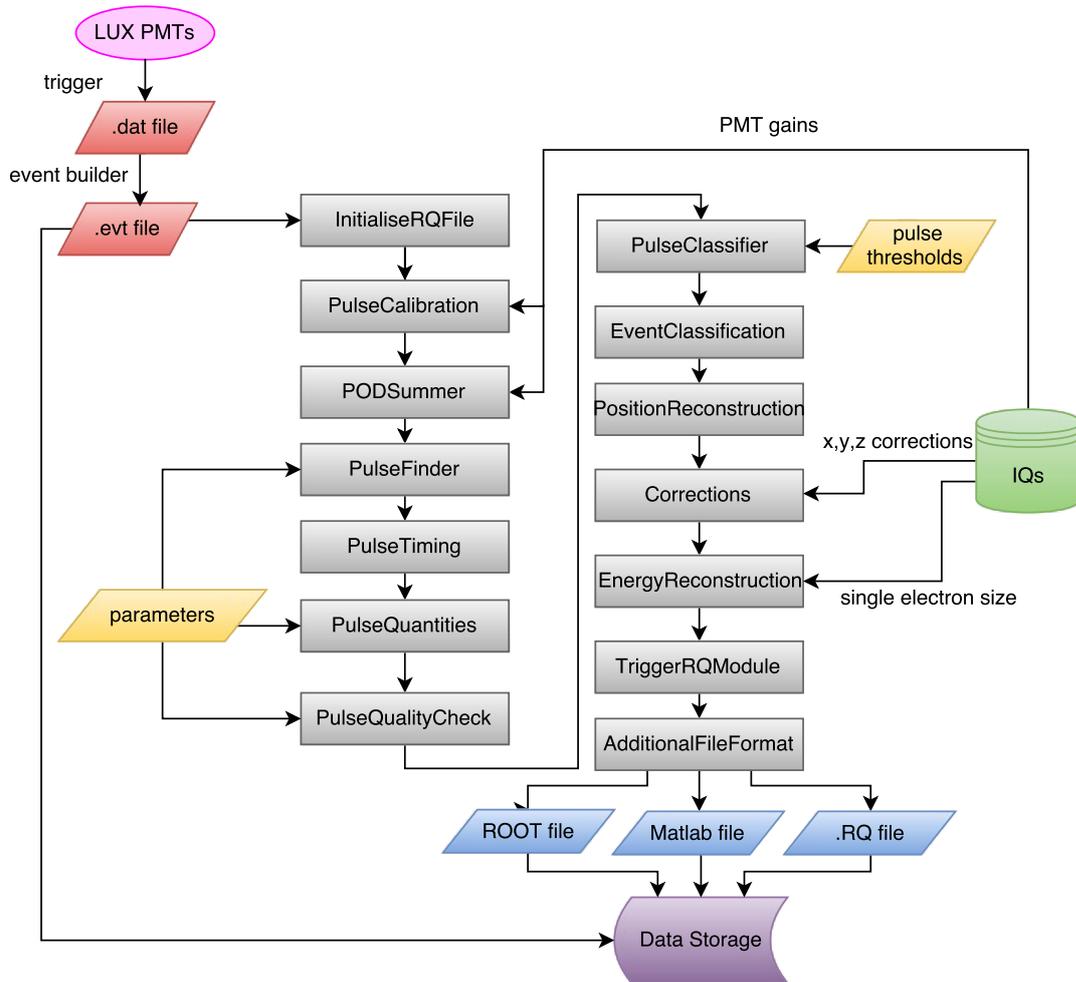
Figure 3.2: Flowchart of the LUX Data Processing Framework, demonstrating its modular structure. The RQ (reduced quantity) file, initially described in section 2.5.5 is continuously updated by each module, and quantities from any previous module are accessible to the next. The file is converted to ROOT and Matlab format ready for analysis at the end. The IQs (Important Quantities) are obtained from previous data analysis and are stored in a MySql database.

module type (in grey) that can be easily swapped using an XML data processing settings file. The structure allows flexibility; easy addition of new modules as well as the possibility to develop alternatives to existing modules to be swapped in and out to test performance.

Additional modules to those in figure 3.2 do exist for more specific data (for example a module to resolve close in time S1 pulses for krypton calibration data, where events contains two S1s), but the following describes the basic modules required for any LUX data processing:

- **PulseCalibration:** takes PMT gains as an input and converts voltages recorded by the PMTs to photoelectrons (phe) per sample, so all other modules can use pulse areas in detected photoelectrons.
- **PODSummer:** sums the calibrated individual PMT channel waveforms to create an overall

waveform for each event.

- **PulseFinder:** identifies and returns the start and end times of generally the largest 10 pulses in the event (see section 3.3). The number of pulses saved can be altered, but 10 is the default. Each subsequent module that returns pulse quantities returns in them in an array per event, with one array entry for each pulse.
- **PulseTiming:** returns the times at which each pulse reached 1%, 10%, 50%, 90%, 99% of its total area, quantities useful for pulse classification as they describe the shape of the pulse.
- **PulseQuantities:** returns information on pulses including its area, its asymmetry between the top and bottom PMT arrays, a measure of its promptness, etc.
- **PulseQualityCheck:** runs some basic checks on the pulse to determine its validity.
- **PulseClassifier:** described in detail in section 3.4, returns a number classifying a pulse of one of several types, most importantly S1 and S2.
- **EventClassification:** determines whether an event is a single scatter, multiple scatter etc. Section 3.5 contains more about this module.
- **PositionReconstruction:** uses algorithms to return an x,y,z position of a pulse.
- **Corrections:** corrects the position and areas of pulses using IQs obtained from krypton calibrations, see section 5.1.1.
- **EnergyReconstruction:** uses single electron size and extraction efficiency to estimate the energy of the recoil in keV.
- **TriggerRQModule:** stores trigger information such as the trigger type and which channel triggered.
- **AdditionalFileFormat:** converts the final RQ file to ROOT and Matlab formats.

The final outputs of the DPF are ROOT and Matlab files that can be used for data analysis. The file formats allow analysts to use whichever programming language they are more comfortable with.

3.3   Pulse Finding

3.3.1   Pulse Finding Algorithm

The pulse finder forms the core of event selection in LUX, identifying individual pulses of light and passing them for further analysis. This is a key module that directly impacts the energy dependent efficiency of the experiment and overall science reach, so must be consistent, high performance and well-validated. The LUX pulse finder used for all main data processing is called TransparentRubiksCube (TRC), and is a simple algorithm using summed areas and thresholds. The input parameters are shown in table 3.3.

TRC is a Matlab module with a C++ mex file containing the finding algorithm, which is called multiple times as the module runs. The pulse finding process is as follows:

Table 3.2: Input parameters for the TRC pulse finder and their default values that have been selected to give the highest pulse finding efficiency.

Name	Value	Units	Purpose
fullBoxSamples	400	samples	starting length of the pulse filter
noiseThre	0.1	phd/sample	noise threshold
maximumGap	50	samples	length of time pulse must be below threshold to cut it off
nLookAhead	1	samples	how far ahead to average a pulse
nLookBehind	1	samples	how far behind to average a pulse
extendPulse	30	samples	number of samples to add on to the end of a pulse*
max_num_pulses	10	-	number of pulses to keep

1. A sliding boxcar filter is passed over the full event waveform, finding the region that maximises the enclosed area. The box starts with the length given by *fullBoxSamples*.

2. The maximum point in the boxcar is found, and the waveform is smoothed over the range of samples spanned by *nLookAhead* and *nLookBehind*. The smoothing occurs sample by sample, moving backwards along the pulse. The start time of the pulse is set at the end of a quiet time defined by the average height falling below *noiseThre* for the length of time given by *maximumGap*.

3. This is repeated moving forwards in time from the maximum, setting the end time as the beginning of the quiet time. Both the start and end time can lie outside of the original boxcar.

4. Pulses are then only kept if they are above 3 samples long, or if their average phd per sample is greater than *noiseThre*.

5. The pulse that has been found is zeroed, and the previous steps run again, in order to find the second largest pulse in the event.

6. After the two largest pulses are found, the waveform is split into region A and B, where A is the region before the start of the latest in time of the two pulses and B is the region afterwards. Region A is searched first to avoid missing an S1 due to the 10 pulse limit. Without the search regions, an S1 could happen in an event with lots of large pulses after the S2, e.g. in a multiple scatter.

7. Region A is searched until either reaching *max_num_pulses* or there are no pulses left to be found. If *max_num_pulses* has not been reached, region B is searched.

Table 3.3: Output RQs produced by the TRC pulse finder (after the updates described in section 3.3.2).

Name	Type	Purpose
pulse_start_samples	integer array	sample where the pules starts
pulse_end_samples	integer array	sample where the pulse ends
index_kept_sumpods	boolean	1 if a pulse is found, 0 if not
num_pulses_found	integer	total number of pulses found in the event

3.3.2 Improvements to the Pulse Finder

Original versions of TRC produced more RQs, including many now produced by the PulseTiming and PulseQuantities module. An effort was made to extensively check, validate and clean up the module and alter it to only return pulse start and end times. This was to improve the functionality and transparency of this key module, as well as introducing compatibility with BlackBox, an alternative pulse finding algorithm. This compatibility allowed the pulse finders to be swapped to compare performance.

A new Matlab module called PulseTiming_BasicSet was written to calculate the area fractional timings (AFTs) defined in table 3.4, and any other quantities were recalculated within PulseQuantities_MinimumSet. Whilst comparing the output of the newly calculated RQs with the original output of TRC, several bugs were identified and fixed or removed. The first of these involved the AFTs; if a pulse did not reach 90% or 99% of its area before its end, the RQs *aft_t90_samples* and *aft_t99_samples* were left as 0, and then later set to the pulse start by error checking. Pulse classification uses an asymmetry ratio calculated with the AFTs, with *aft_t99_samples - aft_t1_samples* as the denominator. As these could be equal due to the bug, this gave an infinite ratio and a classification of 5, or unknown (see section 3.4). This happened mainly to single photoelectron pulses but also a few S1s, and these were classified correctly after the recalculation of the AFTs within the new BasicSet module. The fix was to set the AFTs above 50% to the pulse end rather than the start during error checking. Secondly, the height fractional timing (HFT) calculations were happening during the pulse finding algorithm before the start and end times were defined, meaning only the original boxcar was searched. This meant the whole pulse was not searched and the algorithm did not always start from the true maximum. As the HFT calculation now happens outside of TRC, this is no longer a problem.

Another fix occurred during the redefinition of the pre and post pulse areas; these are the integrated areas contained within the 50 samples either side of a pulse. Because of the zeroing of found pules in TRC, if a pulse was close to one previously found, its pre and post pulse areas contained zeroes from this pulse. Furthermore, if the pulse was close to the region A/B boundary then area beyond this was also not included. Now, these RQs contain the area of the 50 samples either side of the pulse no matter what lies in those regions.

Finally, the region A/B search logic originally used the maximum of a now obsolete RQ called *aft_box_start_samples*, which was initialised to 0. If the first two pulses found had a negative sample number as their start value, then the maximum of the array was the initialised 0, and the region searched was incorrect. Using *pulse_start_samples* instead prevented this problem as the array is initialised to $-999,999$.

3.3.3 Multiple Scatter Treatment

Whilst a rarely interacting WIMP is never expected to multiple scatter within the detector volume, a neutron is highly likely to. Therefore, identification of multiple scatters at low, WIMP-search

energies is essential for neutron-induced NR background rejection. At high energy, multiple scatter identification is still important, as it affects the accuracy of what needs to be a high precision background model. Such a model allows a mapping of backgrounds that can be extended to low energies to determine the expectation in the WIMP-search energy range. Therefore, the pulse finder should be able to identify multiple scatters across all energies.

The nature of the TRC algorithm leads to a drop in performance at higher energies. For events with large S2s, there are many occasions where the pulses do not return to below *noiseThre* for a sufficient time (i.e. the length of *maximumGap*) before the next pulse starts. This means the pulse end time is not set to a point between the pulses, and they become merged. To rectify this, a new algorithm was added to perusePeeks.C, the code containing the pulse finding algorithm. This addition is described below.

1. The pulse finder runs as given in section 3.3.1 to the end of stage 3. At this point, if the length of the pulse is greater than 600 samples, the splitting algorithm proceeds.

2. Starting at the maximum point in the pulse, the waveform is first smoothed forwards in time. The averaging is much greater than in the original pulse finding algorithm to account for the large pulse size and to prevent small fluctuations causing too much splitting; it uses 24+*nLookAhead* (i.e. 25) samples.

3. When the pulse average reaches 40% of the maximum height, the sample number this occurs at is recorded.

4. If the average then rises, so its average is greater than the average of the previous sample, and this average reaches 10% of the maximum height, the sample number of this is also recorded.

5. The pulse average must continue rising for 50 samples to ensure it is a real second pulse and not just a fluctuation. If this is satisfied, the region between the falling 40% point and the rising 10% is searched and the minimum located. The pulse end time is set to this minimum.

6. The algorithm is repeated but looking backwards, averaging over 24+*nLookBehind* samples, in case there is a preceding pulse. If all the requirements are matched as before, the pulse start time is modified, cutting out the largest peak.

The minimum length of 600 samples was chosen by examining the distribution of pulse length. Above this length, the number of S2s is over an order of magnitude larger than pulses classified as unknown, and over two orders of magnitude larger than the number of S1s. Furthermore, S1s have a very clear shape and are unaffected by the algorithm, unless they are merged with another pulse. The minimum length requirement also protects low energy S2s being split unnecessarily if they are particularly jagged.

The threshold values of 40% and 10% and the pulse rise time minimum length of 50 samples were chosen by examining high energy multiple scatters by eye. Several sensible combinations of values were tested, and the final choice gave the best results, which are shown in section 3.7.1.

59

3.3.4 Noise Threshold

The *noiseThre* parameter in TRC was initially set at 0.15 phd/sample, and had been chosen almost arbitrarily from looking at noise in waveforms. After PMT gains were recalculated, pulse areas decreased and handscanning of events with less than *max_num_pulses* identified there were small SPEs sometimes being missed by the pulse finder, despite the arrays being empty. A noise threshold that is set too high has two effects; it causes pulse merging as the baseline does not dip below the threshold for long enough between pulses, and missed SPEs. The second occurs because the pulse finder makes a decision on keeping a pulse based on the pulse passing one of two cuts:

$$pulse\_end\_samples - pulse\_start\_samples > 3 \tag{3.1}$$

or

$$\frac{pulse\_area\_phe}{pulse\_end\_samples - pulse\_start\_samples} > noiseThre \tag{3.2}$$

The second cut was modified to also require the pulse length to be greater than 1 sample (10 ns) to prevent spikes being found as pulses. If an SPE is very narrow and has a very small area, it may not pass the cuts if *noiseThre* is set too high. Data processed with 0.05 <*noiseThre*< 0.15 was obtained and pulse areas plotted between 0 and 2 phd, see figure 3.3. A noise peak is clearly identifiable below 0.2 phd for the lower thresholds. A decision was made to reduce *noiseThre* to 0.1, as this was a sufficient decrease to identify small SPEs whilst not introducing a noise peak.

This issue arose again for run 4 after a large handscan identified that pulses were being cut off too early, this time suggesting that the noise threshold was too low. Changes in the electric field and PMT gains acquired from calibrations had changed pulse sizes between the runs. Data was again obtained processed with 0.05 <*noiseThre*< 0.20, but the decision was made to put the parameter back to 0.15. This was an optimisation between reducing the over-splitting of pulses observed whilst preserving SPEs. Looking at the number of pulses classified as 'unknown' gave a good measure of the amount of noise being picked up as pulses, and this decreased as *noiseThre* increased, reached a minimum at 0.18 and then began to increase again as pulses are merged together. However, the number of S1s per event begins to rapidly increase after *noiseThre*> 0.15 as pulses are oversplit and fake S1s identified, so returning to *noiseThre*=0.15, whilst not optimal for removing noise, seemed the safest option.

3.3.5 S1 Priority System

A modification was made to the pulse finder to avoid the possibility of an S1 not being found due to the 10 pulse limit—an issue that became more problematic with the removal of the low energy S1 threshold, see section 3.6.1. The algorithm goes ahead if 10 pulses have been found and the remaining event area is less than the bad area cut maximum for low energy events, which is 80 phd. The bad area cut is defined as:

$$full\_evt\_area\_phe - (pulse\_area\_phe[S1] + pulse\_area\_phe[S2]) < 80 \text{ phd} \tag{3.3}$$
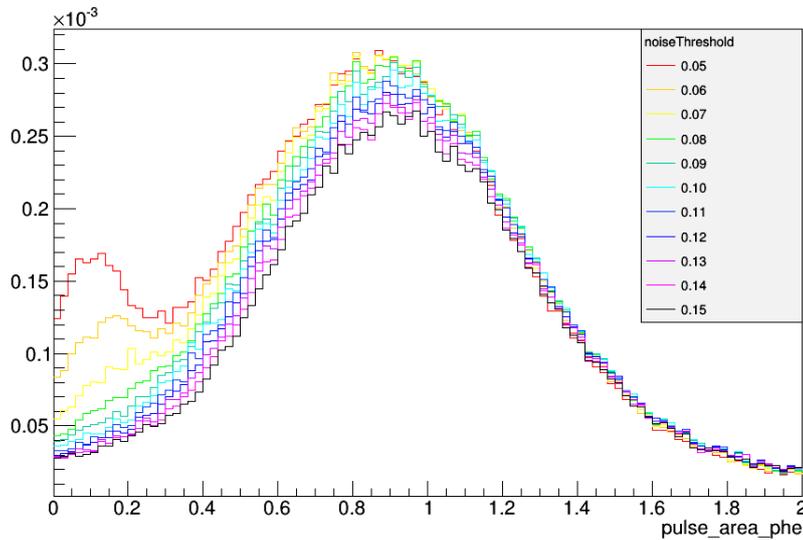
Figure 3.3: Distribution of SPE size within run 3 data. Different colours show different values of *noiseThre*, and a noise peak gradually begins to appear below 0.2 phd as the threshold is decreased.

The algorithm runs as follows:

1. Flag pulse if the pulse length is greater than 50 samples and its area is greater than 5 phd - it is not a S1-like shape.
2. Flag pulse if only 1 PMT passed the *peak_height_phe_per_sample* > 0.09 requirement.
3. If the last pulse found has been flagged, run pulse finding algorithm again in region A.
4. If a pulse is found, check whether *peak_height_phe_per_sample* > 0.09 was satisfied by more than 1 PMT.
5. Check if the pulse length is less than 100 samples. If so, prioritise the pulse.
6. Repeat until a pulse is prioritised
7. When a pulse has been prioritised, do the usual quality checks, and if it passes, replace the start and end times of the last pulse found with those of the prioritised one

The first two steps check whether the last pulse found looks like a non-signal pulse we would be willing to throw away for an S1. Step 4 applies the 2 PMT coincidence requirement for S1s (see section 3.4.3), and step 5 uses the assumption that low energy S1s are very short and never over 100 samples. This prevents the prioritisation of long noise tails. This algorithm was tested alongside other changes, the results are presented in section 3.7.2. In summary, from ~55,000 golden events in tritium calibration data, 54 new golden events with very small S1s were gained due to the priority system, and a total of 219 golden events were no longer found to be golden, as they had a second S1 present. These are obviously small changes (<1%) in a large sample, but as the new events are all low energy, any improvement is important for low WIMP mass signal sensitivity.

3.3.6 Empty Pods

A handscan of Americium-Beryllium (AmBe) neutron calibration data (detailed in section 3.9.5) identified that the only failure mode of the DPF to identify golden events was a misclassification of S1s due to the merging of an S1 and a SPE following shortly afterwards. This feature is not uncommon for PMTs due to afterpulsing, ion feedback and similar effects. This behaviour was noticed in the pulse finder previously; due to the *maximumGap* parameter, if there was not 50 samples (500 ns) of quiet time in between two pulses, they would be merged, even if there were 'empty pods' in between. Here, empty pods is used to refer to an actual gap in the waveform, where no PMT signal above noise levels was recorded. As a gap in the waveform almost always signifies two separate pulses, the pulse finder was modified to split pulses under two conditions; the usual quiet time using *noiseThre* (< *maximumGap*), and if the waveform stayed completely empty for 10 samples. This was tested alongside other modifications, and the results can be found in section 3.7.2.

3.4 Pulse Classification

3.4.1 Introduction

Having found pulses, it is then required that they are classified to identify their nature (e.g. S1, S2) for further analysis. Misclassification can lead to the loss of true signal events, the introduction of false signal events and incorrect characterisation of backgrounds. Therefore, as with pulse finding, classification must be made with high accuracy in rare event searches. The following sections describe the pulse classification system used within the LUX DPF.

3.4.2 Quantities used in Pulse Classification

Table 3.4 lists the quantities, all calculated within the DPF, that are required to classify a pulse. The main LUX pulse classifier is called MultiDimensional, as it performs a multidimensional analysis using a combination of cuts using these quantities or new ones calculated from them. Most of these utilise the differences in pulse shape.

3.4.3 S1

An S1 signal is scintillation within the liquid caused by an electron or nuclear recoil. Its main characteristic is its sharp appearance due to the prompt nature of the scintillation. The first of the S1 classification cuts is:

$$\frac{s2filter\_max\_area\_diff}{pulse\_area\_phe} < \begin{cases} 0.4, & 100 < pulse\_area\_phe < 500 \\ 0.01 + (-0.5 + e^{-1.2 \times pulse\_area\_phe}), & \text{otherwise} \end{cases}$$

$$(3.4)$$

Table 3.4: Definition and use of quantities in pulse classification.

Name	Definition
pulse_area_phe	Full pulse area
aft_t1_samples	Sample where area reaches 1% of total
aft_t10_samples	Sample where area reaches 10% of total
aft_t25_samples	Sample where area reaches 25% of total
aft_t50_samples	Sample where area reaches 50% of total
aft_t75_samples	Sample where area reaches 75% of total
aft_t99_samples	Sample where area reaches 99% of total
prompt_fraction	Ratio of 10 sample window centered on aft_t1_samples to whole area
prompt_fraction_t10	Ratio of 10 sample window centered on aft_t10_samples to whole area
s2filter_max_area_diff	Difference of two filters 50 and 200 samples long each maximising area
top_bottom_asymmetry	Ratio of difference in areas in top and bottom PMTs to total
pulse_height_phe_per_sample	Average height of the pulse over its length
peak_height_phe_per_sample	Maximum height over pulse, per PMT
skinny_peak_area_phe	Maximum area 10 samples, per PMT

This cut is seen as the black line in figure 3.4(a), which consists of ER calibration data pulses, with various pulse types shown in different colours. There is an additional allowance window between 100 and 500 phd where the value can rise to 0.4 - this is in place because of the tendency of the pulse finder to merge the very close in time S1s that occur during a krypton calibration, see section 5.1.1. Next, the pulses must satisfy:

$$prompt\_fraction\_t10 < 0.56 + (-1.2e^{-0.26 \times pulse\_area\_phe + 0.2}) \tag{3.5}$$

seen in figure 3.4(c), and the following cut on area fractional timing ratios:

$$\frac{aft\_t50\_samples \text{ - } aft\_t1\_samples}{aft\_t99\_samples \text{ - } aft\_t1\_samples} < \begin{cases} 0.34, & 100 < pulse\_area\_phe < 500 \\ e^{-0.025 \times pulse\_area\_phe} + 0.29 - \\ 0.8e^{\log_{10}(pulse\_area\_phe) - 6}, & \text{otherwise} \end{cases} \tag{3.6}$$

This cut is shown in figure 3.4(g). Furthermore, as S1s are found to have more light in the bottom array due to reflection at the liquid-gas interface, bounds are placed on *top_bottom_asymmetry*:

$$top\_bottom\_asymmetry > -0.55 - (0.5 \times \log(pulse\_area\_phe))^{-0.7} \tag{3.7}$$

$$top\_bottom\_asymmetry < -0.35 + (0.3 \times \log(pulse\_area\_phe))^{-2.2} \tag{3.8}$$

see figure 3.4(e).

Finally, a two-fold PMT coincidence requirement is applied. For a pulse to be classified as an S1, at least two PMTs must satisfy:

$$peak\_height\_phe\_per\_sample > 0.09 \text{ and } skinny\_peak\_area\_phe > 0.3 \tag{3.9}$$

Furthermore, if the S1 pulse comprises of only two PMT channels, they must not be partner channels. Each PMT is electronically paired with one other PMT, and a signal on one can induce

a false signal on the other. These requirements ensure that only genuine scintillation photons are counted as an S1, rather than any coincident light in the detector.

An example of an S1 is shown in figure 3.5(a).

3.4.4 Single Photoelectron

A single photoelectron, or SPE, as its name suggests, is the signal generated from the impingement of one photon on a PMT's photocathode. An example of an SPE is shown in figure 3.5(c). An SPE and an S1 must pass almost all the same cuts. For an SPE, there is no requirement on *skinny_peak_area_phe*, and only one PMT must satisfy the cut on *peak_height_phe_per_sample*.

3.4.5 S2

The S2 signal is generated by the scintillation of the xenon gas when electrons extracted from the liquid are accelerated. An S2 and a SE must pass all the same cuts, as the S2 is simply multiple SEs. They are distinguished by a threshold in pulse area at which the SE becomes defined as an S2. S2s and SEs must satisfy:

$$prompt\_fraction > -0.015 - 0.6 e^{-0.3 \times pulse\_area\_phe} \tag{3.10}$$

and

$$prompt\_fraction < \begin{cases} 0.2, & pulse\_area\_phe > 1000 \\ 0.7 \left(1 - \dfrac{0.9}{1 + e^{-3(\log_{10}(pulse\_area\_phe) - 2)}} \right) & \text{otherwise} \end{cases} \tag{3.11}$$

as shown in figure 3.4(d). The population of pulses included by the change in the cut after 1000 phd tend to be merged multiple S2s.

S2s and SEs must also pass:

$$\frac{s2filter\_max\_area\_diff}{pulse\_area\_phe} > \begin{cases} 2E - 4, & pulse\_area\_phe < 1 \\ 0.12 - 1.35 \log(pulse\_area\_phe)^{-1.5}, & \text{otherwise} \end{cases} \tag{3.12}$$

and

$$\frac{s2filter\_max\_area\_diff}{pulse\_area\_phe} < 1 \tag{3.13}$$

which is shown in figure 3.4(b). This cut helps to remove large S1s that become more symmetrical at high energies. Furthermore,

$$top\_bottom\_asymmetry > 0.05 - \frac{0.5}{1 + e^{-2(pulse\_area\_phe \text{ - } 4.8)}} - (0.4 \log(0.5 \times pulse\_area\_phe))^{-1.5} \tag{3.14}$$

$$top\_bottom\_asymmetry < 0.15 \times \left(1 - \frac{0.5}{1 + e^{-2(pulse\_area\_phe \text{ - } 5)}} \right) + (0.5 \log(pulse\_area\_phe))^{-0.5} \tag{3.15}$$

(a) S1 cut on the S2 filter

(b) S2 cut on the S2 filter

(c) S1 prompt cut

(d) S2 prompt cut

(e) S1 top/bottom cut

(f) S2 top/bottom cut

(g) Area fractional timing cut

(h) Pulse height cut
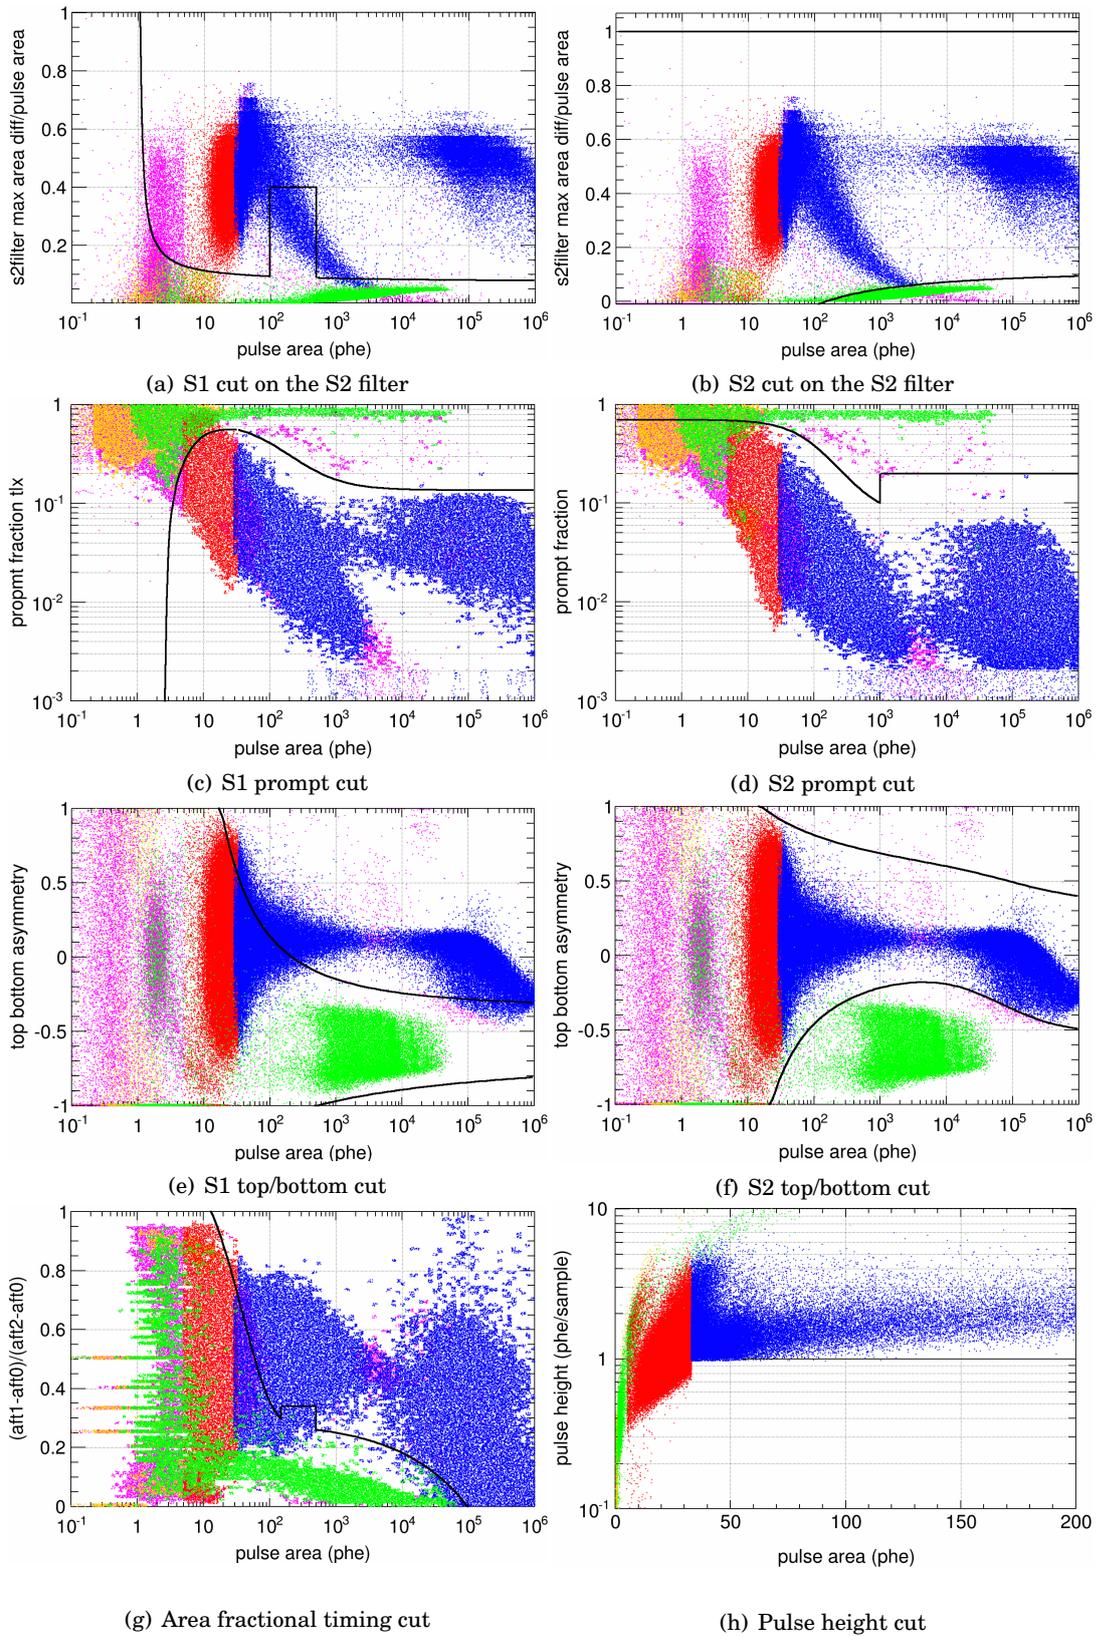
Figure 3.4: Classification quantities against pulse area in ER calibration data. Shown (after full classification) are: S1s (green), S2s (blue), SPEs (orange), SEs (red) and other (magenta). Lines show the cut region. Square regions in the top and bottom left plots indicate an allowance for merged krypton S1s.

seen in figure 3.4(f) requires that S2s are more evenly distributed between top and bottom array than S1s, with slightly more light in the top array as they happen in the gas.

The final requirement for an S2 is an area greater than the current S2 threshold. This has varied throughout LUX's lifetime, and is discussed in section 3.6.2. An example of an S2 is shown in figure 3.5(b).

3.4.6 Single Electron

A single electron (SE) pulse is the signal from a single extracted electron entering the gas phase and causing scintillation, an example is shown in figure 3.5(d). The mean single electron pulse size is an important quantity used for energy reconstruction. It also fluctuates with time, but generally ranges from 20 to 30 detected photoelectrons. An SE must pass all the same cuts as the S2, and the two are distinguished by a threshold, nominally 100 phd.

3.4.7 Merged

In order to aid single scatter selection, a new pulse category was added for pulses that the pulse finder was unable to separate. This happened occasionally with large, close in time S2s and gas events, discussed in section 3.5. An example of a gas event pulse that was found as a single pulse can be seen in figure 3.6(a). These pulses are selected based on their shape and must satisfy:

$$\frac{t_{75}}{0.75 \times t_{99}} < 0.8 \tag{3.16}$$

or

$$\frac{t_{25}}{0.25 \times t_{99}} > 1.97 \frac{t_{75}}{0.75 \times t_{99}} - 1.6 \tag{3.17}$$

where t_x is defined as:

$$t_x = aft\_t10\_samples - aft\_t1\_samples \tag{3.18}$$

i.e. it is the width of the pulse up to the sample where it reaches a fraction x of its area. The ratio of the width containing 25% and 75% of the area to the width containing 25% and 75% of the length can be used to determine whether most of the pulse area is concentrated at the beginning and end of the pulse, forming a shape like in figure 3.6(a).

3.4.8 Other

A final classification exists for pulses that do not satisfy any of the above categories. These pulses are often found to be baseline noise or pulses of very unusual shapes caused by PMT malfunction.

3.5 Event Classification

With pulses identified and classified, all pulses in a full waveform determine the designation of the type of event. It is important to understand the underlying causes and frequency of each type

(a) An S1 pulse. Note that more light is seen in the bottom PMTs due to reflection at the liquid-gas interface.

(b) An S2 pulse. Note the localisation in the top array, where an x shows the reconstructed location of the pulse.

(c) A single photoelectron pulse, in a single PMT on the bottom array.

(d) A single extracted electron pulse, caused by scintillation in the gas.
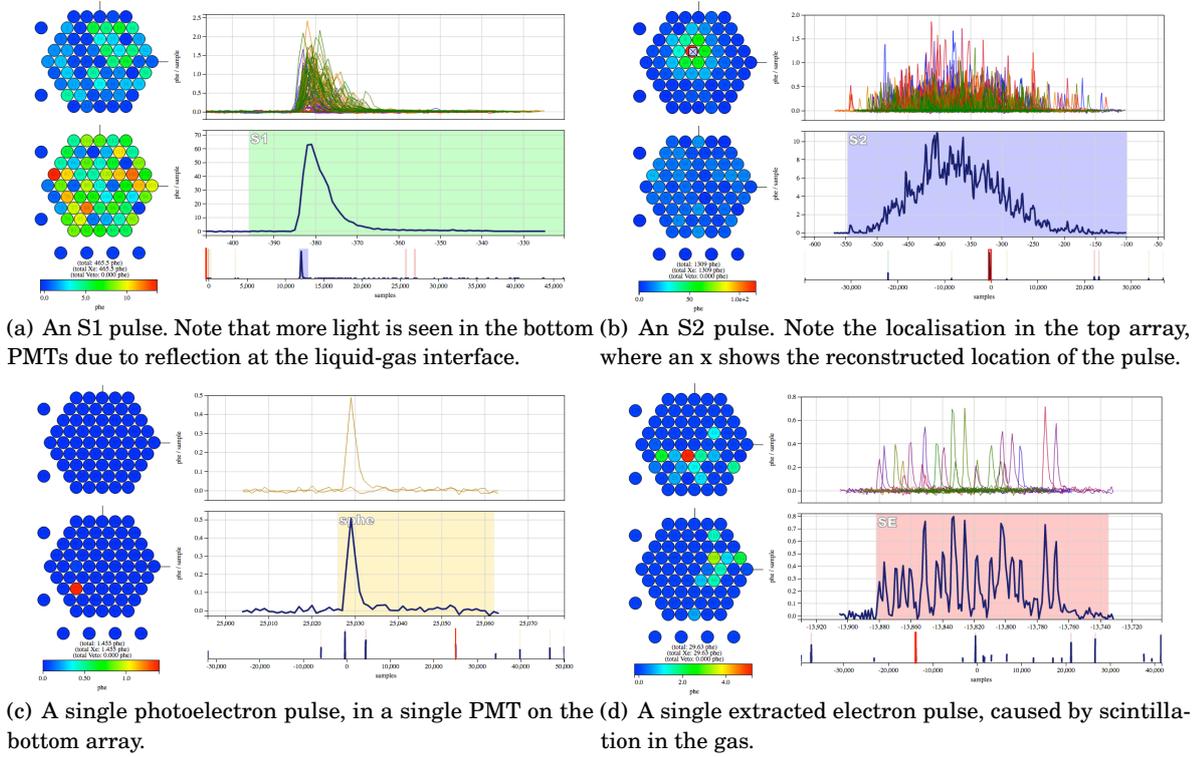
Figure 3.5: Examples of the four main pulse types in LUX data, using the VisuaLUX event viewer. The top panel of each shows individual PMTs in different colours, whilst the bottom shows the summed waveforms. Below, the full event trace is visible with a red box showing the zoomed region. The PMT arrays are shown on the left.

of event, even when they are not signal-like. In general, events in LUX can be split into several main pathologies. The requirements for these are shown in table 3.5.

Table 3.5: Table of common LUX event types and their properties.

Name	Requirements
Single Scatter	1 S1 followed by 1 S2
Multiple Scatter	1 S1 followed by multiple S2s
^{83m}Kr Event	2 S1s followed by 1 S2
Gas Event	S1 and S2-like pulse with continuous light between
Sub-Cathode	1 S1, very large, mainly in 1 PMT
E-train	Multiple SEs in exponentially decreasing frequency
E-burp	Multiple SEs emitted over a time of $\sim 100 \ \mu$s with a rise and fall.

The most important event for WIMP search is the single scatter, as this is the expected signal for a WIMP. There are caveats that allow for a single scatter to contain an S2 before the initial S1, although this is extremely rare by random coincidence, or an S1 after the final S2. These are allowed as they are coincidental to the paired S1 and S2 from a single scatter. Furthermore, if the S1 is comprised of only 2 PMTs, they must not be partner channels, as signal on one may induce signal on the other, falsely creating a coincidence. Electron recoil events (γ and β induced) will

generally create a single scatter. A multiple scatter will contain multiple S2s and is most likely caused by a neutron. Emission of S1 light from multiple vertices is usually too fast to create two separate S1s, but if the vertices are separated in z, the resulting S2s they will be distinguishable in an event due to their difference in drift time. Gas events were mentioned in section 3.4; they are caused by an interaction above the anode, and are characterized by an S1 pulse followed by continuous electroluminescence light for $< 4\ \mu s$, and an S2-like pulse at the end of the drift because of the increase in the field. Sub-cathode events are events that take place below the cathode. A single scatter here will not create an S2 because the field is reversed. Instead, because of the proximity to the PMTs, these events contain a very large S1, with most of its area within one PMT, see figure 3.6(b).

E-trains are caused by the emission of many electrons into the gas phase, usually occurring with an exponentially decreasing frequency after a very large S2. They can go on for tens of milliseconds, so can stretch over several events. They are roughly localised in the position of the original S2.
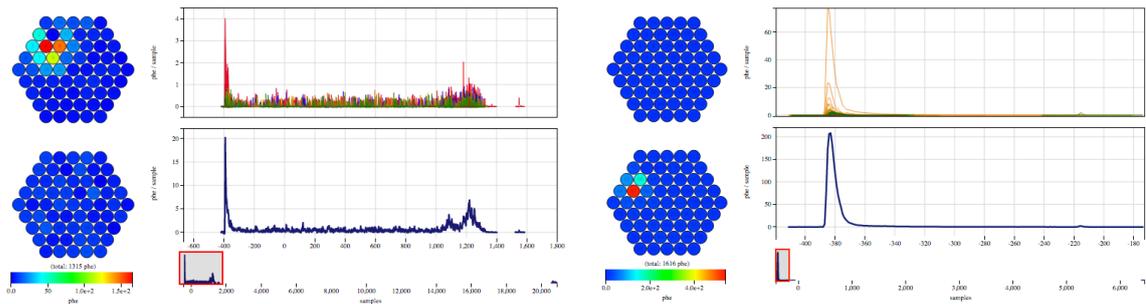
(a) A gas event. Scintillation continues between the S1 and S2-like pulses as electrons move through the gas.

(b) A sub-cathode single scatter, localised beneath PMT 67. There is no drift field so no S2 is produced.
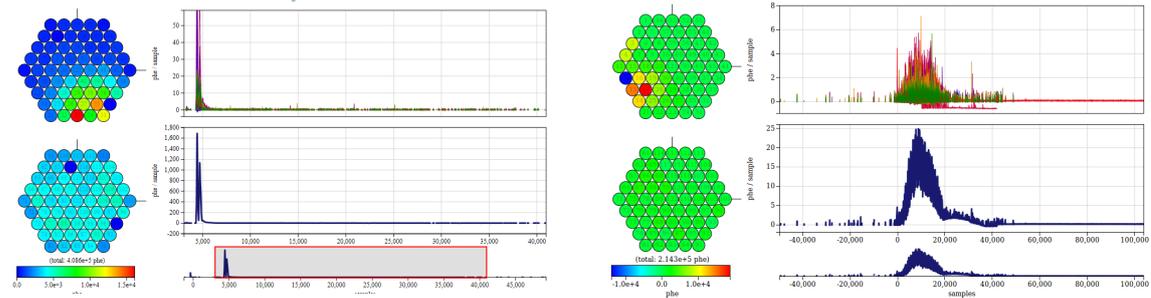
(c) An electron train event, showing the train of single electrons following a large S2.

(d) An electron burp event, where the rate of electron emission increases and then falls off again.

Figure 3.6: Non-signal event topologies visualised with VisuaLUX. Of these, only the gas event tends to be dangerous to WIMP search, as sometimes a low energy S1 and S2 are identified within the waveform.

Finally, E-burps usually also follow a preceding large S2 and are located in the same x,y position. They consist of hundreds to thousands of electrons, but rather than an exponential decay like an e-train, they rise and fall over about $100\ \mu s$. Their physical origin is still unknown.

In the DPF, a module was introduced called EventClassification that uses inputs of *pulse_classification*

Table 3.6: The input and output RQs or parameters required for the module Event_Classification.

	Type	Name	Values
Inputs	**RQs**	*pulse_classification* *pulse_area_phe* *peak_height_phe_per_sample* *skinny_peak_area_phe*	identifies S1s and S2s for checking that S2 area is greater than S1 area for partner channel check
	Parameters	*s2threshold* *echopercent* *echotime*	minimum S2 size maximum allowed echo size, see section 3.6.3 maximum allowed echo time,
Outputs		*golden* *multiple* *selected_s1_s2*	1 if golden, 0 if not 1 if 1 S1 and multiple S2s, 0 if not 1 for pulses that are either the S1 or S2(s) in a golden (multiple) event, 0s for all other pulses

RQs and produces RQs called *golden* and *multiple*. These are booleans that flag whether an event is golden or a multiple scatter. The event classification module uses an algorithm that looks for events with 1 S1 followed by 1 or more S2s. In general, a selection of events is done "by hand" in analysis as this allows more flexibility but the *golden* and *multiple* RQs are useful for quick plotting and running checks rather than an in-depth analysis. Table 3.6 lists the required input RQs and parameters and the output RQs. The input parameters are provided via the xml global settings file.

3.6 Low Energy Thresholds

3.6.1 S1

For the first LUX WIMP search analysis, a minimum pulse area of 2 phd was required for an S1. This was in line with the two-fold coincidence requirement placed on S1s (discussed in section 3.4.3), but it did remove some S1s whose area was lower either due to statistical fluctuations or position corrections. A study was conducted to analyse the impact of removing this low threshold and allowing smaller S1s. At this stage a new RQ had become available for use, calculated with the use of digital photon counting, a 2D array of the form: *spike_count*[nPMT], where generally nPMT = 122. This was a new method of pulse size estimation that counted individual PMT spikes within a pulse for each PMT. For low energy pulses, this provided a more accurate measurement of the pulse size than the usual method of integrating to find the area under the pulse. The spike count estimation becomes less accurate as pulse sizes increase, due to pile up in the PMT channels causing single, larger spikes containing more than one photon. This can be corrected for, and this was implemented within golden event selection code, with the form:

$$A_N + B_N(pulse\_area\_phe) + C_N(pulse\_area\_phe)^2 + D_N(pulse\_area\_phe)^3 \qquad (3.19)$$

(a) Effect of corrections on S1 spike counts and areas.

(b) Corrected S1 spike counts against area.
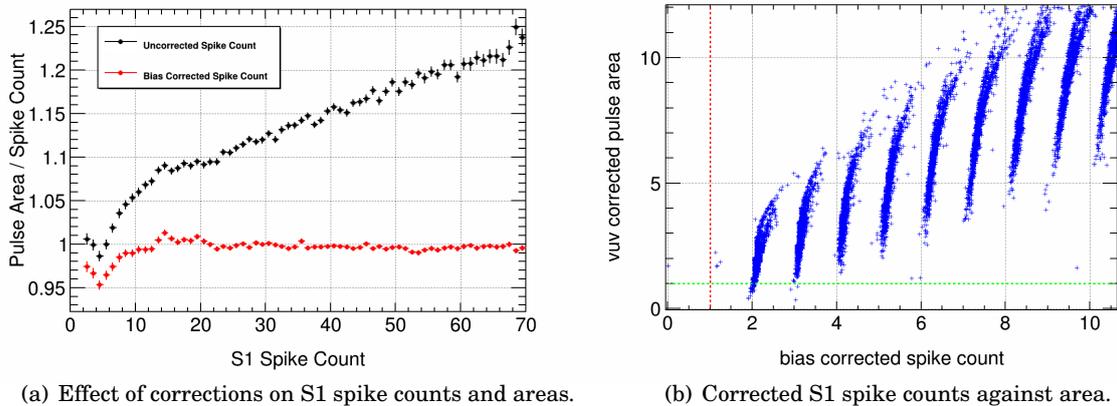
Figure 3.7: Left: Comparison of S1 spike counts with the ratio of pulse areas to spike counts. Values before (black) and after (red) application of pile-up (bias) corrections are shown. Right: distribution of corrected spike counts and areas, comparing the effects of a cut on spike count of 1 (red dashed) and a cut on pulse area of 1 (green dashed).

where A, B, C and D are the offset, linear, quadratic and cubic corrections respectively, that each depend on the number of spikes, N. The corrections return 0 if the pulse area is above 80 phd, as this is considered to be where the spike count becomes less accurate than the pulse area. These corrections was tested in tritium data, see figure 3.7(a), where it can be seen that the spike count and pulse areas deviate most below 20 phd. In this region, the spike count is more successful at representing the true size of the pulse. Applying the pile-up corrections vastly improves the agreement between spike count and pulse areas.

To test whether a cut on S1 spike count rather than area would be more appropriate, a study was performed to assess the affect of allowing S1s below 2 phd to pass selection cuts in calibration and WIMP search data. Figure 3.7(b) demonstrates that in tritium calibration data, when applying a spike count cut of 1, there is only 1 event that does not pass, as it has a spike count of 0. This was an S1 that occurred at the same time as noise in a different PMT, causing the total area of the pulse to be larger than the cut-off of 80 phd in the pile-up corrections. Applying an alternative cut of 1 phd on pulse area excluded many more S1s. Note that all these S1s had passed the 2 PMT coincidence requirement in the pulse classifier, and so there should be no reason to exclude them. From then on, all low energy S1 cuts were applied on the S1 spike count, and the threshold was lowered to 1.

Whilst S1s with corrected areas below 2 phd but passing the spike count cut were true S1s by inspection, this study did bring to light a problem involving the 10 pulse limit on the pulse finder. Despite the pulse finder having a 'safeguard' against missing an S1, the region A/B logic, it is still possible in an event with several small SEs to miss the S1. The smaller the S1, the greater the chance of this happening, so reducing the low threshold opened the pulse finder up to new inefficiencies. To solve this, a new algorithm was added to the pulse finder that prioritised S1-like pulses over SE-like or SPE-like pulses. This algorithm is described in section 3.3.5.

3.6.2 S2

For the first WIMP search results paper, see section 5.3, a low S2 threshold of 200 phd (~8 extracted electrons) was used to select data [106]. This excludes a small number of single-extracted-electron-type events (having poor event reconstruction) and wall events with poorly reconstructed positions and small S2s (due to light loss), as edge events with small S2s could be falsely reconstructed in the fiducial volume. Within the pulse classifier however, the low S2 threshold was chosen rather arbitrarily as ~4 extracted electrons, or 100 phd. After the first analysis, however, there was a call to look at lowering this threshold for several reasons. First and foremost, accepting smaller signals increases the sensitivity of LUX to low mass WIMPs, which have a softer recoil energy spectrum. Furthermore, with the introduction of the deuterium-deuterium (D-D) neutron calibration, smaller signals were required to gain the desired low energy measurements of L_y and Q_y. Finally, there was a drive to perform an S2-only analysis, which by excluding the S1 has a significant effect in increasing low WIMP mass efficiency, but again, a sensitivity to small S2s is required. Therefore, a series of studies on calibration data were performed to investigate the effects of lowering the S2 threshold within the pulse classifier below 100 phd.

As S2s and SEs are classified by all the same criteria except for area, lowering the minimum S2 size is the same as lowering the maximum SE size. With the gains and processing used for the 2013 analysis, the average SE size was measured as 25 ± 7 phd. Therefore, the boundary between SEs and S2s was experimentally lowered to 33 phd, or just above the SE mean $+1\sigma$. The effect this had on the number of single scatters and multiple scatters in a dataset was investigated, see figure 3.8.
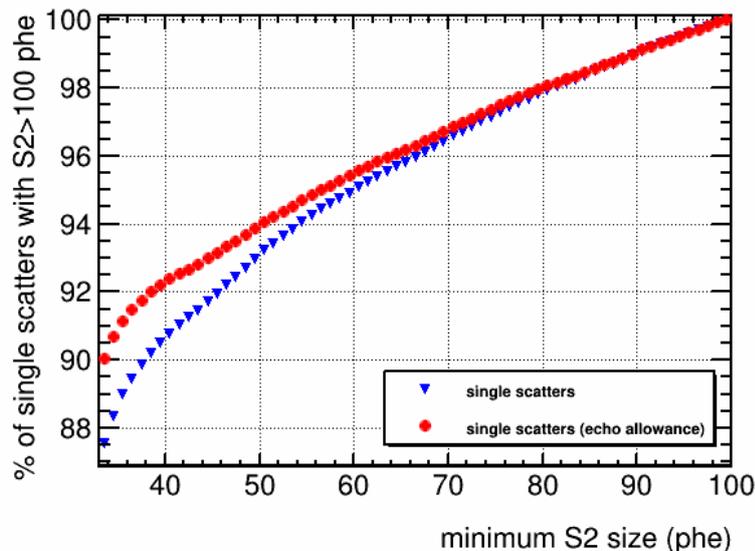
Figure 3.8: Percentage of single scatters with S2 > 100 phd still identified as a single scatter as the S2 threshold changes. The line appears fairly straight except at very low pulse areas. This appears to be where the single electron background begins to adversely effect the single scatter identification.

As there are no negatives to classifying some of the single electrons as S2s, it was decided to move the threshold to 33 phd in the classifier and then apply a separate threshold in analysis for flexibility. The *spike_count* RQ discussed above was found to also be of use here. Applying the requirement that an S2 must satisfy:

$$\sum_{i=1}^{122} spike\_count[i] > 55 \tag{3.20}$$

to be counted as a valid S2 within the golden selection removed almost all of the SE background. The pathology of pulses that are removed by this cut are single electrons with a single large PMT spike—possibly caused by afterpulsing—that raises the pulse area above 55 phd, but not the spike count, hence the success of the cut.

3.6.3 The S2 Echo

After lowering the S2 threshold, it was apparent that a large number of single scatters were being misclassified as multiple scatters, despite the fact that the data was tritium calibration data. Tritium is a β-decay and electron recoil events are almost always single scatters. The events that were 'lost' were investigated, and found to have a common pathology. This was the existence of a small S2 immediately following a larger one, in what should be a single scatter event. These pulses were called 'S2 echoes'. They tended to be localised in the same location as the original S2, see figure 3.9 for examples. These were found to be present in all types of data (nuclear recoil, WIMP search) and so are expected to be a detector related effect. A safeguard was needed to
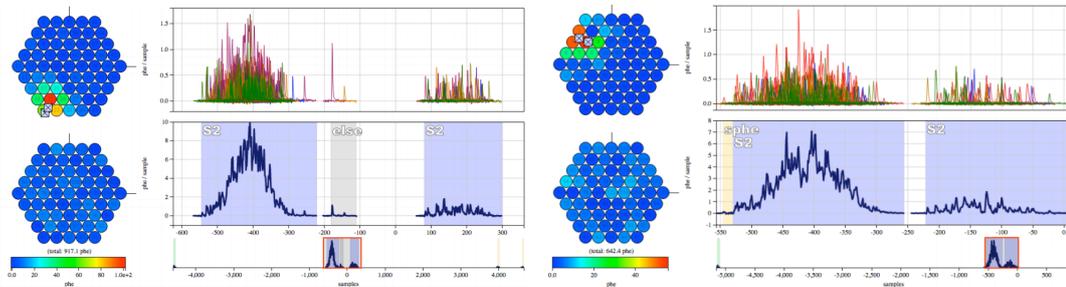
Figure 3.9: Examples of S2 echo events in tritium data, where a small localised S2 follows the actual S2. The grey crosses shown on the PMT arrays are the reconstructed position of the S2s, note how they are close together.

ensure that S2 echoes did not remove real single scatters during selection cuts on WIMP search data. Therefore, an 'echo allowance' was added to the EventClassification module. This added a caveat to the golden event definition: Any S2 following the first S2 that satisfies:

$$pulse\_area\_phe[S2_2] < 0.3 \times pulse\_area\_phe[S2_1] \quad \text{and} \tag{3.21}$$

$$aft\_t1\_samples[S2_2] - aft\_t1\_samples[S2_1] < 1000 \tag{3.22}$$

is ignored as an echo. This allowance recovered 97.5% of tritium events that were 'lost' to multiple scatters after the lowering of the S2 Threshold, suggesting as suspected that these events were not multiple scatters. The phenomenon is still not clearly understood, however, one likely explanation is that some electrons manage to pass through the anode. They then encounter a reversed field and so then drift back towards the anode, possibly emitting additional electroluminescence as they travel back down. As only some small fraction of electrons would make it through the anode, this second electroluminescence would be seen as a smaller S2 pulse shortly after the first S2. A further implication of this is that in cases where TRC is unable to split the S2 and its echo, the event will have an overestimation of S2 area. This could affect discrimination, which depends on the ratio of S2 to S1. Cuts on the S2 width and shape can be used to remove such merged S2s, and these were applied for WIMP search.

3.7 Validation and Testing of Pulse Finding and Classification

Extensive validation and testing has been done on the pulse finder and classification system (PFC). When a modification is made, it must be ensured that there are no negative effects, so several calibration datasets of different types are processed, usually tritium (single scatter ER), krypton (two S1s, 1 S2, ER) and and D-D (multiple scatter NR) datasets - see sections 5.1 and 5.2 for more information on calibration data. A technique that was used for testing of the PFC several times is handscanning. A handscan will usually involve several analysts who are accustomed to looking at LUX data looking at ~1,000s of events by eye in order to check for problems, working through an extensive checklist. Since S1, S2, and other pulses follow well predicted shapes and timing based on the underlying LXe microphysics and detector configuration, including electron chains, handscanning by experienced physicists provides a powerful technique for early validation exercises. The issues that this can identify are incorrect pulse start and end times, missed pulses, merged pulses and split pulses (pulse finder issues) and incorrect classifications (pulse classifier issues).

Several large handscans were performed throughout the run of LUX, with assistance from other data analysts. However, many more smaller ones were performed by the author to test updates and improvements to the PFC. In the following sections, some of the validation tests done for various changes made to the PFC systems are presented.

3.7.1 Multiple Scatter Identification

The efficiency of the multiple scatter identification algorithm (see section 3.3.3) was tested by checking for events that had a different number of S2s identified in data processed before and after the modification. 131 events were handscanned; the results are shown in table 3.7. For D-D calibration data, used to measure L_y and Q_y and the overall nuclear recoil response in-situ, events with 2 S2s are the most important (see section 5.2.4). Before the modification, only 17%

of events with 2 S2s identified by the PFC actually contained 2 S2s, afterwards, 60% did. The efficiency significantly increases with an energy cut, as high energy events are much more likely to be misidentified, and are also not of great importance for WIMP search and low energy calibrations. The algorithm was also tested on tritium calibration data (electron recoils) to check for adverse effects on single scatter identification. 30 events with PFC-identified multiple S2s were correctly classified; all did contain multiple S2s. In NR AmBe data, all 120 handscanned events newly classified as multiple scatters were correctly classified. A list of low energy single scatter nuclear recoils used for an absolute PFC efficiency in AmBe data was also used to check for adverse effects, and the single scatter identification efficiency was not found to decrease. A side-effect of the modification was the increase in the number of events that contained S2s not found by the pulse finder due to the 10 pulse limit. This was not previously a problem as the S2s were merged together into 1 pulse, making it much less likely that the limit was reached.

Table 3.7: Events scanned by eye to test efficiency of S2 splitting algorithm applied to D-D calibration data, where multiple scatters are expected to be common. 104 total events were handscanned. Correct means the number of S2s identified by the PFC was the same as identified by eye. Merged refers to any event where 2 or more S2s identified by eye were merged by the PFC. 10 Pulse Limit refers to events that contained less S2s as identified by the PFC as those by eye because of the 10 pulse limit of the pulse finder. M-C/e-train refers to events that did not contain the correct number of S2s because of an pulse mis-classification or an e-train, the latter often being classified as several S2s.

PFC Assignment	Handscan Result	Before	After
1 S1, 1 S2	Correct	56%	70%
	Incorrect - Merged	27%	23%
	Incorrect - 10 PL	3%	3%
	Incorrect M-C/e-train	17%	3%
1 S1, 2 S2s	Correct	17%	60%
	Incorrect - Merged	60%	5%
	Incorrect - 10 PL	0%	18%
	Incorrect - M-C/e-train	23%	17%
1 S1, 3+ S2s	Correct	9%	26%
	Incorrect - Merged	74%	18%
	Incorrect - 10 PL	0%	9%
	Incorrect - M-C/e-train	17%	47%

3.7.2 S1 Priority System, Empty Pods and the Merged Classification

Two changes to the pulse finder, the S1 priority system (section 3.3.5) and the empty pods modification (section 3.3.6) were tested alongside one change to the classifier (section 3.4.7) in order to save processing time. A benchmark tritium dataset containing about ~50,000 events was processed with the current DPF version both before and after the changes and the results compared. Table 3.8 lists the causes for any events that changed their golden classification.

A similar number of golden events were lost as were gained, but the purity of the gained events was much higher (65%) than that of the lost ones (8%), so overall the changes were positive. The merged classification occasionally caused merged S2s, e-trains or S2s and S2 echoes

Table 3.8: Changes to golden events after reprocessing with modifications to the PFC. Purity refers to how many of the events were golden by eye. Changes caused by classification are in green, the S1 priority system are in pink, and by the empty pods modification in blue. *These were previously classified as S2s.

Lost Golden Events		Gained Golden Events	
Total	541	**Total**	412
Purity	8%	**Purity**	65%
of the non-golden 92%:		*of the non-golden 35%:*	
second S1	44%	MS classified as merged	62%
MS classified as merged*	31%	S2 echo classified as merged*	23%
second S1	14%	surface	15%
S2+echo classified as merged*	11%		
of the golden 8%:		*of the golden 65%:*	
long S2 classified as merged	80%	S1 and S2 split	55%
e-train split	20%	e-train classified as merged*	25%
		S1 found	20%

to be identified as merged, leaving only one S2 in the event resulting in a golden classification. However, any events like these would fail a bad area cut so would not be an issue. All of the lost golden events were high energy and were missed due to related reasons (e-trains) and so would have not been of interest to WIMP search. On the other hand, 75% of the gained golden events were low energy and were identified due to the splitting of a close-in-time S1 and S2, or a very small S1 was identified by the S1 priority system. These are important for low energy efficiency.

A problem with the merged classification later arose after these modifications were applied to krypton data, where events contain two close-in-time S1s that are often merged - so often that the pulse classifier contains allowances for them. The first krypton calibration data processed with the updated pulse classifier had an usually low krypton event rate in analysis, due to merged S1s being classified as "merged" instead of the usual S1 classification. This was quickly remedied by ensuring that the merged classification was only applied to pulses that were over 150 phd in area and did not pass all the S1 cuts.

3.7.3 ^{83m}Kr Distribution Problem

^{83m}Kr events are characterised by having 2 S1s and 1 S2 (see section 5.1.1 for more details on krypton calibrations) and also by their uniformity within the detector. This uniformity is essential for position corrections.

An issue was noticed when run 3 ^{83m}Kr calibration data was processed with a new versions of the DPF. There appeared to be a dependence on drift time; the number of events was falling with drift time more than predicted by a COMSOL Multiphysics [108] simulation. The same dataset processed with an older DPF version matched the simulation, so this dependence clearly was caused by something in the DPF. It was necessary to determine whether this was a problem with the PFC. A problem quickly became apparent that the percentage of events being 'lost' was due to them containing multiple S1s and multiple S2s, see figure 3.10, where drift times of

events with different numbers of S1s and S2s are shown. After handscanning some events, it was determined that the S1 problem was due to events wherein a single electron existed between the S1 and S2. This SE was cut off too early by the pulse finder, leaving behind a spike that looked like an S1 to the pulse classifier. This had never been seen before during testing of the pulse finder, so was surprising, and the effect completely disappeared in an even more recent version of the DPF. After checking the data processing settings, the problematic dataset was discovered to be a test dataset with the *maximumGap* parameter reduced to only 10 samples, which meant small pulses were cut off too early. The problem with multiple S2s was caused by the lowering of the S2 threshold (see section 3.6.2), where a very conservative 33 phd had been used in the pulse classifier, and the ^{83m}Kr data analysis depended on the *golden* RQ. The failure to update the definition of the *golden* RQ was due to an assumption of lack of use; the WIMP search selection had used a separate selection algorithm for flexibility. This highlighted a need for better communication between module developers and analysts, and so the event classifier module was then modified to include the $\sum_{PMTs} spike\_count > 55$ requirement that was found to be needed for golden selection to remove single electrons, see section 3.6.2.
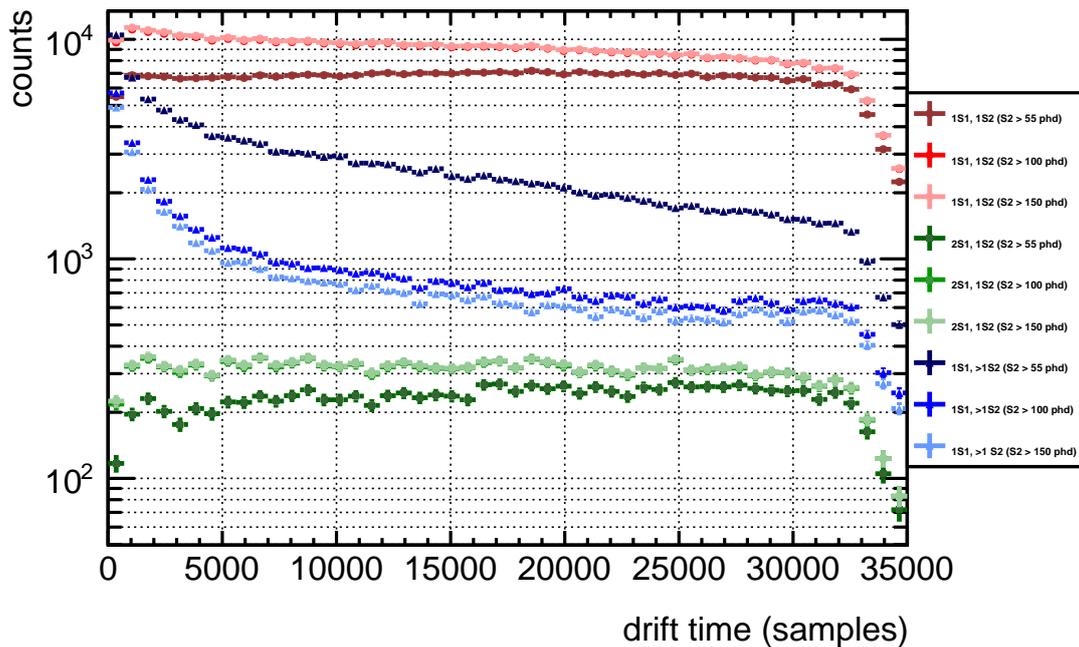
Figure 3.10: Distribution of events in a ^{83m}Kr datasets for various event pathologies. Red colours show events with 1 S1 and 1 S2 and show a gradual decrease with drift time. Blue show events with multiple S2s which are significantly more common at the top of the detector.

A similar issue later occurred with run 4 data, with an under-abundance of events at high drift time. This time, S2 thresholds were applied, and *maximumGap* was correct at 50 samples, so there was no clear reason for the problem. It was determined that again, there appeared to be

an issue with multiple S1s towards the bottom of the detector, where the percentage of events failing due to multiple S1s increased by about 0.3% over the length of the detector. Previous to the empty pods modification discussed in section 3.3.6, ^{83m}Kr S1s were predominantly merged together and the pulse classifier contained windows to allow for this and still classify them as S1s. Therefore, allowing the selection of ^{83m}Kr events with 2 S1s would solve the problem, however, it was not clear as to why the splitting should be more efficient at higher drift times, and it was assumed this must be a bug in the pulse finder. To investigate, the area of the 2 S1s (known as S1a and S1b) were plotted against each other, see figure 3.11(a). S1a should be larger than S1b as the energies of the decays responsible are 32.2 and 9.4 keV respectively. It is clear from figure 3.11(a) that there are two distinct populations; one that appears correct, with the two S1s generally increasing together and having the respective area ranges of approximately 80–250 phd and 20–90 phd. The second population, however, showed consistently small S1b with areas less than 20 phd. As the populations could be quite cleanly split with a cut on S1b of 20 phd, figure 3.11(b) was produced, which clearly demonstrates that the rate of the very small S1b pulses increases with drift time.

(a) S1a area against S1b area for events with 2 S1s.

(b) Drift time dependence of events with 2S1s.
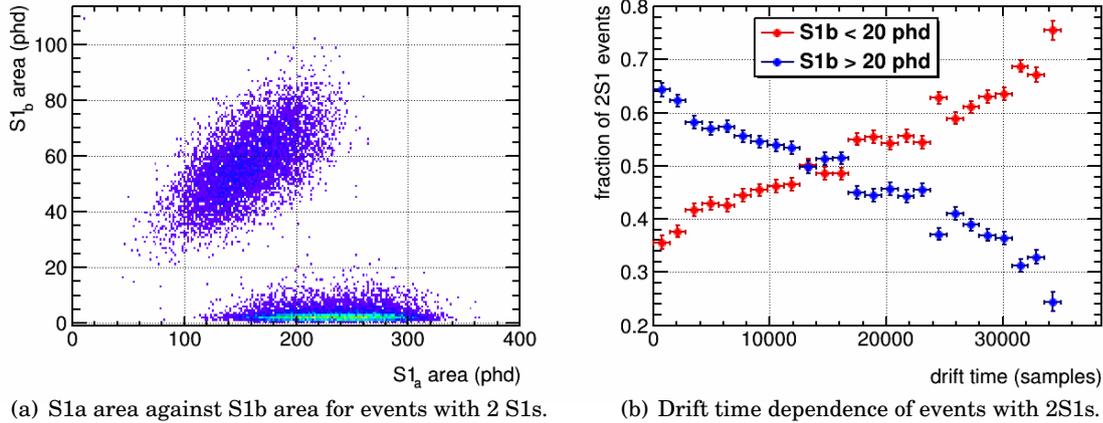
Figure 3.11: Left: S1a S1b area distributions in ^{83m}Kr data that showed an unusual dependence on drift time. Right: Fraction of events with 2 S1s for both S1b above and below 20 phd, demonstrating that the fraction of events containing small S1bs increases with depth in the detector.

Events were then handscanned. It was immediately clear that the events with the small S1b were in fact a merged S1a and S1b followed somewhere in the event by a very tiny unrelated S1, most likely caused by coincident light. This can also be seen from figure 3.11(a) where the bottom population is much larger in 'S1a'. The dependence of these small coincident S1s on drift was simply due to the fact that an increased time between the S1s and S2 leads to an increased probability that some coincident light would be detected. A cut was suggested on the area of S1a to allow these events to pass as valid ^{83m}Kr events and restore the correct distribution of drift time. In this case, there was no problem with the pulse finder, as the coincident light appeared exactly as an S1.

3.8 Impact of DPF Improvements on WIMP Search Data

3.8.1 Introduction

The LUX DPF operates by using periodic 'stable releases' of validated code. The version of the DPF used for the first results paper ([106]) was stable release 1.3 and included the original PFC system, before any of the modifications described in this work. Several large changes and improvements were made to TRC, something that was likely to have an effect on the selection of data for the run 3 reanalysis. After thorough testing, debugging and validation, a new stable release, 2.0 was issued, containing the updated pulse finder. Note that this version did not yet include the changes to the PFC described in sections 3.3.5, 3.3.6 and 3.4.7, which were instead included in stable release 2.1. In the following, these releases are referred to as DP 1.3 and DP 2.0. In order to characterise the effect of the new release, a golden selection was applied to the same WIMP search data processed with both releases, a large handscan of events was performed.

Datasets comprising approximately 10% of the total run 3 WIMP search data were used, which had already been selected for use in an analysis workshop to aid code development and new analysis. The datasets were chosen by simply taking every 10^{th} dataset from a chronological list of data. A set of RQs from both DP 1.3 and 2.0 were produced for each dataset, and only files successfully processed in both datasets were analysed. The data had a total livetime of 9 days.

3.8.2 Golden Selection

The standard golden selection was applied to both sets of data, with a bad area cut of:

$$full\_evt\_area\_phe - (pulse\_area\_phe[S1] + pulse\_area\_phe[S2]) < 100 \text{ phd} \qquad (3.23)$$

and to mimic WIMP search, an S1 energy cut and a fiducial volume cut were also applied:

$$corrected\_radius[S2] < 18 \text{ cm} \quad \text{and} \qquad (3.24)$$

$$3,800 < z\_drift\_samples[S2] < 30,500 \qquad (3.25)$$

In addition, the S2 in a golden event must satisfy the following extra requirements in DP 2.0, since the lowering of the S2 threshold to 33 phd in the pulse classifier:

1. The spike count summed across PMTs is >55
2. The S2 area is larger than the S1 area

Note that the second requirement is not strictly true for all types of event; α-particles have very quenched S2s that are smaller than their S1s. However, for low energy nuclear and electron recoils the S2 is always larger than the S1, so this helps remove any false coincidences. The results of the golden event selection for the entire 10% WS data are in table 3.9. Unique means that the event was only golden in one DP version.

Table 3.9: Breakdown of the golden event selection in both releases. Unique here means the event was only golden in one release. *These events were still golden in both versions but were above 30 phd in one.

	DP 1.3	DP 2.0
Total Golden Events	139,179	74,922
of which are unique	71,896	7,639
of which are in fiducial:	21,436	9,263
of which are unique and in fiducial:	12,184	757
Golden Events with S1 <30 phd	3,498	4,406
of which have S1s migrated below 30 phd*	0	279
of which are unique:	442	1,071
of which are in fiducial:	31	38
of which are unique and in fiducial:	5	12

The total number of golden events has almost halved from DP 1.3 to DP 2.0; however, the number of low energy events has increased. This is in line with expectations after both the introduction of a multiple scatter splitting algorithm to TRC and the lowering of the S2 threshold.

If an event was only golden in one DP version, the number of S1s and S2s in the other version was recorded. The results are shown in table 3.10. Events that had 1 S1 and 1 S2 in both but did not pass the bad area cut in one DP version were also recorded.

Table 3.10: Golden events unique to one DP version, and explanation of why they are not golden in the other.

Why not Golden	DP 1.3 All Energies	DP 1.3 S1 <30 phd	DP 2.0 All Energies	DP 2.0 S1 <30 phd
>1 S1 and >1 S2	83	0	389	0
>1 S1 and 1 S2	1,703	18	656	1
>1 S1 and 0 S2	163	91	1	0
1 S1 and >1 S2	1,876	1	31,981	47
1 S1 and 1 S2 (bad area)	978	86	28,159	335
1 S1 and 0 S2	2,524	684	191	6
0 S1 and >1 S2	3	1	2,056	6
0 S1 and 1 S2	268	174	7,191	51
0 S1 and 0 S2	41	16	1,271	2
Total:	**7,639**	**1,071**	**71,896**	**442**

This confirms that a large proportion (43%) of the golden event loss from DP 1.3 to DP 2.0 is down to the splitting of S2s in multiple scatters that were previously merged by the pulse finder. It is also seen that a large proportion (66%) of the events that were newly golden in DP 2.0 are due to an additional S2. This was found to be down to both the higher S2 threshold and pulse classifier problems in DP 1.3.

3.8.3 Handscan of Low Energy Golden Events

A sample of 393 golden events with S1 <30 phd were handscanned. These were only events that had been categorised differently in the two DP versions. The differences fell into the categories listed in table 3.11. Of these events, 287 were golden by eye.

Table 3.11: Results of the DP 1.3 vs DP 2.0 golden event handscan. Issues identified with the DP 1.3 & DP 2.0 pulse finder (PF) and pulse classifier (PC) are shown, then the overall golden efficiency is given for each. The efficiencies calculated here *Here, 104 events where the S2 was below the DP 1.3 threshold have been removed from the denominator as they are not an actual failure of the PC. The other 2 events where this is the case also had an S1 not found so were still included.

DP 1.3 Missed Events	Number of Events
PF: S1 not found	52
PF: 2 S2s merged	7
PF: Fake S1 in split pulse	3
PF: S1 split	1
PC: S2 previously below threshold	106
PC: S2 misclassified	43 (U), 2(S1), 1(SE)
Total:	215
Correctly classified (golden):	93/183*
Efficiency within this sample:	**50.8%**
DP 2.0 Missed Events	
PF: S1 and S2 merged	7
PF: S1 has fallen below threshold	2
PC: S1 misclassified	4
Total:	13
Correctly classified (golden):	276/287
Efficiency within this sample:	**96.2%**
Non-DP Specific	
Bad area either side of cut	81
Gas event	84

The pulse finder in DP 1.3 showed several issues:

- A quarter of the events missed in DP 1.3 were due to an S1 remaining unfound. An example is shown in figure 3.12(a) This was caused by the region A/B logic bug that was fixed (discussed in section 3.3.2).
- TRC was originally poor at separating close-in-time or high energy S2s. A handful of low energy double scatters were identified that were merged, see figure 3.12(b), that were now separated in DP 2.0.
- In 3 cases a pulse was split in a nonsensical way and an S1 identified from within the pulse, see figure 3.12(c). This is likely to be due to the slightly higher noise threshold used in DP 1.3.
- One case was observed of an S1 split into an SPE and an unidentified pulse, despite the two spikes being within 10 samples of each other, see figure 3.12(d).

(a) An S1 missed by the DP 1.3 pulse finder

(b) Two S2s merged by the DP 1.3 pulse finder

(c) A fake S1 identified within another pulse.

(d) An S1 falsely split into two pulses.
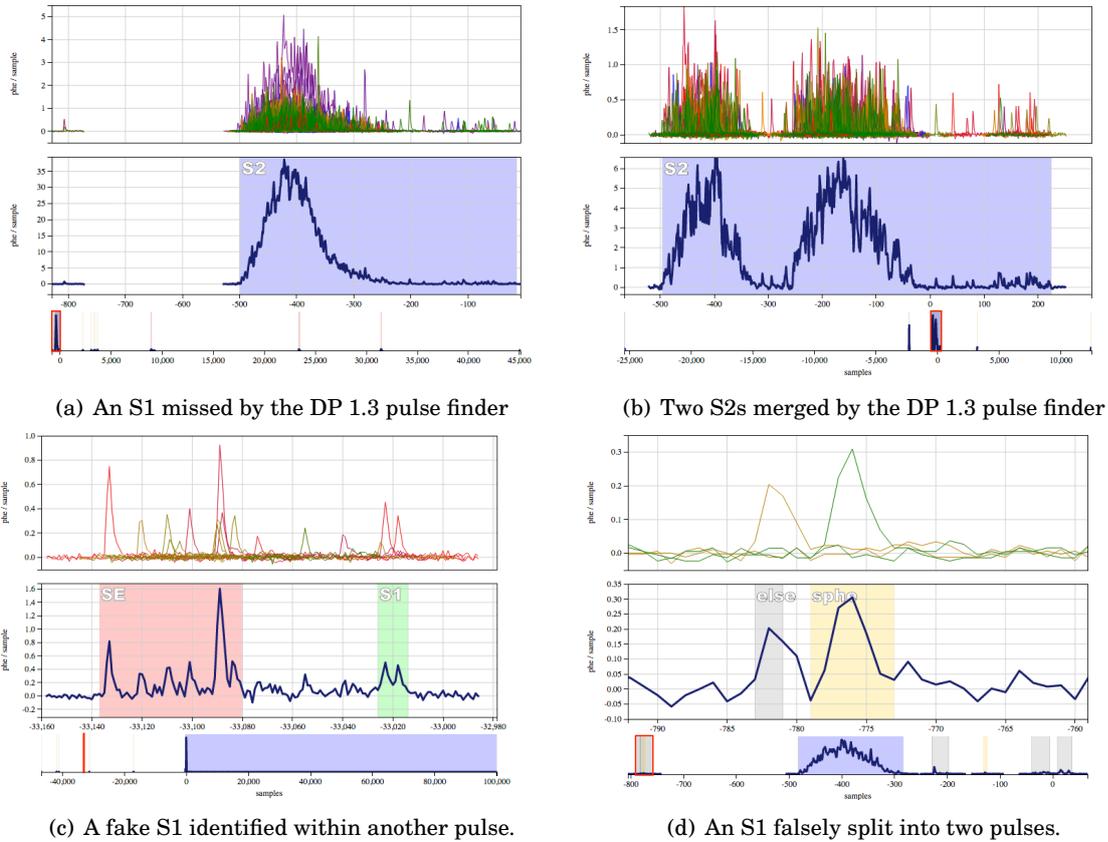
Figure 3.12: Failure modes of the DP 1.3 pulse finder.

Furthermore, the classifier had the following issues:

- Almost half of the events missed were due to the higher S2 threshold of 100 phd. This was not an actual flaw in the classifier.
- Low energy S2s were quite frequently misclassified. They were mainly unidentified, but in two cases had been given the classification of an S1. The misclassified S2s tended to be more asymmetrical, see figure 3.13(a)

For DP 2.0, the problems were as follows:

- The most frequent problem is the merging of very shallow events (with drift times under 1000 samples), see figure 3.14(a). This is a consequence of the lowering of the noise threshold within TRC - the S1 tail does not fall below this threshold for long enough before the S2 starts. These events fail the fiducial cut that defines the volume within the active LXe used for WIMP search.
- 2 small S1s were missed - 1 was not found at all, the other was found but classified as an SPE because only one of its PMTs was above the threshold. This is a consequence of the reduction in pulse areas between DP 1.3 and DP 2.0.
- 4 S1s were classified as unknown. In the case shown in figure 3.13(b) an afterpulse has

81
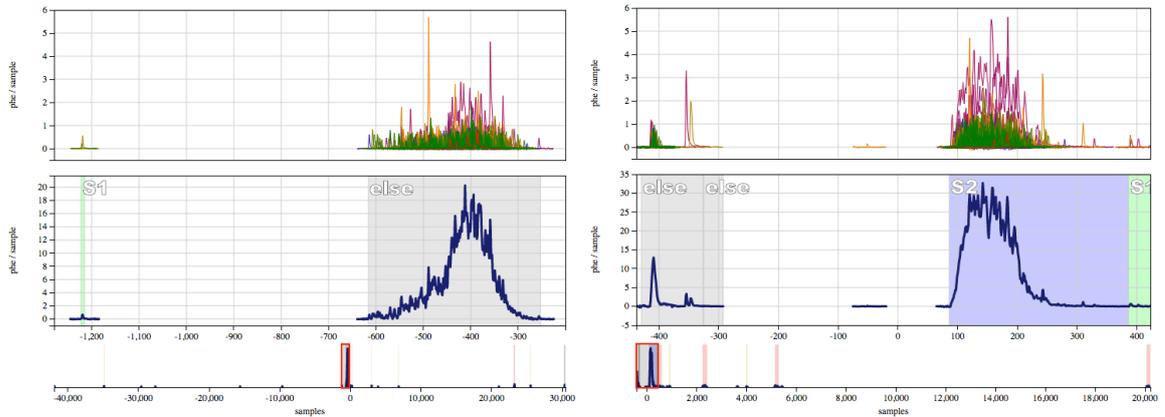
(a) An S2 misclassified as unknown in DP 1.3.   (b) A misclassified S1 in DP 2.0

Figure 3.13: Pulse classifier failure modes for DP 1.3 (left) and DP 2.0 (right)

caused the S1 classification RQs to fall outside of the normal S1 range.
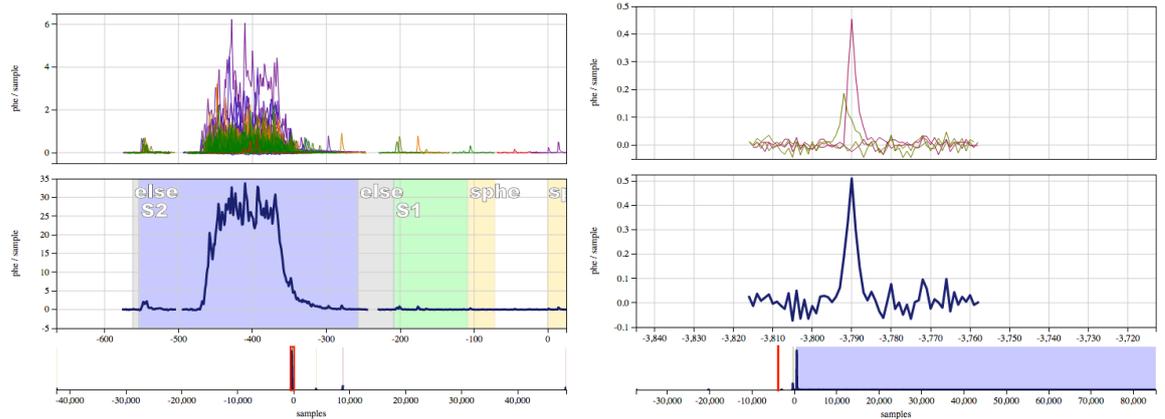
(a) A close in time S1 and S2 merged by the DP 2.0 pulse finder   (b) A small S1 missed by the pulse finder

Figure 3.14: Failure modes of the DP 2.0 pulse finder.

The large number of events passing the bad area cut in one DP version but not the other is due to two effects:

1. The reduction of pulse areas and subsequent lowering of the parameter *noiseThre* in TRC has led to small differences in the start and end points. Some large S2s that previously contained almost the whole event area are now cut off earlier.

2. The addition of the *extendPulse* parameter to TRC in order to catch the tails of SPEs has caused pulses to extend further and so contain more area, causing slightly more area to be inside the good area.

Effect 1 resulted in failure of high energy events to pass the bad area cut, whereas effect 2 acts more on low energy events, causing them to pass the bad area cut more often. At the time of

this study the bad area cut for use with DP 2.0 was still under development, but the cut of 100 phd used in this study ensured a meaningful comparison of clean events. Events moving either side of the bad area cut is not a particular concern. Gas events were often categorised differently due to the modification of the pulse classifier to identify the S1s as unknown in such events.

All low energy fiducial events were also handscanned, with the aim of identifying any non-golden events and determining why certain events were unique to a DP version.

There were two non-golden events present in DP 2.0. One was selected in both DP versions whilst the other lay outside the fiducial volume in DP 1.3. These events both contained two S1s; one below 30 phd and the second much bigger and merged with a large S2. The S2s were of an energy far beyond the region of interest for WIMP search and hence would be removed by the S2 energy cut.

Of the 5 low energy fiducial volume events that were no longer golden in DP 2.0, 4 of these were still classified as golden - 3 were the same category as the non-golden events in the DP 2.0 population: they contained 2 S1s, one of which was merged with a large S2. These events had moved over the fiducial volume boundaries and so were not selected in DP 2.0. One event was golden but no longer passed the bad area cut. The final event had no S1 and multiple S2s in DP 2.0. A handscan of the event revealed it actually contained 2 S1s, but again one was merged with an S2 in both DP versions. The S1 identified in DP 1.3 was not found in DP 2.0 due to the 10 pulse limit - the second largest pulse was at the end of the event and so the arrays filled with e-train pulses. This event would have been cut away by the S2 energy cut. Note that the merged S1 and S2 pulses would no longer be a problem since the addition of the empty pods modification, which would have separated them, leaving the event with 2 S1s and therefore non-golden.

Of the 12 low energy fiducial volume events that are unique to DP 2.0, 11 of these were due to the lowering of the S2 threshold. The remaining event was still golden in DP 1.3 but didn't pass the bad area cut. 2 of the 11 new low energy events had S2s that did not seem to be genuine, see figure 3.15. For the other 9 events, the S2s seemed to be valid by eye. One of the events are shown in figure 3.16.

3.8.4 Conclusions

A marked improvement can be seen from DP 1.3 to DP 2.0 in terms of both the purity of the golden event population and the selection efficiency. For low energy (S1 <30 phd, relevant for WIMP search), from the list of handscanned events, DP 1.3 had successfully identified 50.8% of the golden event population, whilst DP 2.0 identified 96.2% (events where the S2 was below the DP 1.3 threshold have been ignored unless they would have failed for another reason, such as a missing S1), as can be seen in table 3.11. Furthermore, After S2 energy cuts, the event population in the fiducial volume was of pure golden events - there were no finder or classifier issues affecting them. Lowering the S2 threshold has led to the addition of 11 new fiducial volume events.
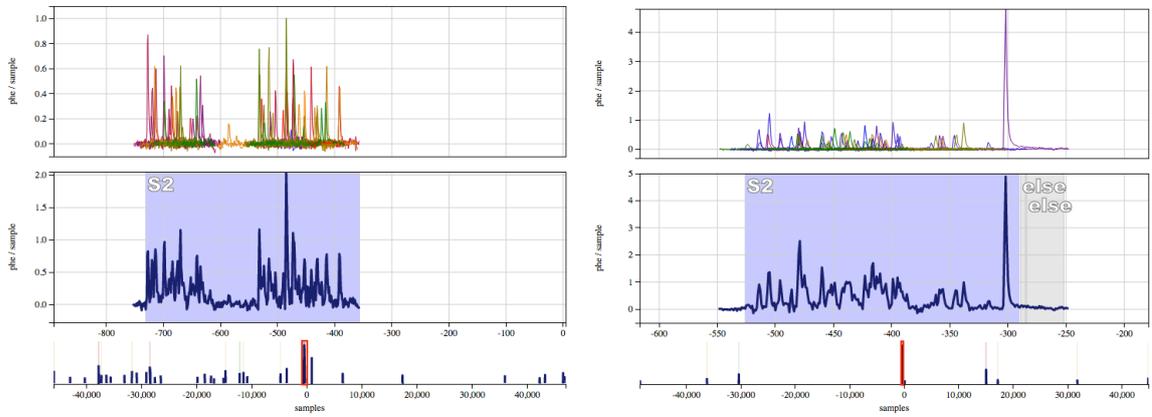
(a) A borderline S2 of 56.4 phd consisting of 2 SEs.

(b) An S2 of 76.9 phd that appears to be two SEs and a PMT afterpulsing spike.

Figure 3.15: S2s within the fiducial volume that are possibly not true S2s.
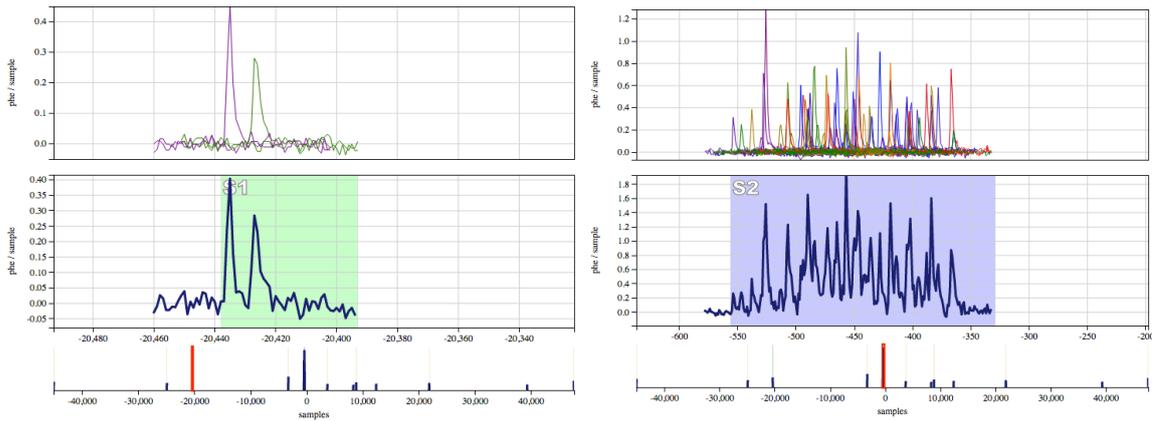
(a) S1 - 3.2 phd.

(b) S2 - 90.4 phd.
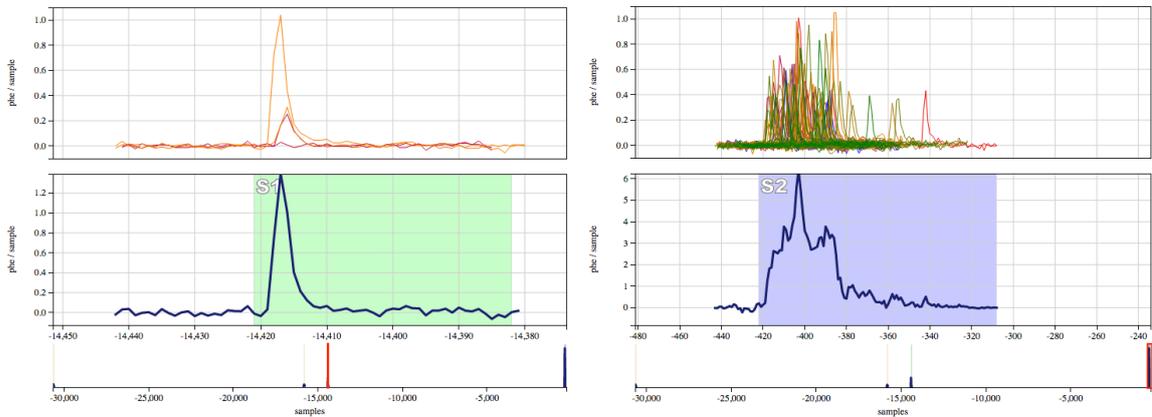
(c) S1 - 1.2 phd.

(d) S2 - 70.2 phd.

Figure 3.16: Two examples of the S1 and S2 of new low energy fiducial golden events.

3.9 Nuclear Recoil Efficiencies

The nuclear recoil efficiency determines the ratio of detected WIMP-like-nuclear recoils to those interacting in the experiment, and so is an essential quantity for WIMP search and limit setting. The efficiency depends on quantities such as the cut-off and values used for L_y and Q_y, but also relies heavily on the pulse finder and classifier's performance. The following sections present a summary of an NR efficiency study using LUXSim NR simulations. S1-only, S2-only and S1+S2 (golden efficiencies) are presented as functions of detected photoelectrons (phd) and recoil energies. The efficiencies here were relevant to the run 3 reanalysis and appeared in the reanalysis results paper [46].

3.9.1 Nuclear Recoil Simulations

Nuclear recoil simulations with a flat energy distribution from 0 - 100 keV, produced with LUXSim and NEST, are used for this study. The relevant modules are the pulse finder TC and classifier Multidimensional that have been the focus of this chapter. The versions used for this study were DPF stable release 2.0.

In order to determine the ability of the LUX detector and the LUX data processing framework to correctly identify a nuclear recoil signal, a set of nuclear recoil simulations were run using FastSim, using the Doke plot version of g_1 and g_2 (see section 5.1.4 for an explanation of this plot and the values for g_1 and g_2). These were distributed uniformly throughout the detector and the recoil energies were distributed logarithmically at set energies, with more points at low energy to clearly demonstrate the efficiency in the low energy / low mass WIMP region where the greatest improvement was seen between run 3 and the run 3 reanalysis.

In table 3.12, the variables starting *mc* are true values given in the simulation. The others are RQs calculated by the data processing framework, DP 2.0.

Table 3.12: Energy scales and the calculations/corrections applied in analysis of the NR simulations produced with LUXSim in GEANT4.

Variable	Simulation Definition
True S1 phd	$mc\_photon\_id[0]/1.284776$
True S2 phd	$mc\_photon\_id[1]/\left(1.284776 \times e^{\frac{-1.0}{100}(54.6416 - mc\_z\_cm)}\right)$
S1 spike count	$\left(\sum_{bot}(spike\_count)_{bc} + \sum_{top}(spike\_count)_{bc}\right) \times correction\_s1\_xyz\_dependence$
S1 thresholds	$1 \geq$ S1 spike count ≥ 50
S2 area	$pulse\_area\_phe$
S2 threshold	S2 area > 150 phd$_{data}/0.915 = 163.9$ phd$_{sim}$
Energy (keV$_{nr}$)	mc_E_keV

Note that in table 3.12, the subscript *bc* refers to the bias correction, not bottom-corrected. A pile up correction dependent on the bottom/top array area and number of spikes is applied separately to the bottom and top spike counts. This correction accounts for the probability of single photons overlapping in the same channel - thus only giving an increase of 1 in the spike

count - which increases as the S1 increases in size. Thresholds are applied to the S1 spike count after pile-up and position corrections are applied, but for S2 the low threshold is only applied to the raw pulse area. This is done in real data to avoid including small S2s from the edges of the detector that have been poorly reconstructed in the fiducial volume, making their area position correction unreliable.

The simulation did not account for double photoelectron emission, so the true areas must be corrected afterwards. In LUXSim, the variables mc_photon_id represent the number of photoelectrons produced for each pulse. The probability of 1 photon producing two photoelectrons is averaged over all PMTs (after removing bad channels that are no longer used) to give 0.284776. Therefore, we divide the number of photoelectrons (mc_photon_id) by 1 + this factor in an attempt to obtain the true number of photons hitting the PMT, and so a true pulse area. This is of course an approximation, but since LUXSim does not produce the number of detected photons it was the simplest route and was considered accurate enough for this purpose.

The subscript $data$ on the raw S2 area of 150 phd used for the S2 threshold is to represent that this threshold was determined using raw area in data that was processed with incorrect PMT gains. In the simulation, the same gains were assumed both within the simulation and the processing, so no corrections are needed. This means the threshold must be corrected to enable use within the simulation. The average gain correction of 0.915 was used to determine the S2 threshold of 163.9 phd.

3.9.2 Selection Criteria

The definition of a golden event is as follows:
- 1 valid S1 preceding 1 valid S2
- A valid S1 requires >2 PMT channels, *or* 2 PMTs that are not partner channels
- A valid S2 has *pulse_classification==2*, a total spike count greater than 55 and area larger than the S1
- S1s after the S2 are allowed, as are S2s before the S1

The fiducial volume, defining a mass of 152 kg, is defined using Monte Carlo truth values for R and Z, as events that do not have an S1 or S2 found would not be reconstructed within the DPF:
- 6.5 cm $<mc\_z\_cm$ <48.5 cm
- R <20 cm ($R = \sqrt{mc\_x\_cm^2 + mc\_y\_cm^2}$)

3.9.3 Efficiencies

Each efficiency curve was calculated by the division of a numerator histogram by a denominator histogram. The requirements to fill these histograms can be seen in table 3.13. Initially, a decrease in S1 efficiency at higher energies (>25 keV$_{nr}$ / 30 phd) was seen in the S1 efficiency. This is caused by larger S1s becoming more symmetrical and failing the pulse classifier cut on the AFT

Table 3.13: Definition of efficiency curve denominators and numerators.

Efficiency	Denominator	Numerator
S1	All events in fiducial passing S1 AFT cut	S1 found: *pulse_classification==1*
S2	All events in fiducial	S2 found: *pulse_classification==2*
Golden - raw	All events in fiducial passing S1 AFT cut	1 S1 & 1 S2 (with caveats, see definition)
Golden - thresholds	All events in fiducial passing S1 AFT cut	as for golden—raw but with 1≥S1 spike count ≤50 and S2$_{raw}$ ≥163.9

ratio in equation 3.6. These symmetrical S1s are overestimated within the simulation compared to data and so any event containing such an S1 was excluded from the efficiency study. In data, the number of S1s failing this classifier cut is negligible. Events containing these symmetrical S1s were still used in the S2-only efficiencies as it is independent of the S1 efficiency.
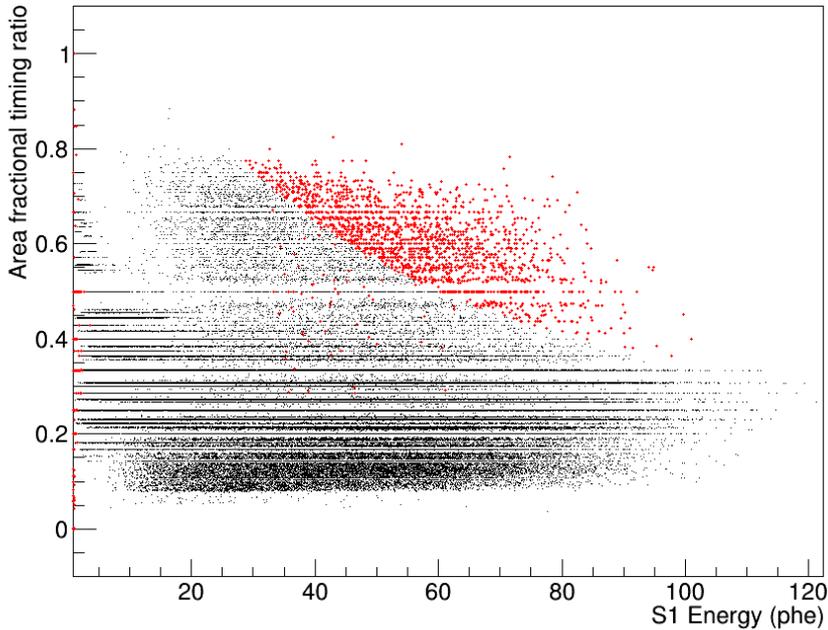
Figure 3.17: Distribution of the AFT ratio in S1s within the NR simulations. Red points shown all S1s not correctly identified by the PFC. The majority lie above the classification cut line, see figure 3.4(g). As these are not present in data, they were excluded from the efficiency denominators for S1 and golden.

The S1-only efficiency can be seen in figure 3.18. The efficiency reaches 100% at about 10 keV$_{nr}$. The two-fold coincidence requirement for S1s affects the low energy S1 efficiency significantly, as the lower the recoil energy, the more likely it becomes that only one photon will be detected. The S2-only efficiency is shown in figure 3.19. The efficiency reaches 100% much earlier than for S1-only, at about 2 keV$_{nr}$, due to the larger size of S2s. This is what allows S2-only analyses to be sensitive to lower WIMP masses. Note that the efficiencies using the Bezrukov model are higher at low energies.
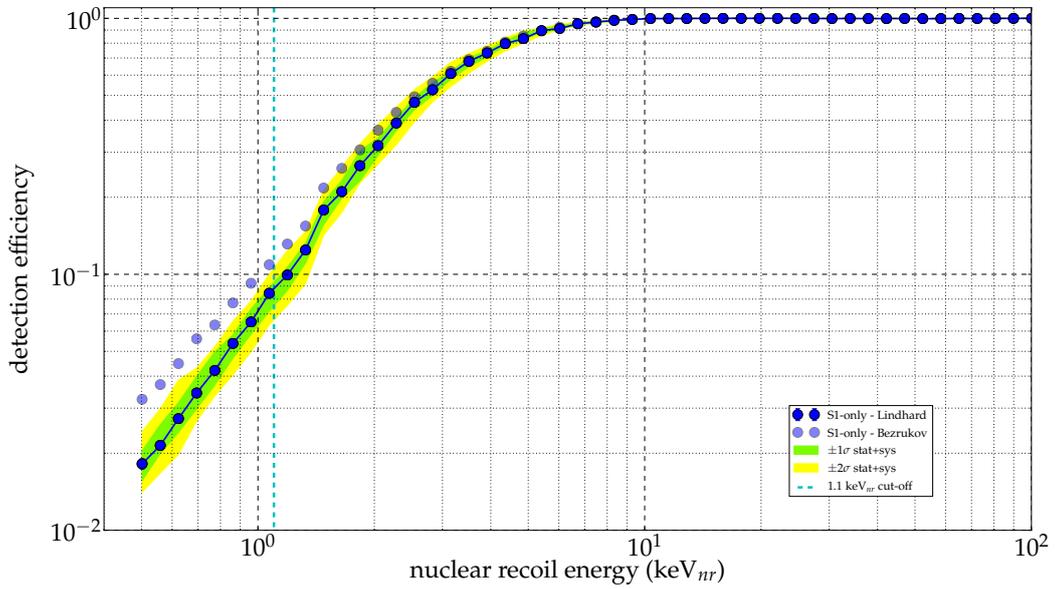
Figure 3.18: S1-only efficiency as a function of true nuclear recoil energy in keV_{nr} from flat NR simulations. 1 and 2-σ bands are shown in green and yellow respectively. Results from simulations run using the Bezrukov model are shown as transparent data points. The cyan line shows the 1.1 keV_{nr} cut-off used for the LUX run 3 reanalysis and run 4, below which no signal yield is assumed.
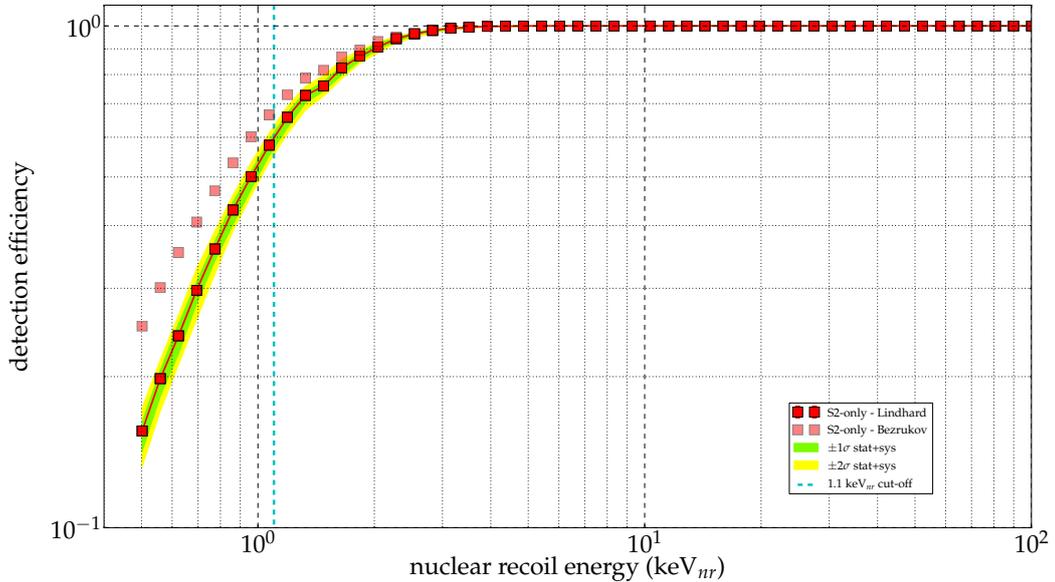
Figure 3.19: S2-only efficiency as a function of nuclear recoil energy in keV_{nr} from flat NR simulations. 1 and 2-σ bands are shown in green and yellow respectively. Results from simulations run using the Bezrukov model are shown as transparent data points. The cyan line shows the 1.1 keV_{nr} cut-off used for the LUX run 3 reanalysis and run 4, below which no signal yield is assumed.

The golden efficiency folds the S1 and S2 efficiencies together and can be seen in figure 3.20(a). Figure 3.20(b) shows the curve including thresholds. The fall-off at higher energies is due to the upper limit on S1 size of 50 phd.
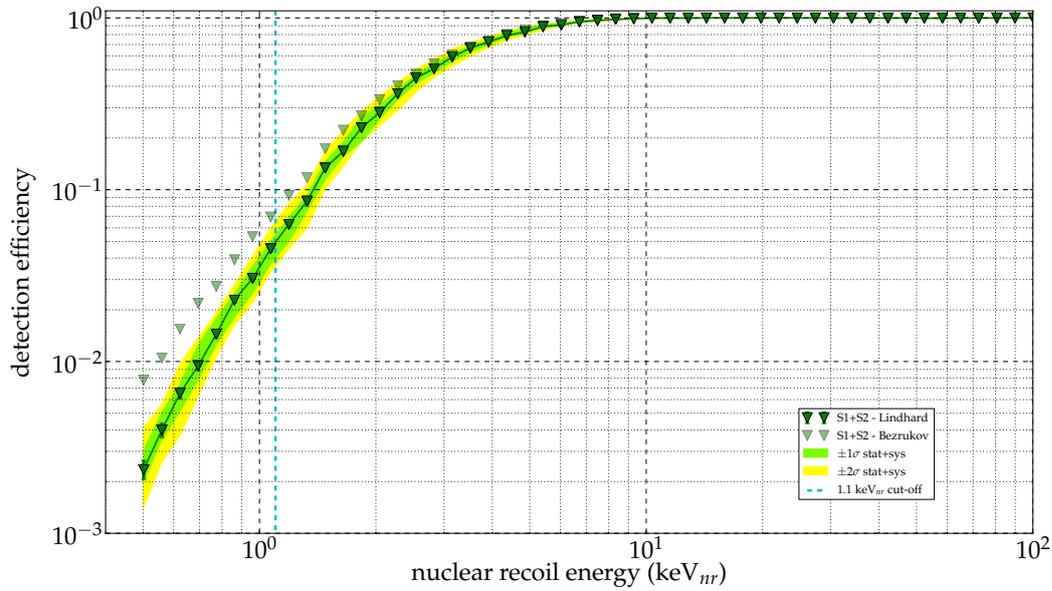
(a) Golden efficiency as a function of recoil energy from flat NR simulations.
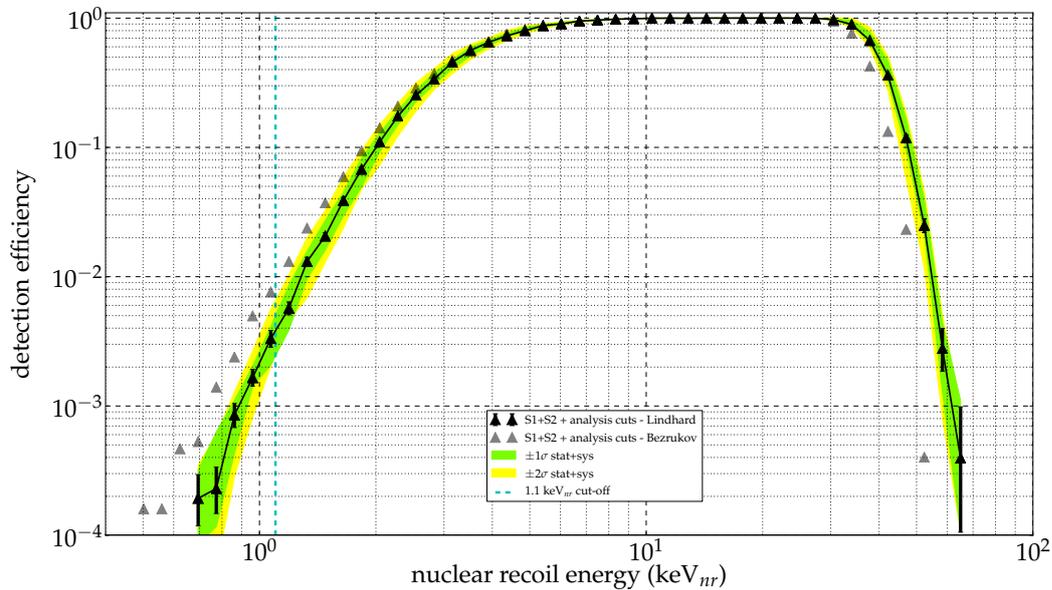
(b) Golden efficiency as a function of recoil energy from flat NR simulations after the application of analysis cuts.

Figure 3.20: Golden efficiencies before (top) and after (bottom) application of analysis thresholds. Results from simulations run using the Bezrukov model are shown as transparent data points. The cyan line shows the 1.1 keV$_{nr}$ cut-off used for the LUX run 3 reanalysis and run 4, below which no signal yield is assumed.

Figure 3.21 shows S1-only, S2-only, golden and the golden with threshold efficiency curves as functions of nuclear recoil energy on the same axes for comparison. The S1-only efficiency dominates the golden efficiency, but adding analysis thresholds has a large effect, mainly because of the S2 threshold of 150 phd.
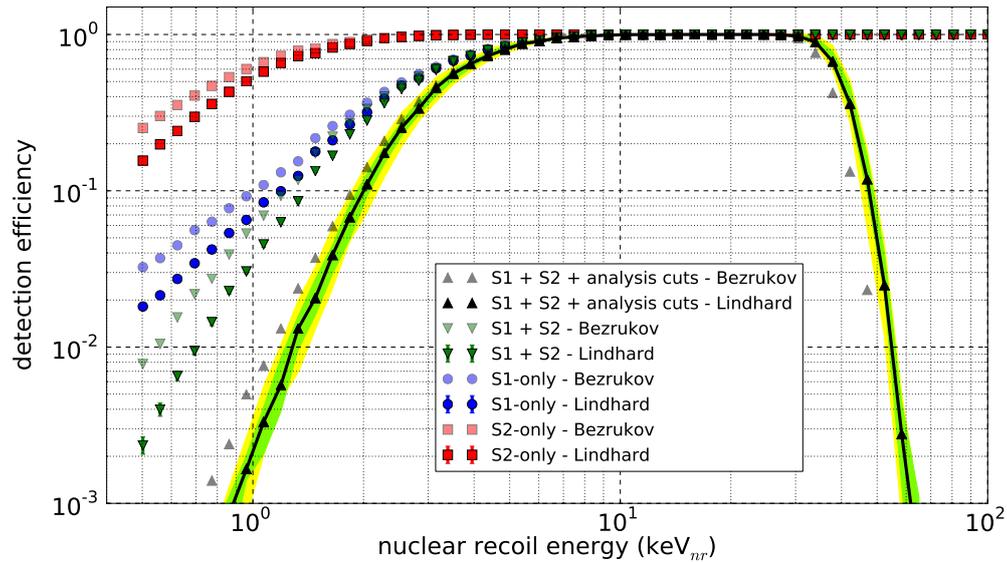
Figure 3.21: Detection efficiencies as a function of nuclear recoil energy. Blue circles: S1-only, red squares: S2-only, green triangles: golden, black triangles: golden with S1 spike count between 1 and 50, and S2$_{raw}$ > 163.9 phd.

Table 3.14 gives the recorded Lindhard efficiency at the three given energies. These were chosen because in the first analysis of the LUX run 3 data a hard cut-off in the light yield was assumed at 3 keV$_{nr}$, using the Lindhard model. It is clear that the analysis cuts have a large impact on the golden efficiency, especially at 1 keV$_{nr}$. However, the S2 threshold is needed to reduce single electron background and poorly reconstructed S2s in real data. These efficiencies were each obtained from a file containing only nuclear recoils of the stated energy and so are not from integrated bins.

Table 3.14: Efficiencies for 1, 2 and 3 keV$_{nr}$ using the Lindhard model.

Energy	S1-only	S2-only	Golden	Golden + cuts
1 keV$_{nr}$	9.27%	63.2%	6.27%	0.26%
2 keV$_{nr}$	36.3%	94.0%	33.5%	12.9%
3 keV$_{nr}$	60.0%	99.2%	59.0%	44.8%

The efficiencies at the lowest point obtained for total light and charge yield using D-D data are shown in table 3.15, for the both Lindhard and Bezrukov models (as described in section 2.3.4). Note that whilst data is indeed available for the yields at 1.1 keV, the models are used instead within NEST to generate light and charge all the way up to 100 keV$_{nr}$ for these simulations.

90

Whilst both models were found to fit the data well, the Lindhard model was chosen for final analysis as it is the more conservative model at low energies.

Table 3.15: Efficiencies for 1.1keV_{nr}. Errors are statistical only.

Model	S1-only	S2-only	Golden	Golden + cuts
Lindhard	$(8.8 \pm 0.2)\%$	$(59.8 \pm 0.4)\%$	$(5.0 \pm 0.2)\%$	$(0.39 \pm 0.05)\%$
Bezrukov	$(11.5 \pm 0.2)\%$	$(68.1 \pm 0.2)\%$	$(7.5 \pm 0.1)\%$	$(0.90 \pm 0.05)\%$

3.9.4 Comparison to Run 3

Efficiencies obtained from simulation data processed with DP 1.3 for the LUX first results paper [106] are plotted here alongside the new efficiencies. The efficiences have been cut-off at the point of their respective low energy thresholds. As the upper S1 threshold applied has been increased from 30 phd to 50 phd, there is an obvious difference in the position of the fall off. The impact of this improvement at low energies contributes in combination with lower L_y and Q_y to improve WIMP search sensitivity for low mass WIMPs. The improvement is clear in figure 5.23, where the mass range 3.3–5.2 GeV became detectable for the first time in LXe.
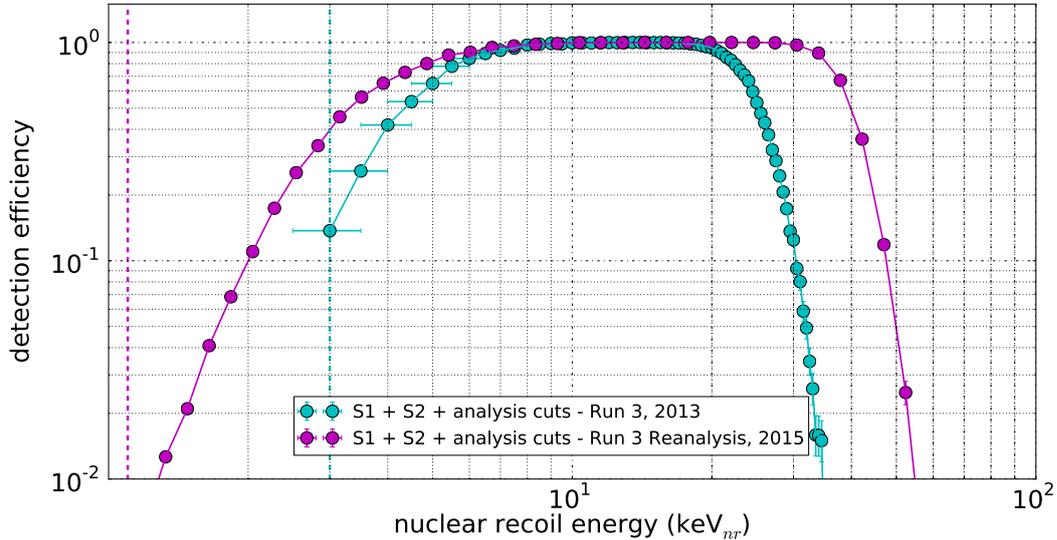
Figure 3.22: Comparison of DP 1.3 (cyan) with DP 2.0 (magenta) golden efficiencies as a function of nuclear recoil energy, after the application of S1 and S2 thresholds

Figure 3.22 demonstrates the significant efficiency improvement seen at low energies.

3.9.5 Absolute PFC Efficiencies

The above nuclear recoil study was done on a simulated flat nuclear recoil spectrum. It is essential to ensure that the PFC does not have failure modes that may not be identifiable with simulation and instead require pathologies from real data.

Therefore, a handscan of an AmBe calibration dataset was performed in order to determine an absolute PFC efficiency. All single scatter events from within a preselected set of events were identified, and the PFC's ability to correctly categorise those events calculated. The events were handscanned as evt files only, so that there was no bias from seeing the overlaid DPF decisions on pulse parameters. A loose set of cuts were applied to select candidate NR events, mainly only on event level RQs to try to avoid any bias. This was done as the full dataset would be hundreds of thousands of events and far too many to look at by eye. The preselection cuts were as follows:

- Low energy cut: $55 < full\_evt\_area\_phe < 5000$

- Average signal per sample:

$$\frac{full\_evt\_area\_phe}{n\_samples\_in\_evt} > \begin{cases} -0.11(A_{max}/full\_evt\_area\_phe) + 0.145, & full\_evt\_area\_phe < 250 \\ 1.4(-0.11(A_{max}/full\_evt\_area\_phe) + 0.145), & full\_evt\_area\_phe < 400 \end{cases} \tag{3.26}$$

The majority of e-train event made up just of SEs and SPEs will be below this threshold.

- Sub-cathode cut:

$$top\_bottom\_asymmetry < \begin{cases} = -0.8, & 100 < pulse\_area\_phe < 1000 \\ < -0.6 & pulse\_area\_phe > 1000 \end{cases} \tag{3.27}$$

This aims to remove events where the vertex was below the cathode so events contain a large S1 in bottom array and no S2. This depends on pulse-level quantities but should not introduce bias as it is unlikely to remove low energy golden events.

A few approximations had to be made in order to calculate the absolute efficiency; areas and spike counts had to be calculated from evt files (for cases where the PFC had not identified the pulse), some flat factors were applied to correct gains for the effect of double photoelectron emission, raw areas were used for thresholds (for S2 and bad area cuts, not for S1 as this may introduce a systematic error on the threshold of ~5%) and finally the S1 partner channel cut was not applied as it was shown in previous studies of the golden event definition that it has very minimal impact.

The bad area cut applied to events was:

$$A_{bad} < \begin{cases} 80 & \text{if } A_{good} < 630 \\ 80 + 0.095(A_{good} - 630) & \text{otherwise} \end{cases} \tag{3.28}$$

where $A_{good} = pulse\_area\_phe[S1] + pulse\_area\_phe[S2]$ and $A_{bad} = full\_evt\_area\_phe - A_{good}$. Around 2,000 events were selected for handscanning. The results are shown in table 3.16.

Table 3.16 shows the results of the handscan and the failure modes of the DPF. After application of WIMP search cuts, the only failure mode remaining was a missed S1. This S1 was misclassified as an SE after being merged with a closely following SPE. This was a problem that was fixed by the 'empty pods' modification mentioned in section 3.3.6.

Table 3.16: Absolute efficiency of the LUX DP 2.0 PFC using handscanned AmBe data. The total column shows the number of events identified by eye, 1S1 + 1S2 is a successful golden event identification and the other four categories are failure modes.

	Total	1S1 + 1S2	>1 S2	No S1	No S2	>1 S1
All	216	189 (87.5%)	6 (2.8%)	11 (5.1%)	7 (3.2%)	3 (1.4%)
Bad area < 80 phd	178	161 (90.4%)	1 (0.6%)	9 (5.1%)	5 (2.8%)	2 (1.1%)
Drift > 33 us (5 cm)	138	130 (94.2%)	1 (0.7%)	7 (5.1%)	0 (0.0%)	0 (0.0%)
S2 < 3000 phd	137	129 (94.2%)	1 (0.7%)	7 (5.1%)	0 (0.0%)	0 (0.0%)
S1 < 30 phd	91	86 (94.5%)	0 (0.0%)	5 (5.5%)	0 (0.0%)	0 (0.0%)
S2 > 200 phd	71	70 (98.8%)	0 (0.0%)	1 (1.4%)	0 (0.0%)	0 (0.0%)
S1 > 1 phd	71	70 (**98.8%**)	0 (0.0%)	1 (1.4%)	0 (0.0%)	0 (0.0%)
S1 > 2 phd	68	67 (98.5%)	0 (0.0%)	1 (1.5%)	0 (0.0%)	0 (0.0%)
S1 > 3 phd	62	61 (98.4%)	0 (0.0%)	1 (1.6%)	0 (0.0%)	0 (0.0%)

Table 3.17 takes the events identified as golden by the DPF and shows the results of a handscan to identify the purity of the events. The only failure mode was multiple S2s which were not split by the pulse finder. The multiple scattering splitting algorithm is less efficient for low energy S2s as the pulse is less smooth, and of the total sample 6 (3.1%) of the golden events were actually close-in-time multiple scatters. For the WIMP search cuts used for the run 3 reanalysis, the absolute efficiency of nuclear recoil identification in an AmBe dataset was 98.8%, and the purity of events selected by the DPF was 98.6%. Both cases had just one incorrect event.

Table 3.17: Purity of golden events selected by DP 2.0. The only failure mode was the merging of multiple S2s.

	Total	1S1 + 1S2	>1 S2
All	195	189 (96.9%)	6 (3.1%)
Bad area < 80 phd	166	162 (97.6%)	4 (2.4%)
Drift > 33 us (5 cm)	134	131 (97.8%)	3 (2.2%)
S2 < 3000 phd	133	130 (97.7%)	3 (2.3%)
S1 < 30 phd	89	88 (98.9%)	1 (1.1%)
S2 > 200 phd	73	72 (98.6%)	1 (1.4%)
S1 > 1 phd	73	72 (98.6%)	1 (1.4%)
S1 > 2 phd	70	69 (98.6%)	1 (1.4%)
S1 > 3 phd	64	63 (98.4%)	1 (1.6%)

The equivalent absolute efficiency for an AmBe handscan done previously for DP 1.3 was 94.0%, with 5 missed events; 1 missed due to a falsely identified S2 and 4 missed due to no S1. The purity was 94.9%, with 4 non-golden events falsely selected due to the merging of S2s. Therefore, my modifications to the pulse finder resulted in a ~4% increase in both efficiency and purity, which is modest but important when attempting to identify a rare dark matter signal.

$$4$$

LUX BACKGROUNDS

For any rare event search, it is essential to fully characterise all sources of background and to understand their appearance in the detector. If the background rate is too high, it becomes impossible to convincingly find a dark matter signal. Furthermore, if the background is low but poorly quantified, statistical significance can not be readily ascribed to any excess that is observed over expectation. This chapter will discuss background sources in LUX, how they are modelled and measured, and my work on a search for ^{210}Bi α-decays.

4.1 Discrimination, the DRU, and Analysis Differences

As previously discussed in section 2.1, the difference in the division in energy between S1 and S2 for electron recoils and nuclear recoils is large enough to allow >99.5% discrimination. However, this $\sim 0.5\%$ leakage is dangerous for dark matter searches. Even more dangerous is backgrounds from neutrons, which can induce nuclear recoils like a WIMP scatter. Neutrons may be distinguished in that they will often scatter more than once within the xenon volume, creating multiple S2s, but some may not and would be indistinguishable from a WIMP. Therefore, it is important to have an accurate and well-validated background model for any dark matter search susceptible to these sorts of backgrounds.

The discrimination space used by LUX is S1 vs log(S2/S1). As will be seen in the next chapter, ER and NR events form two distinct bands when these variables are plotted. The discrimination efficiency is based on the number of events that leak from the ER band into the NR band.

An important quantity for both this section and the later chapter on LZ backgrounds is the DRU; this is a canonical unit in the dark matter community, standing for Differential Rate Unit. 1 DRU is defined as 1 event/kg/day/keV, where the relevant mass is the detector target. DRU can also be differentiated into DRU_{ee} for electron recoils using keV_{ee} and DRU_{nr} for nuclear recoils

Table 4.1: Important quantities and their difference between LUX WIMP search analyses. *Problems with the electric field in run 4 meant a smaller and time-varying fiducial volume had to be used. The radial cut was always 3 cm from the measured position of the wall at the time of data taking, see section 5.3.5 for details.

Quantity	Run 3	Run 3 Reanalysis	Run 4
Livetime (days)	85.3	95.0	332.0
S1 range (phd)	2–30	1–50	1–50
L_y cut-off (keV$_{nr}$)	3	1.1	1.1
Drift cut (μs)	38–305	38–305	40–300
Radial cut (cm)	18	20	~21*
Fiducial mass (kg)	118.3 ± 6.5	145.4 ± 1.3	98.4–105.4*

using keV$_{nr}$ (see section 2.3.3 for a definition of these energy scales). In general a dark matter signal is expected to give a rate on the order of mDRU$_{nr}$ or nDRU$_{nr}$.

Additionally, there were several important differences between analysis for the run 3, the run 3 reanalysis and the run 4 WIMP search results. These are detailed in table 4.1, and will be useful in this and the following chapter.

4.2 Internal Backgrounds

4.2.1 Cosmogenic Xenon Activation

Activation of xenon by cosmic rays produces radioactive isotopes, decays of which will produce a background. The LUX xenon was stored at 200 m above sea level for 8 months at Case Western Reserve University, then was shipped to SURF in batches. Half was stored at the surface laboratory at an altitude of 1.6 km altitude for a further month, the other half for only 7 days. To model activation, sea level results from ACTIVIA were used and corrected with an altitude scaling factor of 3.4.

All the predicted isotopes were visible and measurable in early LUX data, see figure 4.1. The decays are as follows:

$$^{127}_{54}\text{Xe} + e^- \rightarrow \,^{127}_{53}\text{I} \qquad\qquad (t_{1/2} = 36.35 \text{ days}) \qquad (4.1)$$

$$^{129m}_{54}\text{Xe} \rightarrow \,^{129}_{54}\text{Xe} \qquad\qquad (t_{1/2} = 8.88 \text{ days}) \qquad (4.2)$$

$$^{131m}_{54}\text{Xe} \rightarrow \,^{131}_{54}\text{Xe} \qquad\qquad (t_{1/2} = 11.93 \text{ days}) \qquad (4.3)$$

$$^{133}_{54}\text{Xe} \rightarrow \,^{131}_{55}\text{Cs} + e^- \qquad\qquad (t_{1/2} = 5.25 \text{ days}) \qquad (4.4)$$
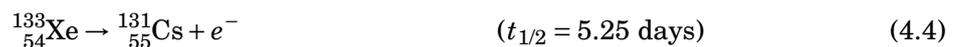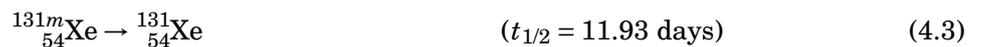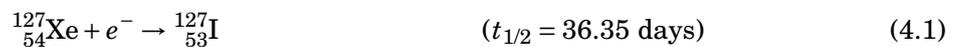
The EC and excited state decays involve internal transition and the emission of γ-rays, thus producing an ER background.

Once the xenon is underground, the activation from the reduced cosmic ray flux becomes negligible. The activated xenon began to decay away as soon as the xenon was brought underground and so only presented as a background for early run 3 data. ^{127}Xe EC transitions are of energies

Table 4.2: Table of Xe radioisotopes comparing predicted decay rate with those observed after 90 days underground, when activation has stopped. Predictions were made assuming appropriate exposure at sea level, then differing mass exposures at 1.6 km. An overall $\times 8$ factor was applied to match measurements, which was thought to arise from the error on the thermal neutron flux [109].

Isotope	Half-life (Days)	Decay Rate (μBq/kg)	
		Predicted	**Observed**
^{127}Xe	36	420	490 ± 95
^{129m}Xe	8.9	4.1	3.2 ± 0.6
^{131m}Xe	12	25	22 ± 5
^{133}Xe	5.3	0.014	0.025 ± 0.005
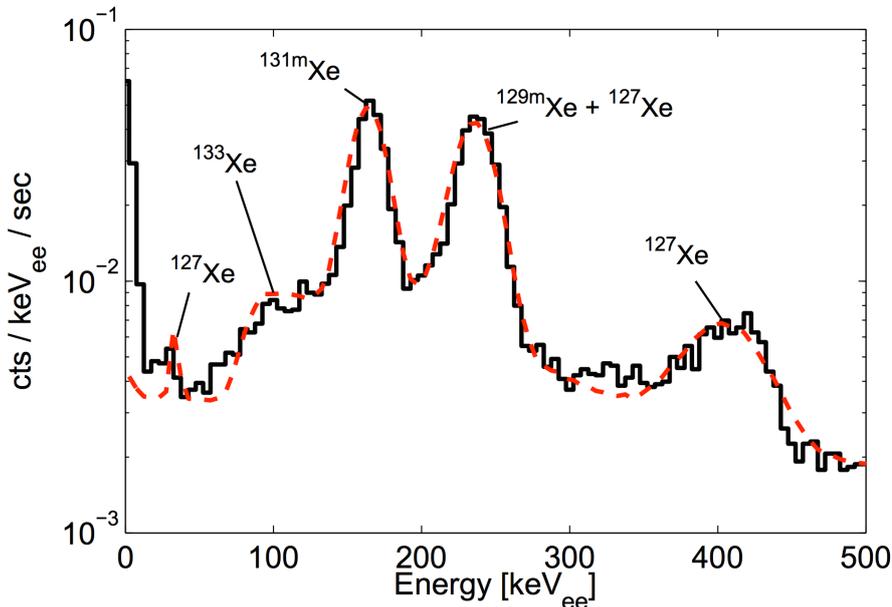
Figure 4.1: Zero-field LUX data taken 12 days after the xenon was brought underground, showing the location of several cosmogenically activated Xe peaks [109]. The red line shows a best-fit spectrum with exponential background. Peak resolution was simulated by measuring from ^{129m}Xe and ^{131m}Xe peaks and extrapolating as $\sqrt{E}$. Relative peak amplitudes relate to branching fractions within the various decays.

43.9 keV, 287.3 keV and 459.4 keV, see figure 4.1. The lowest energy transition may leak into the WIMP search region, leading to a predicted rate of $0.5 \pm 0.02_{\mathrm{stat}} \pm 0.1_{\mathrm{sys}}$ mDRU$_{ee}$ in the WIMP search energy region and fiducial volume, during the 85.3 livedays of run 3.

4.2.2 Neutrons and γ-rays from Detector Materials

As well as the cosmogenic activation of certain elements within detector construction materials, there is a level of intrinsic radioactivity expected in all substances. For LUX, the primary isotopes of concern are ^{40}K and the ^{238}U and ^{232}Th decay chains.

All materials used for construction were screened by high purity germanium detectors at the Soudan Low-Background Counting Facility (SOLO) for their ^{238}U and ^{232}Th content. The

^{238}U and ^{232}Th chains are particularly dangerous as they contain a series of α, β and γ decays, which can cause both electron and nuclear recoils within the WIMP energy region of interest, the latter through (α, n) reactions within detector materials. The uranium and thorium decay chains present a measurement challenge as they contain a number of isotopes with half lives ranging from milliseconds to billions of years. In general, we can assume such chains are in secular equilibrium; this is a situation that arises when a daughter nuclei has a half-life much shorter than its parent, seen in several mother-daughter pairs within both chains. The production rate of the nuclei is constant on the timescales considered and is equal to its decay rate, so the quantity of the nuclei reaches an equilibrium value. If a chain is in secular equilibrium then a measurement of the activity of one part of the chain can be used to infer the activity of the rest of the chain. However, as the chains contain gaseous radon which can emanate out of the material, there is often a break in equilibrium at this point in the chain. Chemical processing can also create a break. Therefore, the chains are often split into early and late; for uranium, isotopes above ^{226}Ra are known as the early chain (^{238}U$_e$), and ^{226}Ra and below are counted as the late chain (^{238}U$_l$). Another break in the uranium chain often occurs at ^{210}Pb (t$_{1/2}$ = 22.3 years) and is sometimes given separately when available. For the ^{232}Th chain, we define the early part of the chain (^{232}Th$_e$) as coming from isotopes above ^{224}Ra and the late part of the chain (^{232}Th$_l$) as coming from isotopes from ^{228}Th and below. See figure 4.2 for a full schematic of both chains. Materials expected to contain ^{40}K and ^{60}Co were also screened for these isotopes. Section 4.2.4 describes ^{60}Co, which is produced cosmogenically in copper. ^{40}K is present at a level of 0.012% in natural potassium, has a long half life of 1.251×10^9 years and undergoes all three types of β-decay, albeit the β^+ decay is very rare:

$$^{40}_{19}\text{K} \rightarrow {}^{40}_{20}\text{Ca} + e^- + \bar{\nu}_e \qquad \text{(BR: 89.28\%)} \qquad (4.5)$$

$$^{40}_{19}\text{K} + e^- \rightarrow {}^{40}_{18}\text{Ar} + \nu_e \qquad \text{(BR: 10.72\%)} \qquad (4.6)$$

$$^{40}_{19}\text{K} \rightarrow {}^{40}_{18}\text{Ar} + e^+ + \nu_e \qquad \text{(BR: 0.001\%)} \qquad (4.7)$$

The β^--decay has an endpoint of 1.311 MeV, whilst the EC decay includes a γ transition of 1.461 MeV. This γ-ray is what is measured to detect ^{40}K in radioassays. Screening results where available for high mass LUX components can be seen in table 4.3.

The expected ER background from γ-rays from detector components can be seen in figure 4.3(a). This was generated with LUXSim (see section 2.6) and came from energy deposits from generated γ-rays normalised to the energies, branching ratios and activities determined by screening. The rate from γ-rays alone was predicted as $1.8 \pm 0.2_{stat} \pm 0.3_{sys}$ mDRU$_{ee}$ [109]. Figure 4.3(b) also includes ^{127}Xe, ^{214}Pb and ^{85}Kr. The background arising from ^{127}Xe is not uniform, despite the isotope being uniformly distributed, but instead decreases exponentially with distance from the active region edges. This is expected as the low-energy decay can occur in coincidence with a 375 keV γ-ray, which if detected, will veto the event. This γ-ray has a mean free path of
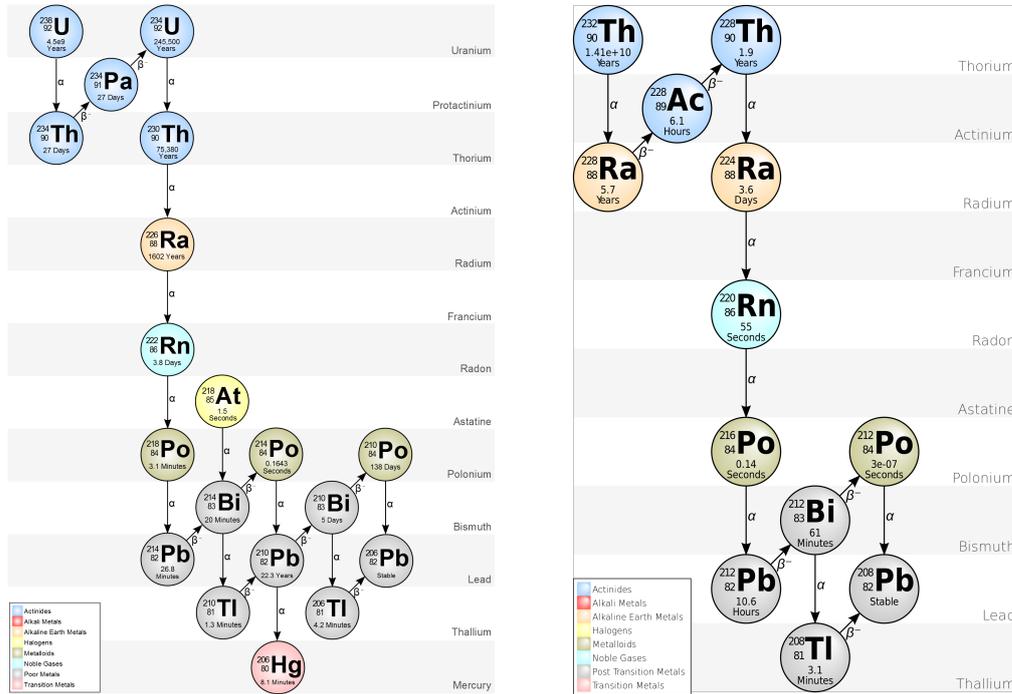
Figure 4.2: Isotopes of the ^{238}U and ^{232}Th decay chains, with their half lives and decay types. It should be noted that certain decays shown here are extremely rare; ^{210}Pb $\rightarrow$ ^{206}Hg $+{}^4_2\alpha$ has a branching ratio of $1.9 \times 10^{-6}\%$ and ^{210}Bi $\rightarrow$ ^{206}Th $+{}^4_2\alpha$ has a branching ratio of 1.32×10^{-4} %.

2.6 cm within the xenon, and so has a greater chance of escaping near the edges of the active region, increasing the background rate here. The total predicted here (including the previous γ-rays) is $2.6 \pm 0.2_{stat} \pm 0.4_{sys}$ mDRU$_{ee}$ [109]. Finally, figure 4.3(c) shows the data, which generally agrees well with the model, except for on the walls of the detector, where there is non-modelled contribution from ^{210}Pb decays. The total observed rate was $3.6 \pm 0.3_{stat}$ mDRU$_{ee}$ [109].

The neutron background for LUX is dominated by neutrons from the PMTs. These are generated through (α,n) reactions by α-decays from ^{238}U and ^{232}Th chain isotopes. The rate of emission was calculated at 1.2 n/PMT/year [109]. Furthermore, the fluorine in PTFE has a high (α,n) cross section. ^{210}Po α-decays on the detector walls may induce neutrons in the PTFE, so this was assessed and predicted as 8.8 n/year. This is only 6% of the total emission from the PMTs, so was not included in the background model.

For neutrons from the PMTs, an energy spectrum for (α,n) reactions was generated using the Neutron Yield Tool [110]. Neutrons from this spectrum were simulated with isotropic emission from both the top and bottom PMT arrays in LUXSim. Results from the simulation led to a predicted number of single scatter events within the WIMP search region of interest of 2.4 - 25 keV$_{nr}$ of 0.06 in the 85.3 livedays of the original run 3 analysis.

Table 4.3: γ-ray radioassay results for LUX detector construction materials samples [109]. Results in the $^{238}U_e$ column have high uncertainty as γ-rays in the early part of the chain with low branching ratios. The ^{238}U values are taken from ^{226}Ra and are the usual numbers reported for ^{238}U-chain when early and late chain is not explicitly stated. ^{232}Th results are taken from a series of lines, usually from ^{228}Ac, ^{212}Bi, ^{212}Pb and ^{208}Tl. *This result is the cosmogenically activated ^{46}Sc, as ^{60}Co is not present in titanium.

Component	Unit	Counting Results (mBq/unit)				
		$^{238}\mathbf{U}_e$	$^{238}\mathbf{U}_l$	$^{232}\mathbf{Th}$	$^{40}\mathbf{K}$	$^{60}\mathbf{Co}$
PMTs	PMT	<22	9.5 ± 0.6	2.7 ± 0.3	66 ± 6	2.6 ± 0.2
PMT bases	base	1.0 ± 0.4	1.4 ± 0.2	0.13 ± 0.01	1.2 ± 0.4	<0.03
Field ring supports (inner panels)	kg		<0.5	<0.35		
Field ring supports (outer panels)	kg		<6.3	<3.1		
Reflector panels (main)	kg		<3	<1		
Reflector panels (grid supports)	kg		<5	<1.3		
Cryostats	kg	4.9 ± 1.2	<0.37	<0.8	<1.6	4.4 ± 0.3*
Electric field grids	kg		1.4 ± 0.1	0.23 ± 0.07	<0.4	1.4 ± 0.1
Field shaping rings	kg		<0.5	<0.8		<0.3
PMT mounts	kg		<2.2	<2.9		<1.7
Weir	kg		<0.4	<0.2		<0.17
Superinsulation	kg	<270	73 ± 4	14 ± 3	640 ± 60	
Thermal insulation	kg		130 ± 20	55 ± 10	<100	

(a) γ-ray sources only (LUXSim). (b) γ-rays, ^{127}Xe, ^{214}Pb and ^{85}Kr (LUXSim). (c) Measured background (data).
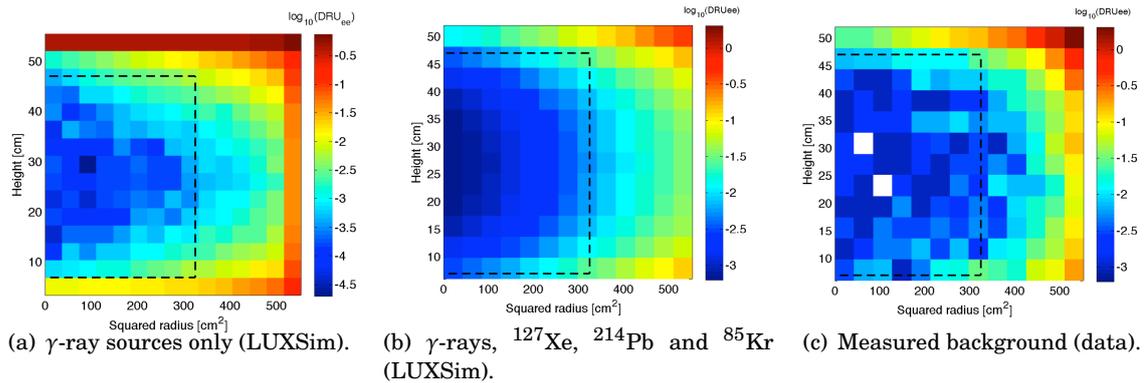
Figure 4.3: Background rates from γ-rays, ^{127}Xe, ^{214}Pb and ^{85}Kr, and measured background data in LUX. The data contains ^{210}Po decays on the PTFE walls that are not included in the model. The dashed black line shows the run 3 fiducial volume. Data is taken from the range 0.9 - 5.3 keV$_{ee}$ and rates are in $\log_{10}(\text{DRU})_{ee}$.

4.2.3 Radon

^{222}Rn and ^{220}Rn from the ^{238}U and ^{232}Th chains can present a large problem for low background experiments. The isotopes are gaseous and so once formed can emanate out of the material they were produced in and into the active detector volume. Therefore, components close to the liquid xenon must be carefully screened to ensure low ^{238}U and ^{232}Th chain activities and thoroughly cleaned of any dust that may also contain contamination to keep radon emanation to a minimum. However, LUX construction materials were not directly screened for radon contamination. The danger of radon lies in the decays of its daughter isotopes; ^{222}Rn produces ^{214}Pb and ^{214}Bi and ^{220}Rn produces ^{212}Pb, which undergo 'naked' or 'semi-naked' β-decay. Naked β-decays emit no
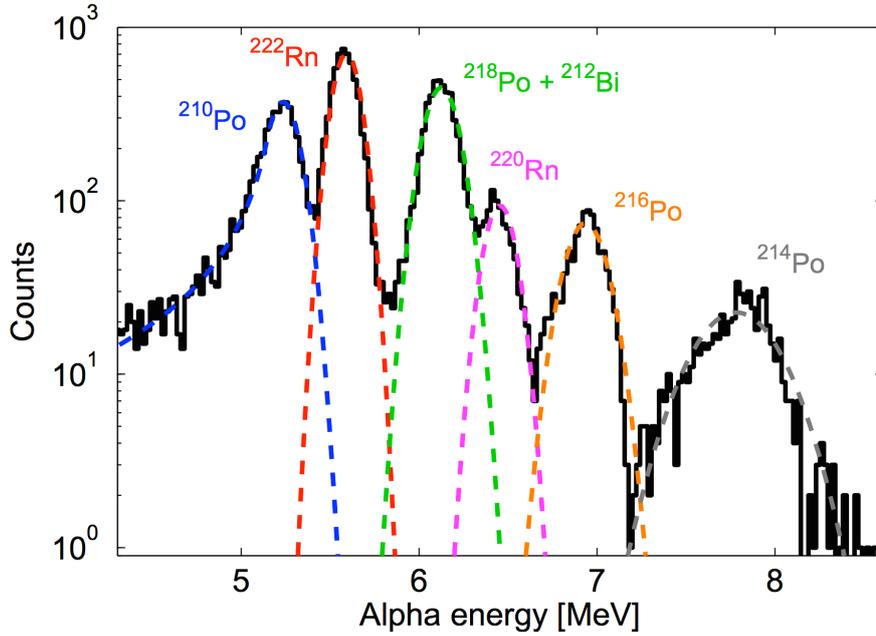
h]

Figure 4.4: High energy peaks from various α-decays taken over 6 livedays, spaced periodically throughout the WIMP search run. All but ^{210}Po are fitted with Gaussians; ^{210}Po is fitted with a Crystal Ball distribution to include a power law tail, which accounts for the energy loss of α-particles as they transit through materials close to the active region [109].

γ-rays, only a β-particle, making it impossible to veto the event. Semi-naked decays do emit a γ-ray, but it is of a sufficiently high energy that it is likely it would escape the active region, again removing the ability to veto the event.

In LUX, radon daughters can be identified by looking for their α-decays, as these are very high energy events with S1s of 40,000 to 90,000 phd. These are easily distinguished from γ-ray S1s, which don't go above 15,000 phd. Section 4.5 demonstrates a more in depth study on α-decays in LUX. Radon daughters below long-lived lead isotopes are found to deposit on the walls of the detector and the cathode, as they are positively charged ions ~90% of the time. This process is known as radon daughter plate-out, and is especially true of isotopes below ^{210}Pb as it has a half life of 22.3 years.

Each peak in figure 4.4 is fitted with a Gaussian and has a mean of the total Q-value of the decay, except for ^{210}Po, which is fitted with a crystal ball distribution. This is because ^{210}Po is deposited on the walls of the detector and decays there, meaning if the α-particle enters the active volume, the recoiling ^{206}Pb nucleus stays in the material surface, depositing its energy there instead. Therefore, it is assumed the energy is the α-particle energy only. α-particles that travel through some of the wall before entering the xenon will lose some energy there, leading to the observed low energy tail of the crystal ball distribution.

The α-particles seen in figure 4.4 are of a much greater energy than the region of interest, and so themselves are not a danger to WIMP search. However, the radon daughters do also generate

a low-energy electron recoil background that does pose a problem. The α-decay rates, shown with their measured energies in table 4.4, allow constraints to be placed on the activities of radon daughters, for example ^{214}Bi and ^{214}Po are bounded by ^{218}Po and ^{214}Po rates as 3.5-14 mBq. ^{212}Pb has been bounded using ^{216}Po as < 2.8 mBq. However, not all of these are of concern; ^{212}Pb has a longer half-life (10.6 hours) and so is expected to have moved from the active region before its decay, and as ^{214}Bi β-decays to ^{214}Po, which is very short-lived ($t_{1/2}$ = 164 μs) and α-decays, the two are expected to overlap within an event window, so the β-decay alone is not expected as a background. This leaves ^{214}Pb as the main concern.

^{214}Pb decays as follows:

$$^{214}_{82}\text{Pb} \rightarrow ^{214}_{83}\text{Bi} + e^- + \bar{\nu}_e \quad \text{(Q= 1024 keV)} \tag{4.8}$$
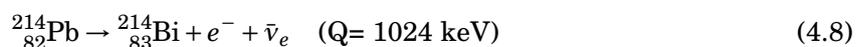
In order to place an upper limit on ^{214}Pb the range 300-350 keV$_{ee}$ was used, as this lies between two ^{127}Xe peaks, and the ^{214}Pb β spectrum rises there. The predicted range of event rate for ^{214}Pb within the WIMP search energy region was 0.11 - 0.22 mDRU$_{ee}$ for run 3 [109].

Table 4.4: Radon chain daughters and their energies and event rates measured in the active volume during the first LUX WIMP search run of 85.3 livedays. As in figure 4.4, the data was taken over 6-livedays periodically spaced throughout the run. All energies are Q-values except for ^{210}Po, which is just the α-particle energy, as expected when the emission occurs on the walls of the detector.

Decay Chain	Isotope	Energy (MeV)	Measured Energy (MeV)	Half-life	Event Rate (mHz)
^{238}U	^{222}Rn	5.59	5.59 ± 0.08	3.8 d	17.9 ± 0.2
	^{218}Po	6.16	6.12 ± 0.10	3.1 m	14.4 ± 0.2
	^{214}Po	7.84	7.60 ± 0.2	160 μs	3.5 ± 0.1
	^{210}Po	5.30	5.22 ± 0.09	140 d	7.2 ± 0.2 (cathode) 14.3 ± 0.2 (walls)
^{232}Th	^{220}Rn	6.41	6.47 ± 0.09	56 2	17.9 ± 0.2
	^{216}Po	6.91	6.95 ± 0.1	0.15 s	14.4 ± 0.2
	^{212}Bi	6.21	6.12 ± 0.10	61 min	3.5 ± 0.1
	^{212}Po	8.83	-	0.3 μs	-

4.2.4 Cosmogenic Activation of Materials

LUX was assembled in the Sanford Surface Laboratory over a period of two years; during this time any construction materials were exposed to cosmic rays. For example, the muon flux has been measured at $1.149 \pm 0.017 \times 10^{-2}$ s^{-1}cm^{-2}sr^{-1} [86]. This causes activation of certain elements within detector materials; of most concern for detector backgrounds are ^{46}Sc, generated by either muon capture or an (n,p) reaction in the titanium cryostat, which decays by β-decay, with

$t_{1/2} = 83.83$ days, endpoint energy 0.357 MeV:

$$^{46}_{22}\text{Ti} + \mu^- \rightarrow ^{46}_{21}\text{Sc} + ^1_0 n + \nu_\mu \tag{4.9}$$

$$^{46}_{22}\text{Ti} + ^1_0 n \rightarrow ^{46}_{21}\text{Sc} + ^1_1 p \tag{4.10}$$

$$^{46}_{21}\text{Sc} \rightarrow ^{46}_{22}\text{Ti} + e^- + \bar{\nu}_e \tag{4.11}$$

also emitting 2 γ-rays of energies 889 keV and 1,121 keV. For the LUX cryostat, a sample of titanium was screened at the SOLO facility after spending two years underground. It was then allowed to spend 6 months on the surface at Sanford, before being sent back to SOLO for another measurement, which yielded 4.4 ± 0.3 mBq/kg [109], consistent with predictions from the ACTIVIA package [111]. It was then estimated that the total ^{46}Sc activity in the cryostat was 1.3 Bq immediately after moving underground.

The other isotope of concern is ^{60}Co, generated by a (n, α) reaction, which decays by β-decay, with $t_{1/2} = 1,925$ days, endpoint energy 0.318 MeV:

$$^{63}_{29}\text{Cu} + ^1_0 n \rightarrow ^{60}_{27}\text{Co} + ^4_2 \alpha \tag{4.12}$$

$$^{60}_{27}\text{Co} \rightarrow ^{60}_{28}\text{Ni} + e^- + \bar{\nu}_e \tag{4.13}$$

Similarly to ^{46}Sc, the decay is usually accompanied by one or two γ-rays. LUX contains 620 kg of copper, and the expected activation rate was 62 ± 29 kg^{-1} day^{-1} at sea level, which then must be scaled up by with the $\times 3.4$ neutron flux altitude correction factor, to 210 ± 100 kg^{-1} day^{-1}. The exposure above ground was 800 days, leading to a decay rate of 1.0 ± 0.5 mBq/kg once the copper was taken underground.

4.3 External Backgrounds

4.3.1 Muon-Induced Neutron Background

Cosmic ray muons can generate neutrons in materials by reactions such as:

$$\mu^- + X \rightarrow X' + n \tag{4.14}$$

$$\mu^- + p \rightarrow n + \nu_\mu \tag{4.15}$$

$$\gamma + X \rightarrow X' + n \tag{4.16}$$

$$\pi + X \rightarrow X' + n \tag{4.17}$$

$$\tag{4.18}$$

where X and X' are different nuclei. The simplified reactions represent, in order: a spallation reaction induced by a a muon-nuclear interaction via the exchange of a virtual photon, resulting in nuclear disintegration, muon capture, photo-nuclear interactions in electromagnetic showers that were triggered by a muon, and similarly production in hadronic cascades initiated by a muon.

The Davis laboratory at SURF recieves a muon flux of $(4.4 \pm 0.1) \times 10^{-9}$ cm^{-2} s^{-1}[112], with an average muon energy of 321 GeV [113]. The neutron flux induced by these muons in the cavern rock has also been measured at several underground sites, although not at the 4850 level, however the flux and energy distributions can be fitted using depth-dependent functions [112]. Using simulations, this gives a neutron flux of $(0.54 \pm 0.01) \times 10^{-9}$ cm^{-2} s^{-1}. Attenuation by the water tank reduces the integrated neutron flux to 1×10^{-7} s^{-1}, which would produce a single scatter nuclear recoil background rate of 60 nDRU in the 3.4 - 25 keV$_{nr}$ range, within a 100 kg fiducial volume. There is further muon induced production of neutrons within the water, and this was also estimated using Monte Carlo simulations using GEANT4. The integrated flux at the LUX outer cryostat was calculated as 6.3×10^{-7} s^{-1}, resulting in 120 nDRU$_{nr}$. The total of 180 nDRU$_{nr}$ tells us that we may expect 0.1 WIMP-like event per 100 kg in a year of data taking. Therefore, these were not a problem in run 3, but run 4 is over 300 days long. However, the LUX water tank is instrumented with veto PMTs that can identify Cherenkov radiation from passing muons. This has been used successfully in run 4 to veto events occurring directly after such Cherenkov light is detected.

4.3.2   Cavern Rock

The rock in the Davis cavern contains naturally occurring uranium, thorium, their daughters, and potassium, all of which emit high energy γ-rays. The primary purpose of the LUX water tank is to prevent these γ-rays from reaching LUX; the attenuation length of a 1 MeV γ-ray in water is ~14 cm, and there is 3.5 m of water between the edge of the water tank and the edge of the LUX cryostat, reducing the flux to a factor 8×10^{-9}. Further shielding is provided by a steel pyramid located beneath the detector in the cavern floor, which is the rock closest in proximity to the detector. This reduces the flux by a further factor of 40. The background from cavern γ-rays is expected to be less than 6×10^{-3} events in the 1.3 - 8 keV$_{ee}$ WIMP search energy range in 30,000 kg·days, assuming a highly radioactive cavern composed entirely of rhyolite, producing 9 γ/cm^2/s at the water shield outer edge [114].

4.4   Run 4 Backgrounds

For run 3, as previously stated, ^{210}Po decays on the wall were not included in the background model. A fiducial cut of R<18 cm meant that wall events were cut away for WIMP search results. For run 4, however, a model was created to allow the fiducial volume to be extended towards the wall of the detector.

β-decays of ^{210}Bi and ^{206}Pb recoils from ^{210}Po α-decay can be mis-reconstructed as some distance from the wall, especially if their S2 sizes are small, which is likely due to charge loss in the PTFE. This can also cause events to leak downwards in the log(S2/S1) discrimination space, which becomes dangerous as once below the NR band mean, they become WIMP candidates. To

account for this, the leakage of wall events below the NR band mean was parameterised as PDFs depending on radius, drift and S2 size:

$$P_w(r, S1, S2, z) = P_r(r|S2) \times P_d(S1, S2, z) \tag{4.19}$$

$P_r(r|S2)$ describes the radial distribution of S2s and is characterised using using events below the NR mean with S1 > 55 phd, whilst $P_d(S1, S2, z)$ is inferred using a population greater than 1 cm from the fiducial boundary outwards. This ensures that the model is constructed using population samples outside of the region of interest. The radial leakage becomes much higher at small S2, extending over 2 cm further into the detector when comparing S2s of 100 - 150 phd and 600 - 1000 phd.

The wall model can be used to predict the number of events below the NR mean expected when extending the fiducial volume to higher radii, depending on S2 size and drift time. For run 4, this allowed the fiducial volume to be extended to 3 cm inwards from the measured position of the wall The number of predicted events was 8.4 ± 0.35, but as the profile likelihood ratio (PLR) analysis used to set limits uses spatial information, these events have a low signal likelihood.

Other changes in the background model for run 4 were the removal of ^{127}Xe as an ER background, as it had all decayed away by the beginning of run 4 data taking, and the addition of ^{8}B solar neutrinos as an NR background. The number expected from these neutrinos was low, 0.16 ± 0.03 in the whole exposure, but this was considered significant enough to included in the PLR. See section 5.3.5 for the run 4 results.

4.5 A Search for ^{210}Bi α-decays in LUX

4.5.1 Alpha Decays and Motivation

As previously discussed in section 4.2.3, ^{214}Pb, ^{214}Bi and ^{212}Pb are problematic backgrounds as they undergo naked or semi-naked β-decay, which are more difficult to veto. Concern arose about ^{210}Bi in particular due to the Borexino solar neutrino experiment [115] and the KamLAND-Zen ^{136}Xe $0\beta\beta\nu$ search [116] reporting ^{210}Bi backgrounds within the bulk of their detectors. This is unexpected, as it is below long-lived ^{210}Pb, which should have plated-out on detector walls. Therefore, the possibility of mobility of ^{210}Bi was investigated using α-decays.

^{210}Bi undergoes a naked beta decay with a branching ratio $BR_\beta = 0.9999986$ and an endpoint of 1.16 MeV:

$$^{210}_{83}\text{Bi} \rightarrow \, ^{210}_{84}\text{Po} + e^- + \bar{\nu}_e \tag{4.20}$$

and a much rarer α-decay, with a branching ratio $BR_\alpha = 1.4 \times 10^{-6}$, and an energy of 4.7 MeV:

$$^{210}_{83}\text{Bi} \rightarrow \, ^{206}_{81}\text{Tl} + \, ^4_2\alpha \tag{4.21}$$

Whilst the ^{210}Bi β-decay has not been problematic in LUX, it could pose a serious background for LZ within the fiducial volume, and would only be suppressible by ER/NR discrimination and

PLR modelling. Identifying its presence in LUX at this stage would allow for more stringent cleanliness and radioactivity requirements for the construction of the successor experiment LZ, particularly on ^{210}Pb plate-out levels, in order to mitigate the background.

LUX has been found to have ^{210}Po on the walls of the detector, producing an event rate of 14.3 ± 0.2 mHz, therefore there must be a large quantity of ^{210}Bi present producing it. Trying to determine if the β-spectrum is visible in LUX background data is difficult because of several other backgrounds within the same region, whereas the α-decay would be a much more recognisable signature, albeit very rare. The aim of this study was to set an upper limit on the quantity of ^{210}Bi within LUX using α-decays in order to inform cleanliness and screening protocols for the construction of LZ to ensure suitably low quantities of ^{210}Bi as to not adversely affect WIMP sensitivity.

4.5.2 Data Selection

The data used was taken during run 3 of LUX WIMP search operation and comprised a total of 80.2 livedays. Events were selected using the standard LUX golden selection. To initially select α events, the following cuts were applied:

- S1 > 10,000 phd
- S1 < 80,000 phd
- S2 > 2,000 phd
- S2 < 100,000 phd

This selects a population of events on the MeV scale, whose very large S1s can only have been produced by α-decays.

4.5.3 Energy Calibration and PMT Saturation

For α-decays, it is not possible to use the standard LUX energy calibration, given by:

$$E_r = 13.7\text{eV}\left(\frac{S1}{g_1} + \frac{S2}{g_2}\right) \tag{4.22}$$

This is due to two reasons: α-particle S2s show large quenching, and any α-decays happening on the walls of the detector can show significant charge loss, again reducing their S2s. Therefore, for α-decays in LUX, the calibration is done only using the S1 pulse area. However, this can still pose a problem, as these very large S1s can saturate the bottom array. S1s have more light detected in the bottom array due to the reflection of photons at the liquid-gas interface, and so especially large S1s close to the bottom array may cause saturation. Figure 4.5(a) shows the waveform of an S1 that has caused saturation in the bottom array; PMTs that are saturated become flat at the top of the S1. Note that different PMTs may saturate with different amounts of light.

The saturation effect was investigated by obtaining top and bottom array areas for S1s. Array information is usually available in reduced LUX data for S2s, because of the tendency of S2s to
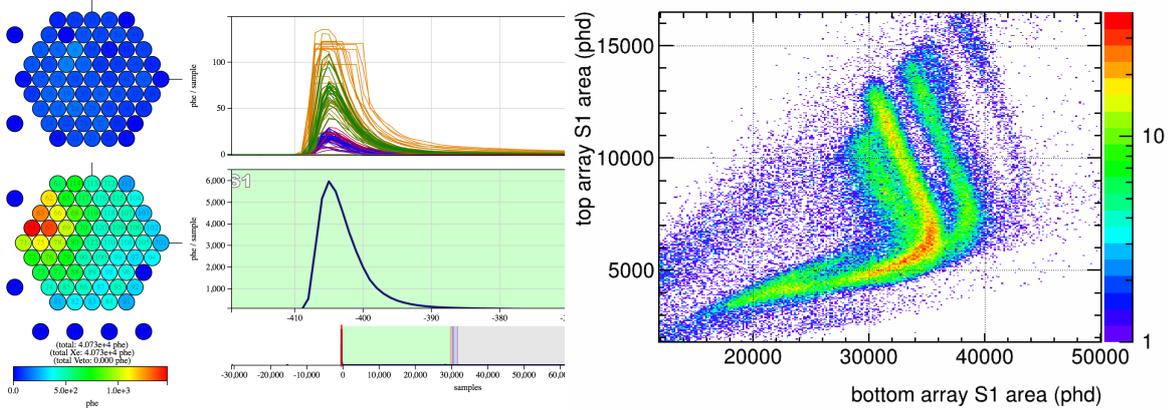
(a) A large S1 from an α-decay showing saturation in several PMTs.

(b) S1 light in the top array vs the bottom array for large S1 areas.

Figure 4.5: Effects of PMT saturation on S1s. Left shows the waveform of a saturated S1. The individual PMT waveforms are shown on the top and many are flattened at their differing saturation values. Right shows S1 areas in the top and bottom arrays. Each line is a different α-particle with the strongest being ^{222}Rn. The lines remain linear and show how the bottom array area increases as the top array decreases as events move downwards in the detector, but there comes a point where the lines curve round, showing saturation in the bottom array. The bottom array area decreases as the solid angle of the light reaching the array decreases, allowing less PMTs to reach saturation.

saturate the top array, but data had to be reprocessed to obtain this information for S1s. Figure 4.5(b) demonstrates how strong this effect can be. The diagonal bands show each α-decay as events move downwards in the detector towards the bottom array; the light collected in the top decreases and that in the bottom increases accordingly. However, there comes a point where the lines all curve around and show a decrease in the bottom array. This demonstrates saturation; the reason the detected light in the bottom array does not just level off at constant is that as events move closer and closer to the PMT array, the solid angle of the light hitting the PMTs decreases, so it is concentrated in less PMTs, meaning less PMTs are able to reach saturation, and the overall light detected goes down.

A scaling factor was determined using the largest peak, which is uniform throughout the detector and known to be ^{222}Rn, with an energy of 5.59 MeV. For run 3 data, a Gaussian fit gave a mean of 42,852 phd, and so all other S1s were scaled by a factor of $5.59/48,852 = 1.3 \times 10^{-4}$ to convert to MeV. For run 4, the mean of the ^{222}Rn peak had decreased to 37,280 phd, so the scaling factor was changed accordingly.

In order to mitigate this saturation effect, S1 areas were plotted against drift time using run 3 data, as in figure 4.6(a). S1 areas stay fairly constant above a drift time of 200 μs, so events were only selected from the top half of the detector.

The saturation effect was reassessed for run 4, as there were changes in the electric field and event reconstruction that may change the impact. Figure 4.6(b) demonstrates that whilst S1 sizes have shrunk overall, the 200 μs drift cut is still sufficient to remove saturated S1 events.

107

(a) Run 3 - 80 livedays

(b) Run 4 - 360 livedays
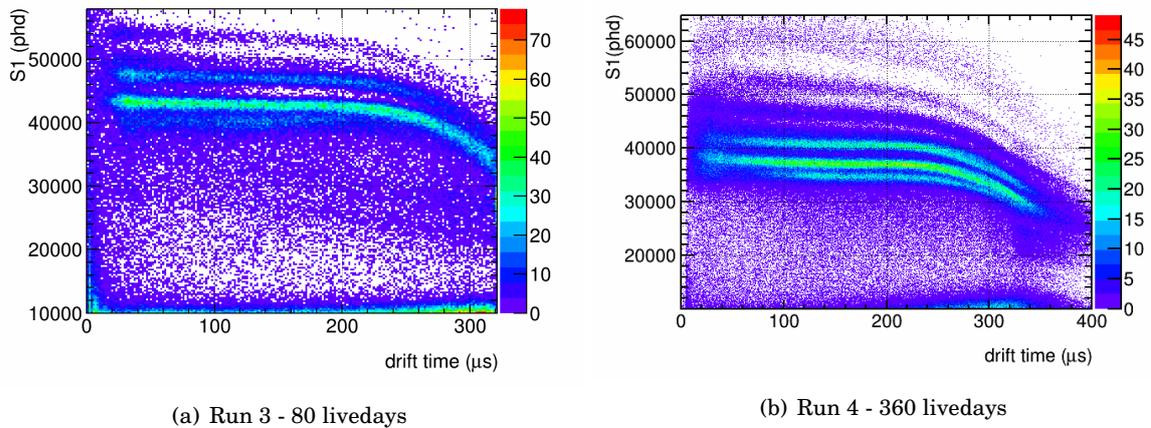
Figure 4.6: S1 areas vs drift time. There is a clear decrease at high drift times due to poorer light collection efficiency and array saturation. The various α populations can be identified as the lines that at low drift time are constant in S1, with the brightest being ^{222}Rn.

4.5.4 Run 3

After applying the drift time cut, the α-particles were found to be well defined in S1 vs S2, see figure 4.7. Cutting into the bulk of the detector with the fiducial radial cut of R<18 cm removes the ^{210}Po population from the lower left, confirming that as expected it has plated-out onto the walls.
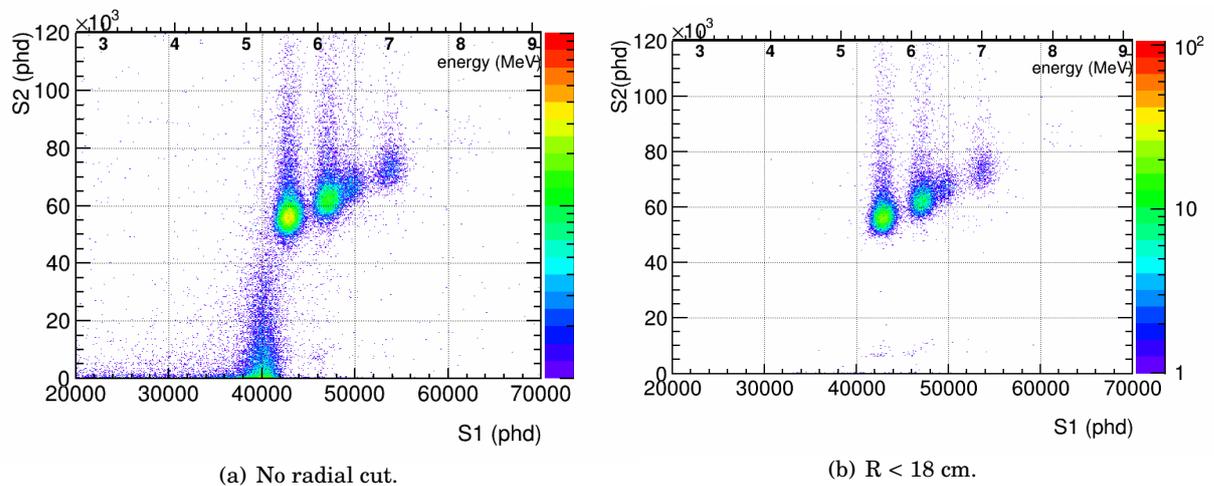
(a) No radial cut.

(b) R < 18 cm.

Figure 4.7: α-particles shown in S1 vs S2. Here, phd stands for detected photons. The effect of the walls on S2 size is noticeable in the ^{210}Pb population. The top axis shows the energy in MeV as calibrated using the ^{222}Rn peak. .

Figure 4.8 demonstrates the energy peaks for both the whole active xenon volume and after a radial cut of R<18 cm. Gaussians have been fitted for all peaks except ^{210}Po, where a crystal ball function has been used. In figure 4.8(b), the ^{214}Po peak was not fitted as there were so few data points.

After the radial cut, there are two events close to the 4.7 MeV expected from ^{210}Bi α-decay, with energies of 4.59 MeV and 4.69 MeV.

(a) *All events after S1, S2 and drift cuts are applied.*   (b) *Remaining events within R < 18 cm.*
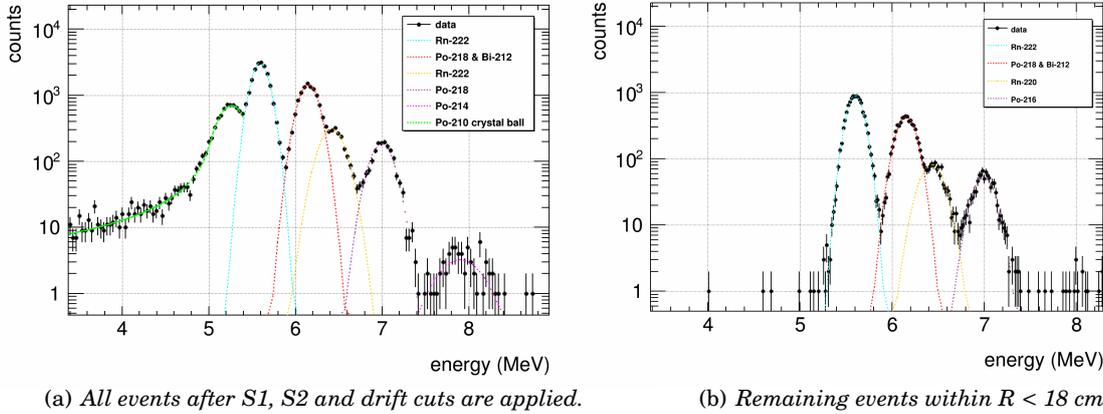
Figure 4.8: α-decay signatures as peaks in energy in LUX run 3 data. The energy is calibrated using the ^{222}Rn peak. All fits are Gaussian except for the ^{210}Po peak, which follows a crystal ball distribution due to energy loss to the walls. Applying a fiducial cut removes the vast majority of the ^{210}Po peak.

4.5.5   Run 4

376.3 livedays taken during run 4 were also used with the same analysis scripts as the run 3 data. The energy calibration Figure 4.9 shows the α populations in S1 vs S2. The effects of the increase in field strength in run 4 are apparent; when comparing with figure 4.7, S1s are smaller and S2s are significantly larger (note the change in the y scale). As already stated, this was taken into account for the energy calibration, as ^{222}Rn S1s had decreased by 13%. This is expected with a higher electric drift field as less recombination takes place, leading to a reduced primary scintillation signal but a larger number of extracted electrons and a higher electron extraction efficiency.

Then α peaks are as expected larger due to the longer livetime of the run 4 data, and the ^{214}Po peak is much better resolved. A peak became visible on the walls at ~8.6 MeV, which can be seen in figure 4.10(a) and is thought to be the α-decay of ^{212}Po from the ^{232}Th chain. This was initially missed in run 3 data due to the lack of statistics in the energy window.

For the run 4 fiducial volume, there were no observed events around 4.7 MeV, compared to the 2 that was present in run 3.

4.5.6   Bi-210 Search

4.5.6.1   Cut and Count

For run 3, after the radial cut, there are two events in the region around 4.7 MeV, one at 4.59 MeV and one at 4.69 MeV. An example is shown in figure 4.11. Both events occur at a radius of 14 cm.
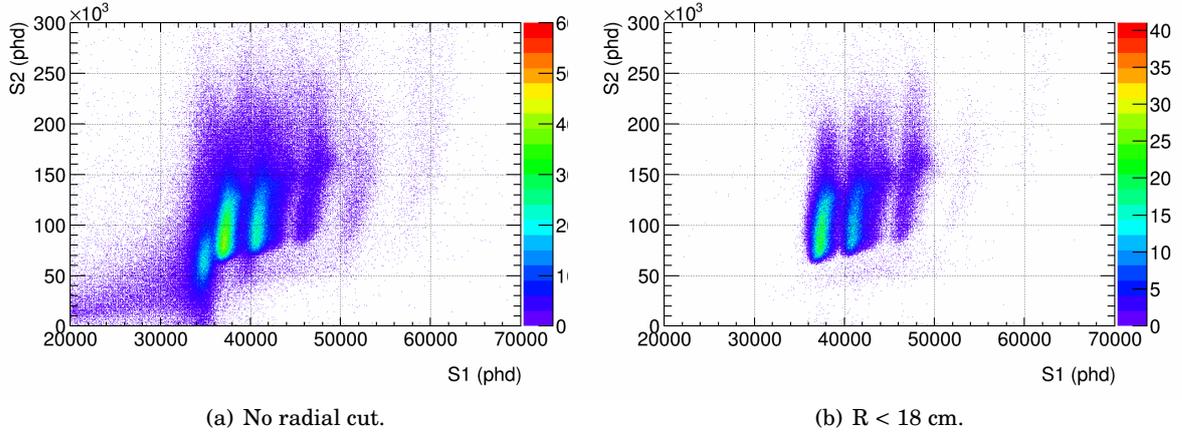
(a) No radial cut.

(b) R < 18 cm.

Figure 4.9: α-particles shown in S1 vs S2. Here, phd stands for detected photons. The effect of the walls on S2 size is noticeable in the ^{210}Pb population, and this is seen to disappear inside the bulk xenon.

(a) *All events after S1, S2 and drift cuts are applied.*

(b) *Remaining events within R < 18 cm.*
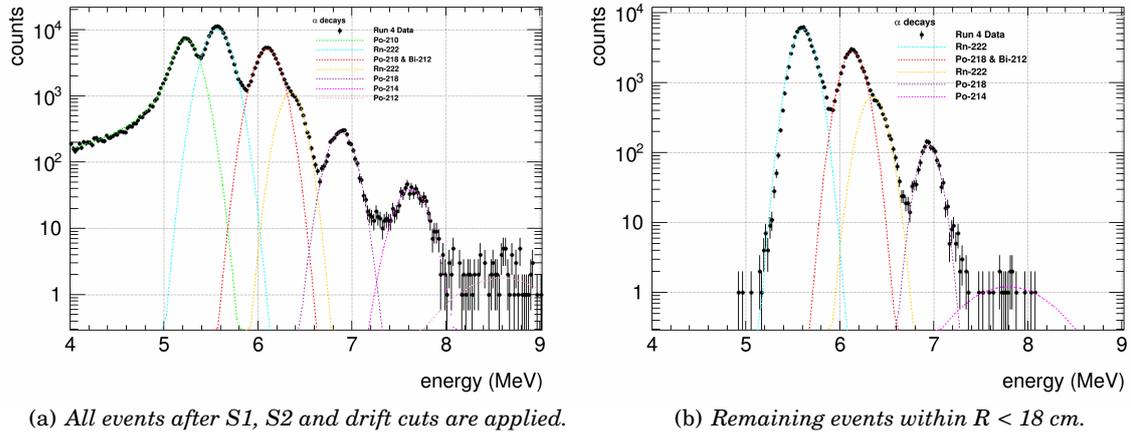
Figure 4.10: α-decay signatures as peaks in energy in LUX run 4 data. The energy is calibrated using the ^{222}Rn peak. All fits are Gaussian except for the ^{210}Po peak, which follows a crystal ball distribution due to energy loss to the walls. Applying a fiducial cut removes the vast majority of the ^{210}Po peak.

Otherwise, the region from 4 to 5.1 MeV is very clean and only contains 4 events total, the other two being at 4 MeV and 5 MeV.

For run 4, there are no events close to 4.7 MeV. In the window from 4 to 5.1 MeV there are 3 events, with energies of 4.93 MeV, 4.97 MeV and 5.05 MeV. These are likely to be down-fluctuations or poorly resolved/resconstucted ^{210}Po events.

4.5.6.2   Fitting

There is no obvious excess in the data that suggests the presence of ^{210}Bi α-decays. If there are any events present, they are hidden within the tail for ^{210}Po. Attempts to fit to the data with crystal ball distributions for both ^{210}Po and ^{210}Bi favour zero signal, where the value of the χ^2/dof of the total fit is 1.7. A crystal ball distribution was also used for ^{210}Bi as these α-decays would
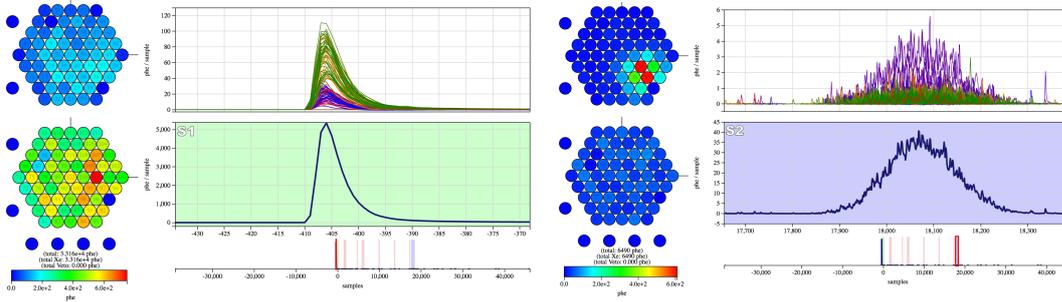
(a) *S1 of area 35,977 phd, equivalent to 4.69 MeV.*

(b) *S2 of area 8,238 phd, significantly quenched.*

Figure 4.11: S1 and S2 waveforms showing individual PMTs (top) and summed signal (bottom) from an event within the fiducial volume at 4.69 MeV. The event is fairly clean, with just a few single electrons emitted between the S1 and S2, and the S1 and S2 appear to be correlated in position. The position reconstruction algorithm used by LUX locates the event at a radius of 14.4 cm. .

also be happening on the walls, and so would have the same tail characteristic of energy loss to the detector materials.

It is possible to force signal into the fit by setting limits on the signal variable. The parameters were varied in order to determine a maximum possible signal. Two examples using run 3 data are shown in figure 4.12. Figure 4.12(a) demonstrates an example fit where the crystal ball parameters other than the mean have been constrained to be similar for both the ^{210}Po and the 4.7 MeV peak. This fit only increases the χ^2/dof from 1.7 (no signal) to 1.8. Figure 4.12(b)
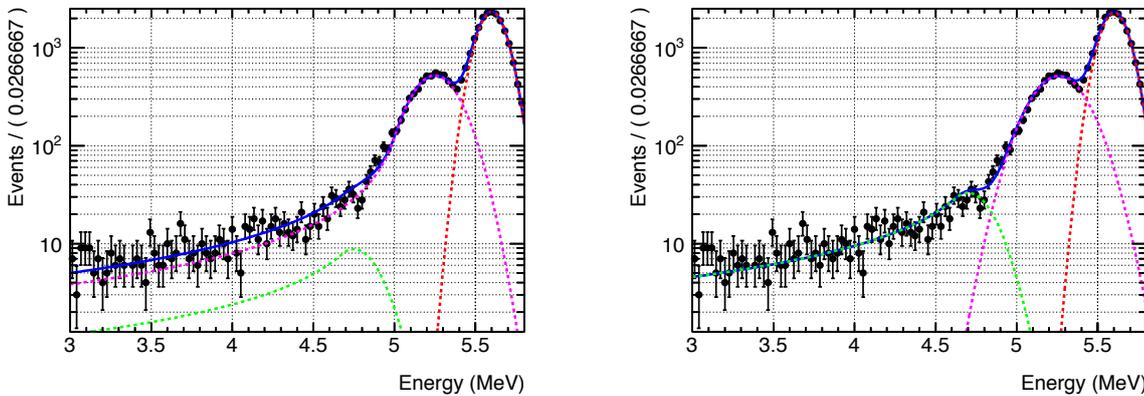
(a) This fit enforces similar parameters except for the mean for both crystal ball functions as both would lose energy to the walls by the same mechanism. χ^2/dof = 1.8

(b) This fit minimises the amount of ^{210}Po in the tail and instead maximises the peak at 4.7 MeV. This is likely to be unrealistic. χ^2/dof = 2.1

Figure 4.12: Examples of fits attempting to force signal in order to find an upper limit. The ^{222}Rn signal peak at 5.59 MeV is fitted with a Gaussian (dashed blue), ^{210}Po fitted with a crystal ball (dashed red) and the forced signal peaked at 4.7 MeV.

shows the maximum amount of signal that can be fitted into the energy range, but it causes the exponential tail of the ^{210}Po to disappear, which is unphysical, and the χ^2/dof is worsened to 2.1. Note that this was not repeated for run 4 data, as it was deemed to be a fairly arbitrary process

that would not have benefited from better statistics. The ^{210}Po tail in run 4 data did not show signs of containing any excess.

Another consideration is the existence of several other α-decays with similar energies to the ^{210}Bi α. Table 4.5 lists these. These are all above radon in the chain and so would be expected to still be located inside detector materials, rather than in the xenon or deposited on the walls. However, it is possible that some of these may contribute to the events in the tail of the ^{210}Po distribution. A very rough estimate of the contribution of these early chain isotopes can be made.

If we assume secular equilibrium within the ^{238}U chain and thus assume the same activity for ^{234}U, ^{230}Th and ^{226}Rn, then we expect 0.0000014 ^{210}Bi α-decays for every three α-decays from the other three isotopes. However, as previously stated, these isotopes are expected to be internal to detector components rather than the deposition on the walls characteristic of radon daughters. In order to make an estimate of the proportion of these decays that might make it to the xenon, we use the range of a 4.7 MeV α-particle within PTFE, 0.002 cm. The LUX PTFE is $\sim$ 1 cm thick, so we can then assume if the radioactive contamination is distributed equally, 0.2% of decays may occur close enough to the inner wall to reach the active xenon. Of these, $\sim$ 50% may be going towards the xenon. Therefore, we now assume there are 0.0000014 ^{210}Bi α-decays for every 0.5×0.002×3=0.003 α-particles from the early chain isotopes. Therefore, ^{210}Bi could comprise only 0.5% of any α-decays around these energies.

Table 4.5: α-decays within the ^{238}U chain in the range 4.6 - 4.8 MeV and their intensities. *Overall α branching ratio of ^{210}Bi is 0.00014%.

Isotope	Energy (MeV)	Intensity (%)
^{234}U	4.60	0.5
	4.72	28.42
	4.78	1.38
^{230}Th	4.62	23.4
	4.69	28.42
^{226}Ra	4.60	5.55
	4.78	94.45
^{210}Bi	4.66	$\sim$ 60*
	4.69	$\sim$ 40*

4.5.6.3   Upper Limits

An upper limit on the activity of ^{210}Bi on the walls of the detector (A, in Bq/m^2), using α-decays can be found using:

$$A_{\text{Bi-210}} < \frac{N}{A \cdot L \cdot f_\alpha} \tag{4.23}$$

where N is the number of events, L is the livetime of the data used in seconds, A is the relevant surface area and f_α is the branching fraction of the ^{210}Bi α-decay.

Table 4.6 shows the results of this calculation using two fits (the unphysical maximum forced signal, the minimum forced signal that only worsens χ^2/dof by 0.1 (for run 3 only) and radial cuts

of 18 cm. The mass of xenon used for the full radius was 163 kg, for an 18 cm radius 88 kg, and the surface area of PTFE covering the height of the detector used was calculated as 0.4562 m^2.

For the cases where $N = 0$, $N = 1$ has been used in the calculation, which uses the assumption that the maximum possible number of ^{210}Bi β-decays has been 714,285, which is 1/BR$_\alpha$ and the number expected for one α-decay.

Table 4.6: Calculated upper limits on ^{210}Bi activity in LUX depending on different assumptions. All limits are given to 2 significant figures. *This is the maximum signal that preserves the exponential tail of the ^{210}Po distribution. **This is the minimum signal that increases χ^2/dof by 0.1. † These numbers use the assumption that 0.5% of decays are ^{210}Bi.

Assumption	Number of events (N)	Activity in LXe (mBq/kg)	Activity on Walls (mBq/m^2)
Run 3: Maximum forced signal* (walls)	650	<410 (<0.3)†	<150,000 (<59)†
Run 3: Minimum forced signal** (walls)	219	<140 (<0.1)†	<49,000 (<20)†
Run 3: Best fit (walls)	0	<0.63	<230
Run 3: R<18 cm	2	<2.3	<450
Run 4: Best fit (walls)	0	<0.13	<48
Run 4: R<18 cm	0	<0.25	<48
Combined Runs	2	<0.41	<79

4.5.7 α-Decay Rates

Table 4.7 shows the measured total rates and energies of each α-decay for both runs. These are lower than those reported in section 4.2.3, but this is understood to be for two reasons: firstly, these are the rates for just the selected volume using a drift time of less than 200 μs. This was not simply scaled up because rates at the bottom of the detector are expected to be higher than at the top as positively charged isotopes will drift towards the cathode and decay there. Secondly, this analysis only used single scatter data, and several of the isotopes may decay in coincidence (within the same event window) with their mother due to short decay times, so will not have been included here.

Note that ^{218}Po and ^{212}Bi are too close in energy to be resolved separately, so are given the same rate from one fitted peak.

Some of the measured energies appear to shift between run 3 and run 4 and this is visualised in figure 4.13. The energy shift appears to be more significant at higher energies, and also when including the walls. As run 4 had a large radial field component (see section 5.3.5), this is expected to be an energy reconstruction issue due to field distortions. The peaks were still fitted successfully for both runs.

To determine the effect of various radial cuts on the rate of α-decays, seven energy histograms were filled for radii between the 18 cm fiducial cut and the wall at 24.5 cm. Each histogram was then fitted for all peaks simultaneously using Roofit. The calculated signal was then converted to a rate by dividing by the livetimes in seconds. Figure 4.14 shows the results for both run 3 and

Table 4.7: α-decay measured energies and rates shown for both run 3 and run 4 data for the detector volume used in this analysis. Errors are taken from Roofit.

Chain	Decay	Energy (MeV)	Measured Energy (MeV)		Measured Rate (mHz)	
			Run 3	Run 4	Run 3	Run 4
^{238}U	^{222}Rn	5.59	5.5951 ± 0.0007	5.5734 ± 0.0003	2.80 ± 0.02	3.99 ± 0.01
	^{218}Po	6.16	6.145 ± 0.001	6.0932 ± 0.0005	1.54 ± 0.02	2.14 ± 0.01
	^{214}Po	7.84	7.93 ± 0.03	7.593 ± 0.008	0.010 ± 0.001	0.023 ± 0.001
	^{210}Po	5.30	5.20 ± 0.05	5.2124 ± 0.0005	1.2 ± 0.02	3.49 ± 0.01
^{232}Th	^{220}Rn	6.41	6.467 ± 0.003	6.408 ± 0.002	0.30 ± 0.02	0.212 ± 0.003
	^{216}Po	6.91	7.001 ± 0.003	6.880 ± 0.002	0.24 ± 0.01	0.118 ± 0.002
	^{212}Bi	6.21	6.145 ± 0.001	6.0932 ± 0.0005	1.54 ± 0.02	2.14 ± 0.01
	^{212}Po	8.83	8.82 ± 0.09	8.62 ± 0.03	$(3 \pm 2) \times 10^{-4}$	$(2.3 \pm 0.3) \times 10^{-3}$

run 4. Overall, they show similar patterns, although there are some differences, discussed below. Figure 4.15 shows an example of the fitted histograms used.

As expected, isotopes above the long-lived lead isotopes (^{222}Rn, ^{220}Rn, ^{218}Po, ^{216}Po, ^{214}Po and ^{212}Bi) are distributed uniformly throughout the detector, increasing gradually as the contained xenon mass increases. ^{210}Po, however, increases rapidly as the wall is approached. For run 3, it starts to appear above 22 cm, whilst for run 4, 20 cm, but in the latter it falls to a lower rate within the bulk xenon. This may just be an artefact of fitting to so few events. Interestingly, ^{214}Po seems to show a stronger radial dependence in run 4 than run 3. However, as previously mentioned, the data used contained only single scatters, many ^{214}Po decays are missed - ^{214}Po decays rapidly with a half-life of 0.1643 seconds and so is often the same event waveform as the decay of its mother, ^{214}Bi. Therefore, statisics are generally low, especially for run 3, see figure 4.8. Collectively, the three isotopes between radon and ^{214}Po live for almost an hour so it is also possible some isotopes are deposited on the walls in that time. ^{212}Po has an even shorter half-life, 3×10^{-7} seconds, so even fewer are captured in single scatter events without the preceding ^{212}Bi β-decay. Enough are seen, however, to see the radial dependence, which is expected as ^{212}Po is post-^{212}Pb in the thorium chain, which has a half-life of 10.6 hours.

4.5.8 Implications for LZ

Alongside the α-decay study above, two other related searches were done: a search for the 46.5 keV γ-line from ^{210}Pb, and a search for the β-decay of ^{210}Bi. An upper limit was placed on the activity of ^{210}Pb on the walls of LUX as 1.67 μBq/cm^2, and an upper limit on the total radioactivity of ^{210}Bi in the detector volume as 0.26 mBq. The upper limit on ^{210}Pb was converted to a total activity of 12.3 mBq using the surface area of the PTFE in LUX (12 panels of 12.674 cm width, cathode to gate distance 48.32 cm, 7,349 cm^2).

The standard radioactivity formula,

$$A = \lambda N \tag{4.24}$$

where λ is the decay constant ($\lambda = \ln(2)/t_{1/2}$) and N is the number of atoms that produces the
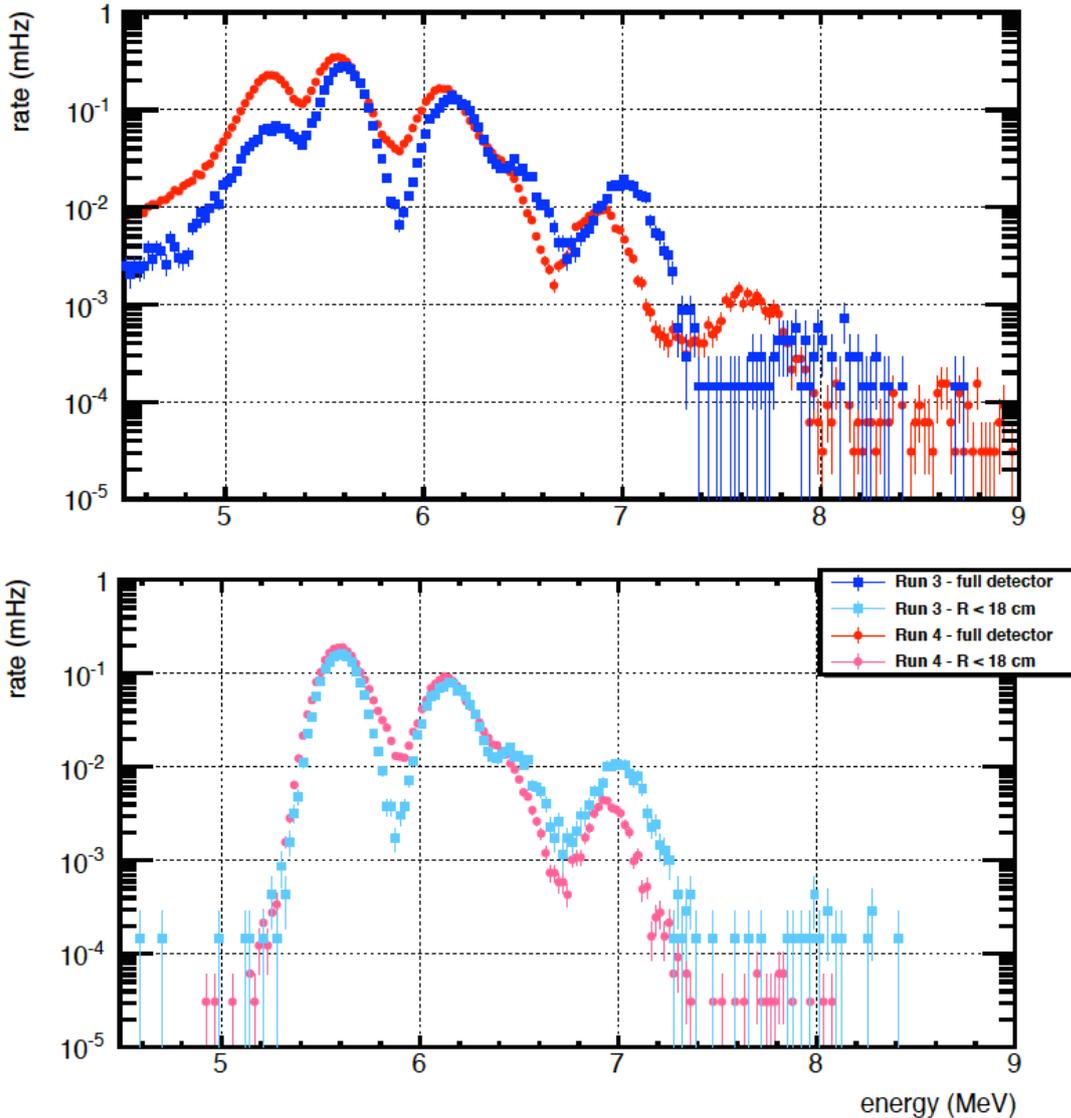
Figure 4.13: α-decay signatures for the full detector (top) and R < 18 cm (bottom) showing data from run 3 and run 4 for comparison. Some peaks shift between the runs, with a bigger effect when including the detector walls, and this is thought to be due to radial field distortions affecting calibration in run 4.

activity A, was used to calculate the number of atoms of each radioisotope present in LUX. This uses the assumption that the population is steady, as would be the case in secular equilibrium. A 'mobility fraction' of ^{210}Bi from the walls to R<18 cm could then be calculated by taking the ratio of the number of ^{210}Bi atoms to its parent ^{210}Pb population. The mobility fraction was also calculated using the number of ^{210}Po atoms, again using the assumption of secular equilibrium. The ^{210}Po activity was calculated from the α-decay study. The measurement from figure 4.10 gave an activity of 3.5 mBq. This was scaled up by 2 for α-particles going into the wall instead of the LXe, ~1.17 to account for the reduction in rate due to S2 loss [117] (i.e. when all of the ionisation becomes trapped in the wall and no S2 is present), and finally by a factor ~1.6 to account for
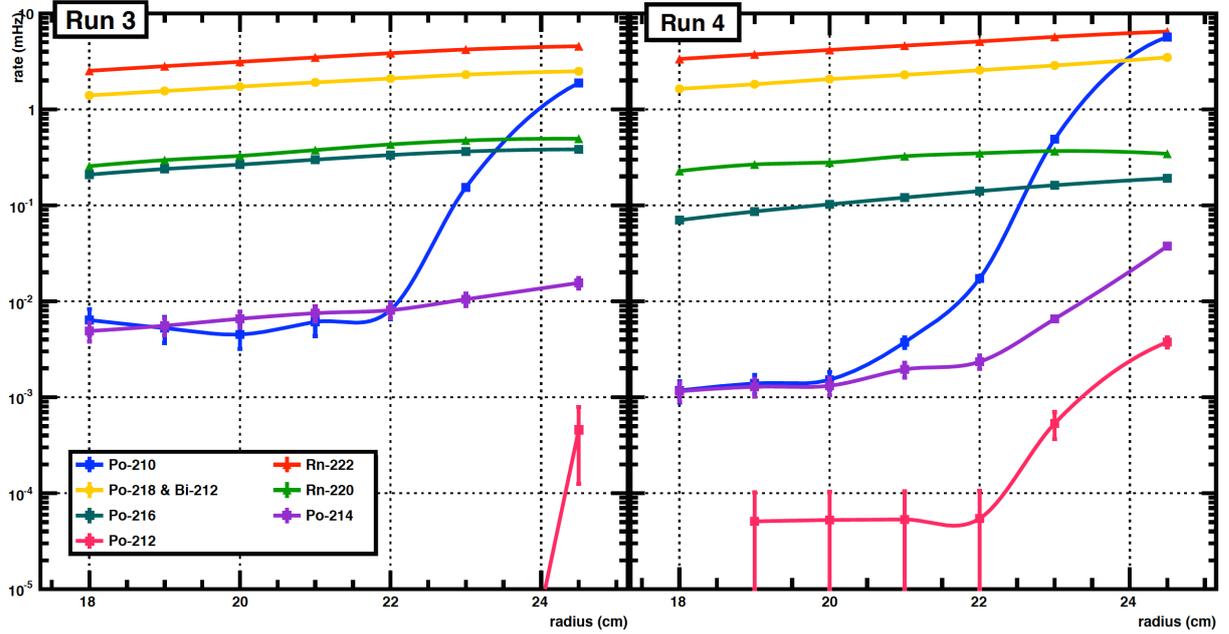
Figure 4.14: α-decay rates contained within the top 30 cm of LUX for different radii. Each point shows the total rate contained within that radius. Run 3 data is shown on the left, and run 4 on the right.

cutting away the bottom 18 cm of the detector to avoid PMT saturation. This resulted in a total activity of 13.2 mBq, in fair agreement with the result shown in [109].

LZ has a requirement to have a maximum of 0.5 mBq/m^2 of ^{210}Pb on its inner surfaces, which have a total surface area of 60 m^2, giving a total activity of 30 mBq. This can be used with the mobility fraction from LUX to estimate the activity of ^{210}Bi within LZ. Table 4.8 shows the total measured activities within LUX of ^{210}Pb, ^{210}Po and ^{210}Bi, the corresponding number of atoms, the mobility fraction of the ^{210}Bi, and the application of these results to LZ. The mobility fraction was calculated using both the ^{210}Pb upper limit and the measured ^{210}Po activity. It has been predicted that a ^{222}Rn contamination level of 2 μBq/kg and the subsequent β-decays of its daughters results in 720 ER counts in the 5,600 tonne-day exposure. Therefore, the ^{210}Bi activities predicted for LZ are also shown in units of μBq/kg in table 4.8, and are converted to an approximate upper limit on ER counts using the ^{222}Rn prediction. For results calculated using the ^{210}Pb activity, secular equilibrium and thus an activity also of 30 mBq on the surfaces is assumed. This assumption is backed up by the similar activities of ^{210}Pb and ^{210}Po (12.3 mBq and 13.2 mBq) on the LUX inner PTFE surfaces.

Similar results are obtained when using ^{210}Pb and ^{210}Po results, which is expected as their activities suggest they are in secular equilibrium. Using the ^{210}Bi α-decay upper limit results in a pessimistic activity of ~12 μBq/kg in LZ, which gives an ER count on the order of 4,000 in LZ's nominal exposure of 5,600 tonne-days. Using the much lower limit obtained with the β-decay results in ~0.1 μBq/kg and ~40 ER counts. This is sub-dominant to the total ER count from both
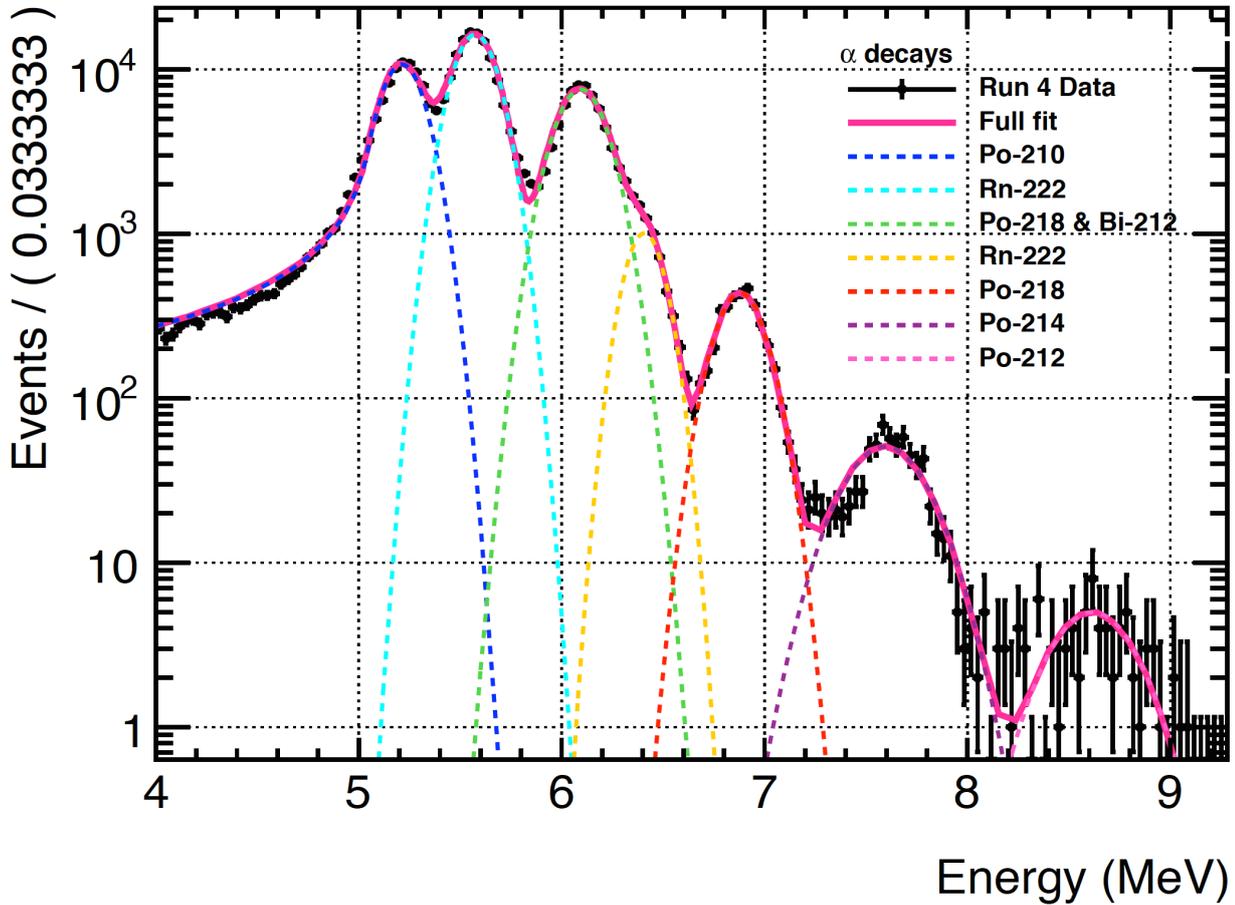
Figure 4.15: Run 4 data for the full detector fitted for seven α-decays using six Gaussians and one crystal ball function.

^{222}Rn and ^{220}Rn of 844, although not at all insignificant.

4.6 Conclusion

It is hard to draw solid conclusions in this study due to the extremely small branching ratio of the ^{210}Bi α-decay. Note that no assumptions are made on the transport efficiency of radon daughters; the activity on the wall from cut and count methods is calculated assuming everything from the wall travels to the fiducial volume before decaying.

For run 3, considering that there are 2 possible ^{210}Bi α-decays in the bulk results in upper limits of <2.4 mBq/kg in the xenon and <470 mBq/m^2 on the PTFE walls of LUX. However, as there are several other α-decays around 4.7 MeV with 100% branching ratios (see table 4.5) means this could be a very conservative upper limit.

Slightly lower limits are achieved by using the best fit to the wall data (0 events) and assuming that <1 ^{210}Bi α-decay has occurred, and allowing the maximum number of ^{210}Bi decays as the

Table 4.8: LUX measurements of ^{210}Pb, ^{210}Po and ^{210}Bi, and their application to LZ for predictions of the ^{210}Bi activity and ER contribution. Mobilities, activities and counts highlighted blue were calculated using the ^{210}Pb activity, whilst pink signifies ^{210}Po was used, using secular equilibrium and thus a ^{210}Po activity also of 30 mBq. ER counts were estimated using the current backgrounds table estimation of 720 ER counts for 2 μBq/kg of ^{222}Rn, from its β-decaying daughters.

		LUX				LZ		
Decay	Region	A (mBq)	$T_{1/2}$	N	Mobility Probability	A (mBq)	A (μBq/kg)	Counts (ER)
^{210}Pb γ	Walls	12.3	22.3 yr	1.24×10^7	-	30	-	-
^{210}Po α	Walls	13.2	138 d	226,922	-	30	-	-
^{210}Bi β	R <18 cm	<0.26	5 d	162	1.3×10^{-5}	0.64	0,113	<41
					7.2×10^{-4}	0.59	0.106	<38
^{210}Bi α	R <18 cm	<28.4	5 d	<17,748	<0.0014	<69.4	<12.40	<4,477
					0.078	<64.6	<11.54	<4,165

number expected per 1 α-decay (1/0.0000014 = 714,285). This provides upper limits of <0.63 mBq/kg in the xenon and <230 mBq/m^2 on the walls. This uses the underlying assumption that we would detect every ^{210}Bi α-decay; to account for decays that go into the walls instead of the xenon this could be doubled to <460 mBq/m^2 on the walls.

For run 4, as the livetime is more than tripled for this analysis, the upper limit within the bulk (as there are no observed events for R< 18 cm) drops to <0.25 mBq/kg whilst treating the data independently. Attempting to fit without forcing signal again favours zero signal, and this gives the lowest limit of all of <0.13 mBq/kg (as there is a larger mass used). However, as small signal could easily hide under the ^{210}Po tail, the cut and count method is a more reliable limit.

Under the assumption the 2 events in run 3 were ^{210}Bi decays, combining the run 3 and run 4 data gives an upper limit of <0.41 mBq/kg in the bulk xenon and <79 mBq/m^2 on the PTFE.

Conclusions can also be drawn on the number of α-decays produced within the fiducial volume and within an energy window. Between 2 MeV - 5.1 MeV, the number of events drops significantly when moving from the walls to a radius of 18 cm (through 4.5 cm of xenon), meaning only 0.03% to 0.16% of α-particles in this energy range are produced in the bulk. The lower limit comes from run 4 where 13 of 41,867 events pass the radial cut, and the upper limit from run 3, where 4 of 2,447 events pass the radial cut.

This study does confirm the strong radial dependence of ^{210}Po decays in LUX due to the deposition of long-lived ^{210}Pb on the detector walls, making it likely that ^{210}Bi is also concentrated there. There is little evidence for mobility of ^{210}Bi within LUX from a study of alpha decays, but due to the rarity of these decays, this can not be considered conclusive.

A significantly lower limit on ^{210}Bi activity within LUX was obtained in a study of the β-decay, where the decay spectra was added to the background model to find the best fit value of 0.26 mBq, within R <18 cm. A search for the 46.5 keV γ-emission of ^{210}Pb was also conducted, resulting in an upper limit of 1.67 μBq/cm^2 on the LUX PTFE. These results have been collated to place upper limits on the amount of ^{210}Bi expected in LZ, using the requirement of a maximum of

30 mBq of ^{210}Pb on the inner surfaces. Using the extremely conservative α-decay upper limit places a upper limit on the activity in the fiducial volume of 12 μBq/kg, which could contribute as many as 4,000 ER counts within the nominal exposure. Using the β-decay upper limit allows a more stringent constraint of 0.1 μBq/kg, leading to about 40 ER counts. This final number is now in the LZ Backgrounds Control Table, see table 7.9 in section 7.2.

5

LUX CALIBRATIONS AND RESULTS

L UX has published several world-leading WIMP-nucleon cross section limits with the initial analysis of its run 3 data of 85.3 livedays, a reanalysis of the same data (with an additional 10 livedays) and its run 4 data of 332 livedays. Furthermore, it has pioneered new and novel techniques for calibrating a liquid xenon TPC down to energies below 1 keV$_{nr}$. Calibration is essential for characterising detector response, and no discovery could be made without doing this thoroughly for both signal and background events across a range of energies. This chapter will detail both the calibration techniques and their results, and WIMP-search analyses and their results.

5.1 LUX Electron Recoil Calibrations

LUX uses a suite of ER calibration sources and techniques to map out the detector response to background events not only in the region of interest but across a wide energy range. These calibrations ensure consistent and correct treatment of WIMP search data, and broaden understanding of LXe physics and detector-specific effects. The next few sections describe each of the LUX ER calibrations and their key results.

5.1.1 Krypton-83m

^{83m}Kr is produced using a Rubidium source:

$$^{83}_{37}\text{Rb} + \text{e}^- \rightarrow {}^{83m}_{36}\text{Kr} \rightarrow {}^{83}_{36}\text{Kr} \tag{5.1}$$

where $t_{1/2}(\text{Rb}) = 86.2$ days. The ^{83m}Kr then undergoes two transitions, emitting a conversion electron of energy of 32.1 keV, with $t_{1/2} = 1.83$ hours, and then one of 9.4 keV, with $t_{1/2} = 0.15$ μs.
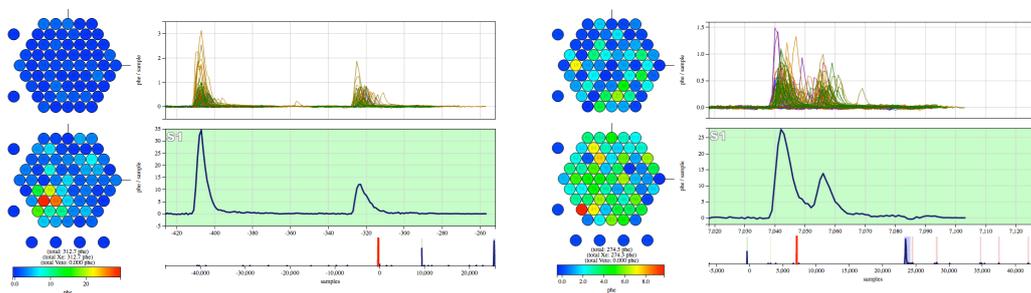
Figure 5.1: Two examples of double S1s from the two decays of ^{83m}Kr in LUX krypton calibration data, demonstrating how they can be both clearly separated when the decay time is long or merged together when it is short.

These two transitions in short succession makes ^{83m}Kr very useful for calibrating LUX at energies down to 9.4 keV.

In LUX, a rubidium source is left to decay in an isolated part of the gas system. When a sufficient amount of ^{83m}Kr has been produced, it can be injected into the TPC. There is sufficient mixing during the 1.83 hour half-life to allow the atoms to spread throughout the active volume for the whole detector to be calibrated for scintillation and ionisation response.

A typical ^{83m}Kr event in LUX will contain two S1s, but only 1 S2. This is because the two decays happen sufficiently close in time that they are in the same position within the resolution of position reconstruction. This is the case both in physical x-y co-ordinates, and in z (equivalent to drift time) so that within smearing by electron diffusion, just one S2 is observed. Two examples of ^{83m}Kr S1s are shown in figure 5.1.

^{83m}Kr injections are performed weekly to determine the electron lifetime and the corrections for photon detection efficiency. These corrections account for the effects of geometric light collection and PMT quantum efficiency. The 9.4 and 32.1 keV depositions can be used to check the stability of the S1 and S2 signals in time.

The electron lifetime is calculated using 60 bins in drift time; in each bin a Gaussian is fit to the ^{83m}Kr S2 signal to determine the mean, and an exponential is fit to the means as a function of z. The characteristic attenuation length is given by:

$$\lambda = \tau v_{drift} \tag{5.2}$$

where v_{drift} is the electron drift velocity and τ is the electron lifetime from the exponential fit.

After this z correction is determined, ^{83m}Kr data can further be utilised to produce x and y corrections for S2s. This is done by an algorithm in the DPF that creates a 25×25 grid on the x-y plane (which corresponds to 2 cm $\times$ 2 cm) of the bottom PMT array and fits a Gaussian to each bin to find the average S2 response. A normalisation function is determined that maps each x,y bin to the centre of the detector at x = y = 0. This can be done regularly using large samples

(~700,000 events) of ^{83m}Kr data to get reliable corrections for datasets taken around the time of the calibration.

For S1 position corrections, a similar method is used, and at least ~400,000 ^{83m}Kr events are needed. The detector is split into voxels. of 25 by 25 by 16 in x,y,z (2 cm × 2 cm × 20 μs), and the normalisation factor to the centre of the detector is determined as before. Unlike the S2 corrections, S1 corrections are not very variable with xenon purity, and so can be done monthly with larger, high statistics ^{83m}Kr injections.

5.1.2 Tritiated Methane

5.1.2.1 Injection and Removal

In order to calibrate LUX's response to electron recoils down to the WIMP search threshold, the source used is tritium, which decays by β-decay:

$$^3_1\text{T} \rightarrow ^3_2\text{He} + e^- + \bar{\nu}_e \tag{5.3}$$

with a half-life of 12.32 ± 0.02 years. 18.6 keV is released in the process, making it an ideal low energy calibration source. LUX uses tritiated methane, CH_3T, which can be injected into the detector. Two sources, with total activities 3 Bq and 200 Bq are contained in 2.25 l stainless steel bottles, mixed with 2 atmospheres of purified xenon. A small amount of the activity can be extracted by allowing the carrier xenon gas to expand into an expansion volume, with the extracted activity proportional to the size of the volume. A methane purifier ensures only CH_3T, CH_4 and noble gases enter the TPC. The CH_3T diffuses into the liquid xenon and becomes distributed uniformly throughout within a few minutes. Methane is removed with a one-pass efficiency of 97% by the getter, and disappears exponentially with a time constant of 5.9 ± 0.07 hours [118]. This fast removal ensues normal WIMP search operation can continue shortly after a calibration.

5.1.2.2 β Spectrum

Data is selected with similar cuts to WIMP search (see section 5.3), requiring single scatters (1 S1, 1 S2), and only S2s above 165 phd are selected to ensure accurate position reconstruction. A fiducial volume is used with drift times from 38 to 305 μs (8.5 to 48.6 cm) and a radius of 20 cm. Events where the end of the S2 is truncated by the end of the waveform are also excluded. The tritium calibration also provides a reliable test of the energy reconstruction in LUX, as the β energy spectrum is well known. The energy reconstruction is performed as in section 2.3.3 using the ER energy:

$$E_R = W \cdot (n_\gamma + n_e) = W \cdot \left(\frac{S1}{g_1} + \frac{S2}{g_2}\right) \tag{5.4}$$

g_1 and g_2 were allowed to float and the data fitted to a tritium spectrum, resulting in best fit values of $g_1 = 0.115 \pm 0.005$ phd/photon and $g_2 = 12.1 \pm 0.9$ phd/electron. The electron extraction
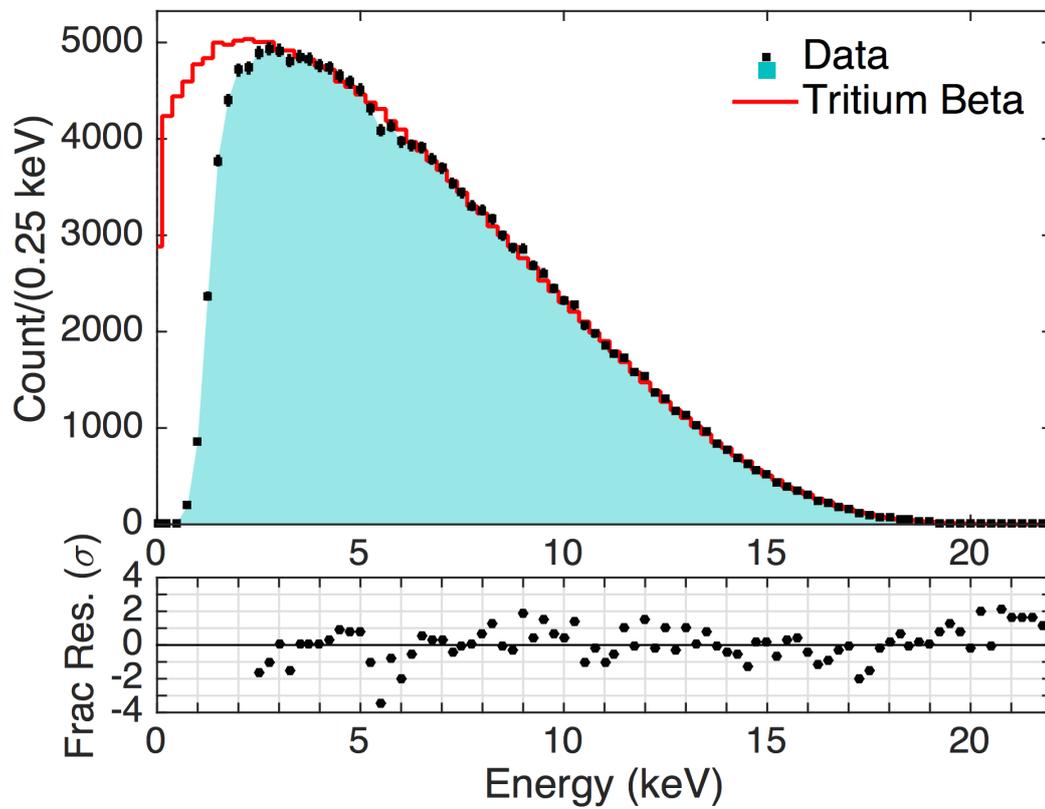
Figure 5.2: Top: The tritium energy spectrum measured by LUX is shown in black, and the tritium β spectrum convolved with detector resolution ($\frac{\sigma_E}{W} = \sqrt{\sigma^2(n_\gamma) + \sigma^2(n_e)}$) is shown in red. Bottom: Bin by bin residuals between data and theory in units of σ [118].

efficiency, defined as the fraction of electrons that are successfully extracted from liquid to gas, can also be measured. Note that this quantity just considers electrons that reach the liquid-gas interface after losses due to electron lifetime/purity. From the tritium calibration, the extraction efficiency is measured as $50.9\% \pm 3.8\%$ [118]. The resulting spectrum is shown in figure 5.2. Detector resolution applied to the tritium spectrum was described using:

$$\sigma_E = \sqrt{\sigma(n_\gamma)^2 + \sigma(n_e)^2} \tag{5.5}$$

with $\sigma(n_\gamma)$ and $\sigma(n_e)$ representing detector resolution for photon and electron counting, and was normalised to data.

5.1.2.3 L_y and Q_y

Tritium data can also be used as a measure of the microscopic processes taking place within the xenon during an electron recoil event. L_y and Q_y were obtained by finding the mean S1 and S2 areas and dividing them by the combined energy E_R in each energy bin of the β spectrum. Data using two field strengths, 105 V/cm and 180 V/cm, was used to check field dependence. Figure 5.3

shows the results alongside predictions from NEST (see section 2.6.5), where the anti-correlation between L_y and Q_y is clear.
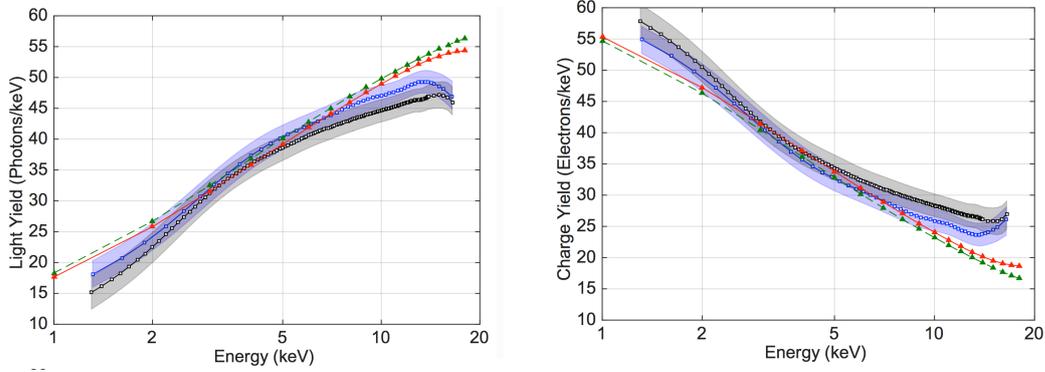
(a) Light yield L_y obtained from tritium ER data.   (b) Charge yield Q_y obtained from tritium ER data.

Figure 5.3: L_y and Q_y as determined with tritium calibration data. Black squares were data taken at 180 V/cm and blue squares at 105 V/cm. Shaded bands show the 1σ systematics due to g_1 and g_2; statistical uncertainties are negligible in comparison. NEST curves are shown in red and green at 180 and 105 V/cm respectively, with triangles placed every 1 keV [118].

Furthermore, the recombination fraction (the fraction of ionised electrons that recombine with a Xe atom) can be calculated using:

$$r = \frac{(n_\gamma/n_e) - \alpha}{(n_\gamma + n_e) + 1} \tag{5.6}$$

where α is the initial ratio of exitons to ions, N_{ex}/N_i. The result is shown on the right of figure
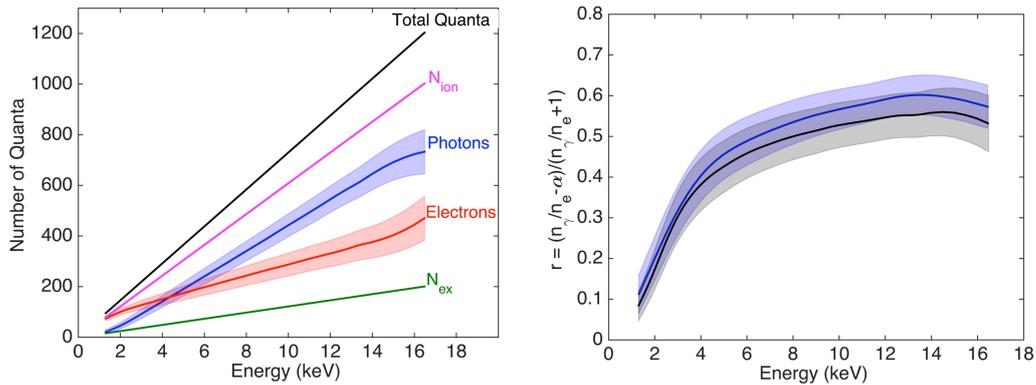
(a) Number of quanta produced in LUX at 190 V/cm as a function of energy.   (b) Recombination fraction as a function of energy at 180 V/cm (black) and 105 V/cm (blue)

Figure 5.4: Number of quanta and recombination fraction as functions of energy. Here, α is taken to be 0.2. Bands are systematic errors on g_1 and g_2. [118].

5.4. r was found to not differ significantly in this energy range between 105 V/cm and 180 V/cm, but recombination is expected to be very small at the lowest energies, approaching zero below 0.4 keV.

5.1.2.4 ER Band

The electron recoil band is a distribution usually presented in $\log_{10}$(S2/S1) vs S1 space, described using a mean and a thickness. It shows a rise at decreasing values of S1, reflecting how the charge and light yields change more rapidly below 6 keV. It forms the basis of discrimination between ER and NR; the NR band lies lower down on the y axis, due to there being less energy channeled into ionisation for NRs than ERs. The ER band characterisation is very important for WIMP search because of leakage. Leaked events look like nuclear recoils and so can be candidate WIMP events. Thus, the leakage fraction, f, defined as the fraction of ER events below the Gaussian mean of the NR band, must be determined. The recoil discrimination efficiency is defined as $(1-f)$ and is found to be 99.81% $\pm$ 0.02% (stat) $\pm$0.1% (sys) for 1<S1<50 phd. The systematic error is the uncertainty in the NR band mean and the effect of field non-uniformity. The width of the band can be parameterised with three components: uncertainties on photon and electron counting σ_{n_γ} and σ_{n_e} and the fluctuations present in recombination $\sigma(R)$. For the S1 and S2, binomial fluctuations are [119]:

$$\sigma^2_{S1,\text{bino}} = (1-g_1)g_1 n_\gamma \tag{5.7}$$

$$\sigma^2_{S2,\text{bino}} = (1-\epsilon)\epsilon(\kappa n_e)SE^2 \tag{5.8}$$

where ϵ is the electron extraction efficiency from liquid to gas, κ is the probability of an electron to not be captured by an impurity, and SE is the number of photoelectrons produced in the PMTs by the electroluminescence of a single electron. Therefore, uncertainties on photon and electron counting can be described by [119]:

$$\sigma(n_\gamma) \propto \sqrt{(1-g_1)g_1 n_\gamma} \tag{5.9}$$

$$\sigma(n_e) \propto \sqrt{(1-\epsilon)\epsilon n_e} \tag{5.10}$$

Furthermore, their variance due to PMT resolution must be included, but this is measured with single photoelectron (LED) data. Therefore, these can be subtracted from the band width and $\sigma(R)$ estimated.

Figure 5.5(b) shows the values of $\sigma(n_\gamma)$, $\sigma(n_e)$ and $\sigma(R)$ increasing with energy. $\sigma(n_\gamma)$ is dominant in LUX at 180 V/cm, which is not unexpected due to the low light collection efficiency; $g_1 = 0.115$ compared to $g_2 = 12.1$. In the WIMP search energy range ($\sim 2-6$ keV), $\sigma(n_e)$ and $\sigma(R)$ are comparable. $\sigma(R)$ appears to grow linearly with energy; from figure 5.4(a) it can be seen that this is because the number of ions available for recombination is increasing linearly, therefore we can describe the fluctuations as:

$$\sigma(R) = (0.067 \pm 0.005) \times N_{ion} \tag{5.11}$$

for the range 2 - 16 keV.

To obtain a description of the width of the ER band, Gaussian fits are performed in 16 S1 bins after subtracting the centroid and dividing by the Gaussian width. Gaussian fits are a good
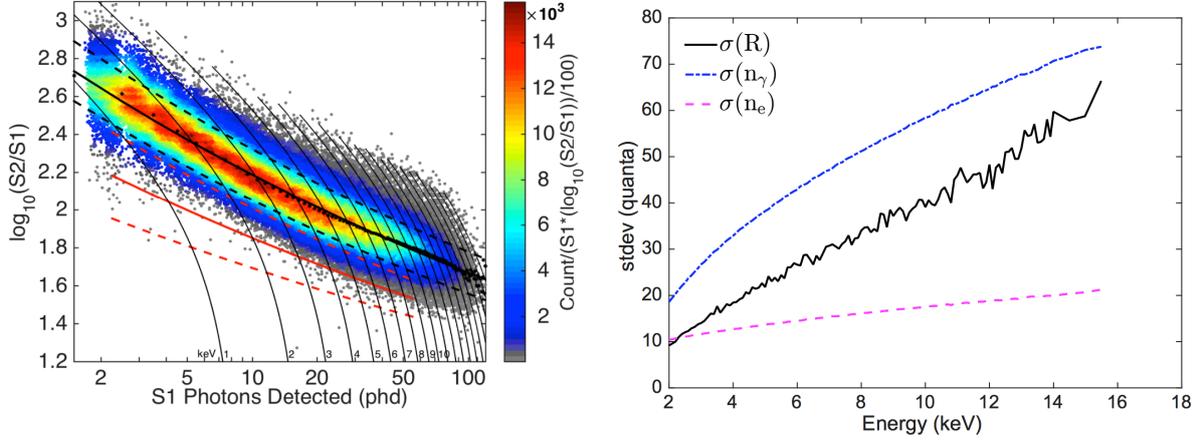
(a) The ER band. Shown are contours of ER energy from 1 to 20 keV, Gaussian means in S1 bins (filled dots), an empirical power law mean (solid black line) and 10% and 90% contours (dashed black line). The NR band from DD data (section 5.2.4) is also shown in red.

(b) Recombination fluctuations in LUX (black), photon counting detector resolution (dot-dash blue) and electron counting detector resolution (dashed magenta), all shown against energy in keV.

Figure 5.5: Left: ER recoil band shown in the discrimination variable log10(S2/S1) vs S1. Right: detector resolutions and fluctuations shown against energy. Data consisted of 170,000 tritium events, with the LUX nominal electric field of 180 V/cm [118].

description of the data from $\mu - 3\sigma < \mu < \mu + 2\sigma$, but either side of that there are non-Gaussian tails, although these are recreated in simulation and originate from photon counting (Poissonian fluctuations). The mean of the Gaussian fit in each bin can be seen as solid black circles in figure 5.5(a).

5.1.3 Cs-137

An external caesium source is used for calibrations for higher energy response and stability checks, using the β-decay of ^{137}Cs and the following decay of the metastable state ^{137m}Ba:

$$^{137}\text{Cs} \rightarrow {}^{137m}\text{Ba} + e^-(512 \text{ keV}) \tag{5.12}$$

$$^{137m}\text{Ba} \rightarrow {}^{137}\text{Ba} + \gamma(662 \text{ keV}) \tag{5.13}$$

The ^{137}Cs source is housed in a source holder that contains a collimater hole and lowered into the LUX water tank via external source tubes, giving the data a position dependence useful for testing position reconstruction. Furthermore, the ^{137}Cs source provides ER data from the backscatter peak at 150 keV to the peak from the emitted photon at 662 keV. The backscatter peak arises from 662 keV photons backscattering of the back of the source holder and passing back through the collimator hole.

(a) Monoenergetic gamma sources in $\log_{10}(S2)$ against $\log_{10}(S1)$.

(b) Doke plot using ^{83m}Kr, ^{137}Cs calibration sources as well as xenon lines and a ^{214}Bi line.
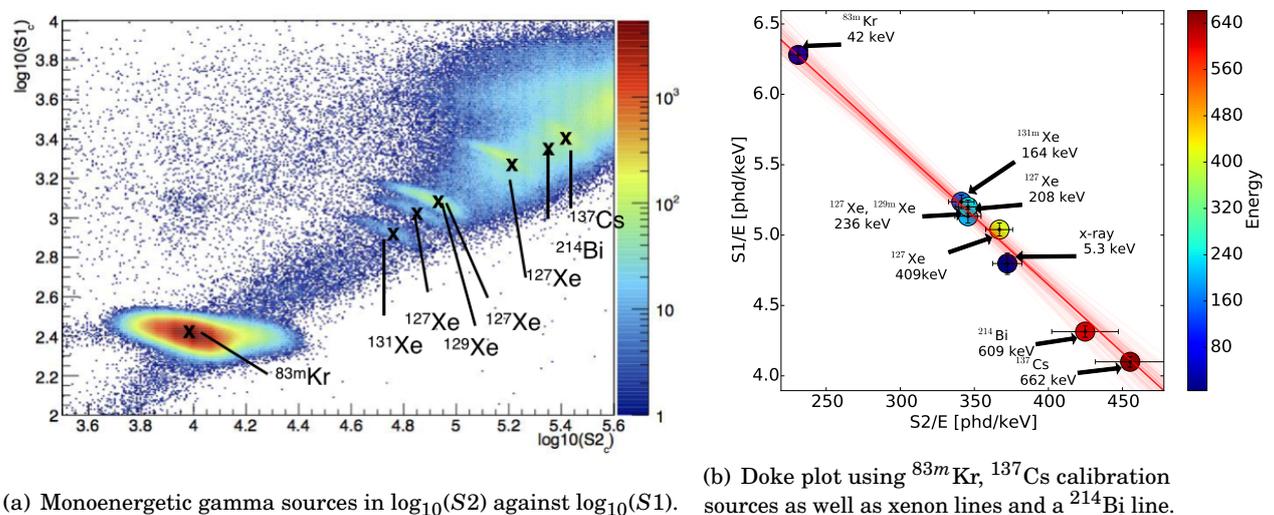
Figure 5.6: The mean of populations of various monoenergetic γ-sources (left) are used to measure g_1 and g_2 in a Doke plot (right) by plotting S1/E and S2/E [119].

5.1.4 γ-ray Calibrations

Combining the ^{137}Cs and ^{83m}Kr calibrations with internal Xe γ-rays can allow an accurate determination of g_1 and g_2 using a Doke plot [78]. Equation 5.4 can be rewritten as:

$$\frac{S1}{E} = \frac{g_1}{W} - \left(\frac{S2}{E}\right)\left(\frac{g_1}{g_2}\right) \tag{5.14}$$

This is possible as each source is monoenergetic and forms a population in S1 and S2, see figure 5.6(a). Then, a plot of $S1/E$ against $S2/E$ will have a gradient of $-g_1/g_2$ and an intercept of g_1/W, as shown in figure 5.6(b). When applied to run 3 data, this resulted in measurements of $g_1 = 0.117 \pm 0.003$ phd/γ and $g_2 = 12.1 \pm 0.8$ phd/e which are consistent with the results from the Tritium calibration mentioned previously.

5.2 Nuclear Recoil Calibrations

NR calibrations are essential for characterising the expected WIMP signal. Multiple techniques have been important in driving LUX sensitivity to higher levels. The older neutron calibrations used to map out the NR band using AmBe and ^{252}Cf sources used spectral matching techniques with a Monte Carlo to compare to data. This worked well, but the new, novel D-D technique gives direct, in-situ measurements, which are far more desirable. The techniques and their results are described in the following sections, as well as the presentation of a simulation method for measuring activities of neutron calibration sources.

5.2.1 AmBe Neutrons

LUX NR calibrations prior to and for the first LUX results in 2013 were done using an AmBe source. AmBe produces neutrons by an (α,n) reaction:

$$^{241}_{95}\text{Am} \rightarrow ^{4}_{2}\alpha + ^{237}_{93}\text{Np} \tag{5.15}$$

$$^{4}_{2}\alpha + ^{9}_{4}\text{Be} \rightarrow ^{12}_{6}\text{C} + ^{1}_{0}n \tag{5.16}$$

Neutron energies range from 0 - 11 MeV, see figure 5.9 for a simulated AmBe spectrum. NR band calibration data points from AmBe neutrons can be seen alongside ^{252}Cf data in the bottom panel of figure 5.14.

5.2.2 Measuring the AmBe Calibration Source Activity

To constrain absolute detector efficiency to nuclear recoils, it is useful to know a calibration source activity accurately. However, such activities, including for the LUX AmBe source, are often unknown. A feasibility study was performed on the activity measurement of the AmBe neutron source using a high purity germanium (HPGe) γ-ray detector, by simulating two different methods. The aim was to determine whether or not the measurement could be performed underground without the associated difficulties and costs of transporting a radioactive source back from underground and to a measurement facility. Such a method is applicable for all neutron sources with emission energies of approximately 0–10 MeV, such as the AmBe and ^{252}Cf sources used commonly in dark matter experiments in addition to D-D generators, and especially where D-D generators are not used.

The first method involved placing the source inside a plastic container filled with water with the intention of thermalising the neutrons. Some neutrons will be captured by hydrogen in the water and release a 2.2 MeV γ-ray. This in principle can be detected by the HPGe detector and the detected number of γ-rays used to infer the activity of the source.

The second method followed a similar procedure, but instead used a MnSO$_4$ bath, followed by a transfer of the solution containing neutron-activated manganese to a Marinelli beaker surrounding a HPGe detector. γ-rays from the decay of ^{56}Mn would be used to estimate the source activity. The method involves more preparation, but increases detection efficiency considerably.

The GEANT4 toolkit was used in this study. An accurately modelled HPGe was shielded by a copper box of $600 \times 600 \times 600$ mm and further by a lead box of $800 \times 800 \times 800$ mm. For method 1, the roof of this shielding was removed in order to leave 400 mm of air between the edge of the shielding and the volume of air containing the detector (labelled as inbox in figure 5.7(a)). The detector itself is placed so that its end-cap is 345 mm from the top of the shielding. A plastic cylinder of radius 245 mm, height 790 mm and 5 mm thickness was placed on top and filled with water. The measurements were taken from a standard plastic container present in the Davis Lab at SURF. The neutron source was placed at the centre of this cylinder, see figure 5.7(a). In
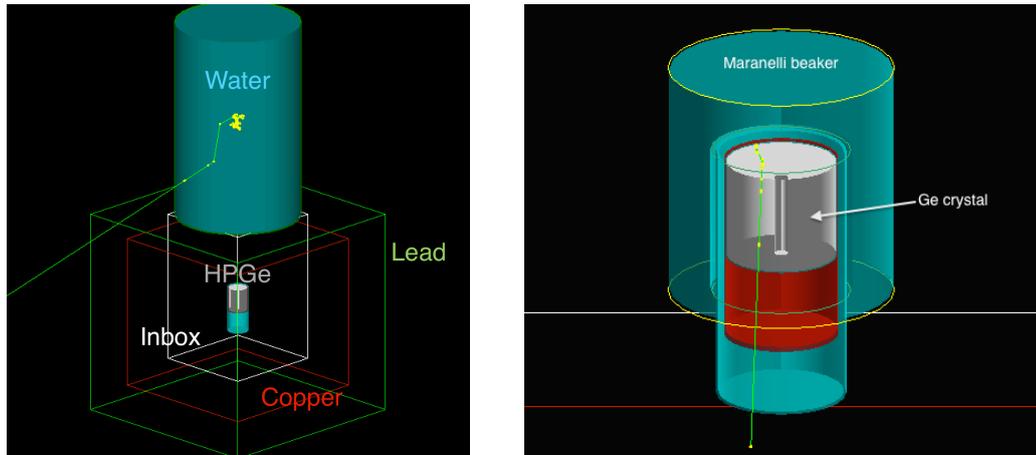
(a) Method 1 - source is located at the centre of the water, which sits on the detector shielding.

(b) Method 2 - a solution activated by the source is contained by the Marinelli beaker

Figure 5.7: GEANT4 visualisation simulation geometries; the shielding and plastic containers have been made transparent. Left shows method 1, right shows method 2.

this simulation, the HPGe geometry is that of a p-type Ortec 2 kg crystal. This would need to be modified depending on which HPGe detector is used at SURF, should these measurements be performed. For method 2, the same HPGe detector was surrounded by a Marinelli beaker of dimensions taken from a real beaker used for these measurements (165 mm height, 157 mm diameter, with a well 104 mm deep and of diameter 96 mm). Marinelli beakers are designed for γ-spectroscopy; made of polypropylene, they contain a well for the detector to be placed inside, allowing for greater sample-detector surface area contact, see figure 5.7(b). In the simulation, the beaker was filled with an $MnSO_4$ solution of concentration 1.25 g/cm^3, and the γ-source was confined to the inner Marinelli beaker volume, meaning the position of each γ-ray fired is generated randomly from within this volume. The source was modelled using a histogram of the energies and intensities of γ-rays emitted by ^{56}Mn, seen later in table 5.2.

The MAEVE HPGe detector used for both methods was modeled using dimensions from ORTEC, with a germanium crystal of diameter 77.6 mm and length 69.8 mm, see figure 5.8. The detector was made sensitive in order to record deposited energy of any γ-rays entering the germanium region. The physics list QGSP_BERT_HP was used, with the addition of neutron thermal scattering, which was added using a modified version of G4HadronElasticPhyiscsHP with a G4NeutronHPThermalScattering process for neutrons below 4 eV. QGSP_BERT includes Geant4 Bertini cascades for primary protons and neutrons at energies below 10 GeV, and the addition of the NeutronHP high precision neutron package transports neutrons below 20 MeV down to thermal energies.

In order to investigate the safety of such a geometry in real life, the plastic of the container was also made sensitive. It recorded when neutrons, γ-rays or other particles left the plastic container and entered the world volume. Their kinetic energy and any energy deposited in the
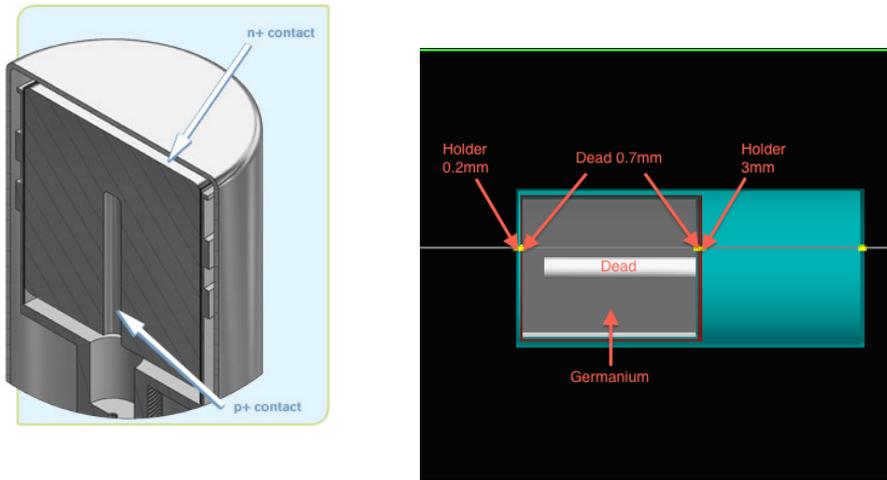
Figure 5.8: Left: A standard coaxial HPGe. Right: Close up Geant4 visualisation of the detector within the simulation.

plastic itself were also recorded. This is important as neutrons escaping could be dangerous as well as possibly damaging to the detector itself - activation of some materials could occur through neutron capture and make the detector unusable. For method 1, the information was also used to analyse how many 2.2 MeV γ-rays were escaping the container in total and to cross-check this with how many the HPGe detector was measuring.

Figure 5.9 is the AmBe neutron spectrum used in the simulation, as measured in the International Organization for Standardization catalogue of reference neutron radiations part 1 (ISO 8529-1). The simulation sampled from a histogram of this spectrum and neutrons were emitted isotropically from the centre of the container.

5.2.2.1 Method 1 Results

It was expected most of the neutrons would be captured by a hydrogen atom within the water.

$$^{1}H + n \rightarrow ^{2}H + \gamma \tag{5.17}$$

where the photon is of energy 2.2 MeV. This process only occurs at low neutron energies and so most must first be thermalised through proton scattering in the water. It is important to try to thermalise most of the neutrons coming from the source. Neutrons can activate parts of the detector and possibly render it unusable.

A solid angle calculation can be used to make a rough prediction of how many 2.2 MeV γ-rays should be hitting the sensitive HPGe detector, treating the container as a isotropic point source:

$$\Phi = \frac{\phi}{4\pi d^2} A \tag{5.18}$$
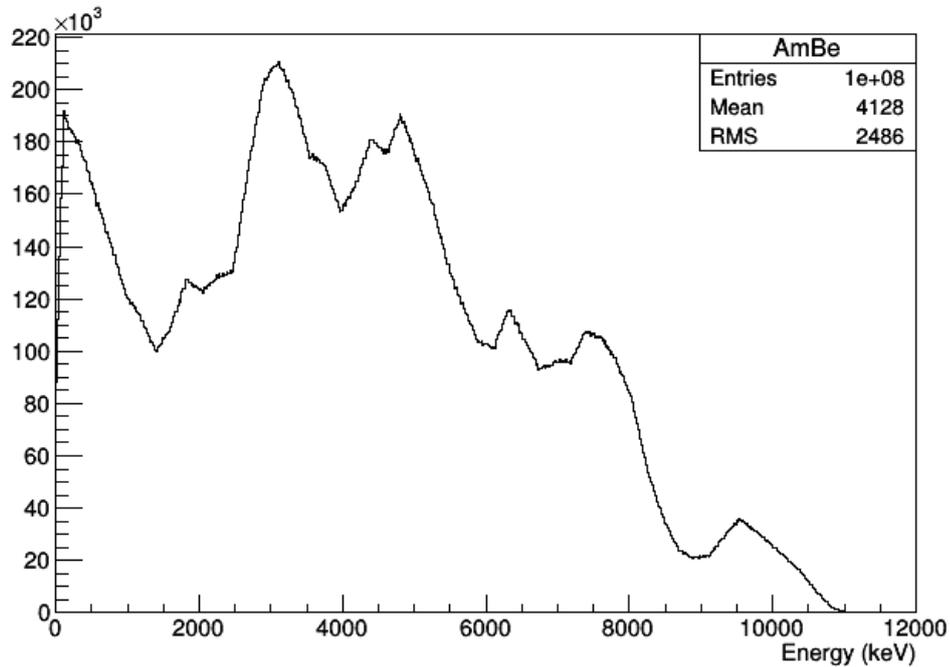
Figure 5.9: AmBe neutron source energy spectrum of 1,000,000 incident neutrons, following ISO 8529-1 (2001)

where ϕ is the initial photon flux, d is the distance from the centre of the container to the HPGe detector, which is 744 mm, and A is the area of the Ge crystal presented to the container, a circle of radius 36.3 mm with a hole of radius 5.15 mm. This results in a total area of about 4,140 mm^2.

In initial tests of 1,000,000 neutrons, 249,152 2.2 MeV γ-rays left the container. The calculation in equation 5.18 leads to an expected 149 γ-rays hitting the face of the Germanium volume, but in reality only 21 were recorded. Not all of these will deposit their full energy in the detector, and detector efficiency will also account for a lower count. This means for an integrated flux of 100,000,000 neutrons for the high statistics run, it can be expected that around 25,000,000 2.2 MeV γ-rays will leave the container, around 15,000 2.2 MeV γ-rays will hit the Germanium, but only around 2,000 will be seen at the 2.2 MeV line in the Germanium sensitive detector.

Table 5.1 lists the counts recorded by the sensitive detectors, as well as the average kinetic energy of the particle. For the container, these were only particles whose next track was in the world volume which ensures they have escaped. The approximate rates of events are calculated assuming an estimated source activity of 100 neutrons/s.

For a run of 100,000,000 neutrons, estimated to be around 11.5 days of data taking, a total of 74,146,030 γ-rays left the plastic container. 24,950,862 of these were at the 2.2 MeV line. The HPGe detector recorded a total of 50,788 (4,388 γ/day) hits from γ-rays, see figure 5.10. Of these, only 1,842 (159 γ/day) were detected by the HGe detector. The expectation from the calculations would be 14,850—but as stated previously this was a rough calculation treating the

Table 5.1: Counts and approximate rates of particles leaving the container and hitting the HPGe detector. Rates are an approximation using a source activity of 100 Bq.

	Particle	Number	Avg. Energy
Container	Neutron	12,342,490 (741 n/min)	2.152 MeV
	γ	74,146,030 (4,449 γ/min)	1.181 MeV
	γ at 2.2 MeV	24,950,862	2.224 MeV
HPGe	γ	50,788 (4,388 γ/day)	0.660 MeV
	γ at 2.2 MeV	1,842 (159 γ/day)	2.224 MeV
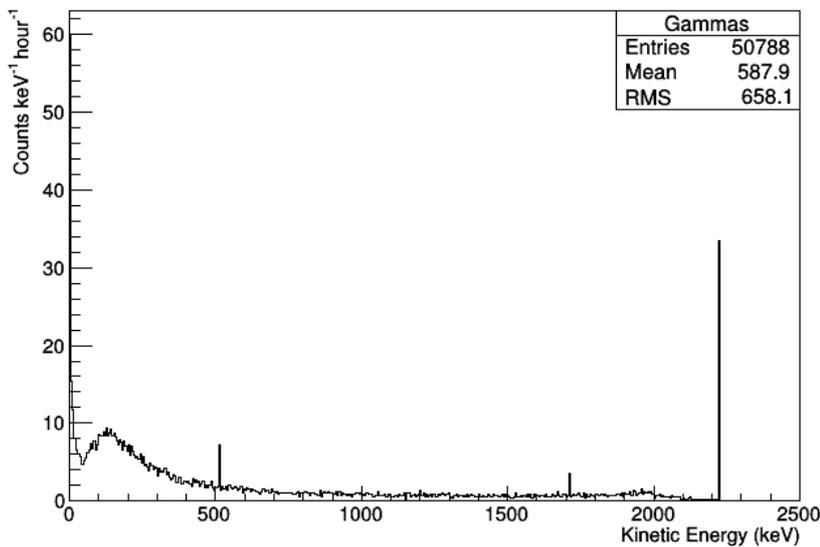
Figure 5.10: Energy deposited by γ-rays in the HPGe detector, from 0 to 2.5 MeV. The 2.2 MeV line is clearly visible.

container as an isotropic point source and assuming perfect efficiency. The comparison to the preliminary 1,000,000 neutron tests predicted around 2,000 hits, reasonably matching the 1,842 seen. The average energy of the γ-rays in the detector is about 0.5 MeV less than those exiting the container—suggesting many may lose energy scattering on the way to the detector. Furthermore, the percentage of the γ flux that is at the 2.2 MeV line is much greater than in the detector - 34% compared to 3.6%. This also helps to explain why the number is approximately 10% of the counts from the solid angle estimate - most of the γ-rays are not depositing the full 2.2 MeV in the detector. 1,842 total counts from 100,000,000 neutrons is approximately 160 counts a day - giving a Poissonian error of $\sqrt{N} \approx 13$. This hypothetically would allow the neutron source activity to be known to $\frac{\sqrt{N}}{N} \approx 8\%$. The data was smeared using a Gaussian function, selecting a random

133

number from the function defined by:

$$f(E_{dep}) = e^{-\frac{(E_{dep})^2}{2\sigma^2}} \tag{5.19}$$

where E_{dep} was the energy deposited in the detector, and $\sigma = 0.002 \times E_{dep} + 0.7464$, which was designed to mimic the resolution of a similar HPGe detector at the Boulby Underground Laboratory Germanium Suite (BUGS). This number was then added to the original E_{dep} and was used to fill a new histogram. This was to give an idea of the level of resolution that may be obtained with this experiment. Increasing the window to accommodate the detector resolution would lead to a slightly larger count of 1,854 due to a few events sitting either side of the original narrow line.

For the high statistics run, 12,342,490 neutrons (12.3%) left the container, with a mean energy of 2.152 MeV, RMS of 2.693 MeV and a cut off around 11 MeV, which mimics the AmBe spectrum (see figure 5.9). This is a rate of 40,800 n/hour. A large proportion (around 5,000,000, 18,000 n/hour) of these neutrons were of energies less than 0.2 eV and so are thermalised.

A solution to the high flux is to use a larger container. A test was run with the radius of the container increased by 100 mm to 345 mm. In this case, 40,061 neutrons left, a rate of 14,400 n/hour, with 15,640 (5,600 n/hour) of these below 0.2 eV. The 100 mm increase in radius leads to a volume increase of 0.144m^3 on 0.143m^3, doubling the volume. This has cut the number of neutrons in comparison to the previous test run by 67%, suggesting if it was considered necessary, the outward flux could be reduced significantly by using larger volumes of water. The larger volume slightly increased the sensitive detector hits from 487 to 561, with an increase of 1 from 21 to 22 at the 2.2 MeV line. A larger increase of the container to a radius of 845 mm, and a height of 1590 mm (just as a proof of principle) resulted in only 214 neutrons exiting out of 1,000,000 - 0.02% of the initial flux - a rate of 78 an hour. Considering this is the rate in all directions this leaves only a tiny flux directed towards the detector. Additional tests of different container sizes demonstrated a power law fall off with the volume of water used. This could allow a volume to be chosen that maximises the safety of the experiment, although the efficiency of the HPGe to detect γ-rays would have to be re-assessed with a container of different geometry.

5.2.2.2 Method 2 Results

^{56}Mn decays by beta decay to ^{56}Fe. The relevant processes of activation and decay are as follows,

$$^{55}_{25}\text{Mn} + ^1_0 n \rightarrow ^{56}_{25}\text{Mn} \tag{5.20}$$

$$^{56}_{25}\text{Mn} \rightarrow ^{56}_{26}\text{Fe} + e^- + \bar{\nu}_e + \gamma \tag{5.21}$$

where the γ-ray is emitted as the excited ^{56}Fe drops to ground state. It should be noted that sometimes the γ-rays are emitted in coincidence - one nuclei can emit two, depending on the nuclear excitement. Table 5.2 shows the energies and intensities of the ^{56}Mn γ-rays.

Table 5.2: ^{56}Mn γ-ray energies and intensities.

Energy (keV)	Intensity	Energy (keV)	Intensity
846.7638	98.85%	2523.06	1.018%
1037.8333	0.04%	2657.56	0.645%
1238.2736	0.04%	2959.92	0.306%
1810.8333	26.9%	3369.81	0.168%
2113.092	14.2%		

The number of ^{56}Mn present at a time t is given by the following equation [120]:

$$N(t) = \frac{R}{\lambda}(1 - e^{-\lambda t}) \qquad (5.22)$$

where R is the production rate and λ is the decay constant, in this case:

$$\lambda = \frac{\ln 2}{t_{1/2}} = \frac{\ln 2}{2.5789 h} = 0.2688/\text{hr} \qquad (5.23)$$

If a production rate can be estimated, this will give an idea of how long the source should be left in solution to maximise the number of ^{56}Mn. Therefore, before simulating the Marinelli beaker, a GEANT4 simulation was performed to measure the activation of a manganese sulphate solution by an AmBe source.

The source was placed at the centre of a container of a MnSO$_4$ solution with a density of 1.25 gcm^{-3}, as described in [121]. Ideally, the solution would be circulated with a pump, but in the simulation it is taken to be stationary and uniform. Concentrations using the solubility of MnSO$_4$ at 5°C, 0.52 g/cm^3, and at 7°C, 0.70 g/cm^3 were also tested, giving densities of 1.52 and 1.70 g/cm^3 assuming no volume changes of the water. The container was the same cylinder of radius 0.645 m and height 1.59 m used in the first method. The outward flux of neutrons from the surface of the container was around 0.1% of the production rate by the source, meaning almost all of the produced neutrons are thermalised within the water.

For the purpose of further analysis, an estimated neutron source rate of 100 neutrons/s is again assumed.

Table 5.3: Neutron capture probabilities and ^{56}Mn production rate from simulation.

Concentration (gcm^{-3})	^{56}Mn Creation Probability	H Capture Probability	Escape Probability	Approx. Production Rate (^{56}Mn hr^{-1})
0.25	44.8 %	55.1%	0.082%	47,472
0.52	61.0%	39.0%	0.048%	64,813
0.70	66.8%	33.2%	0.010%	70,988

As expected almost all the neutrons are captured on Mn or H. The amount of neutrons escaping drops as the concentration increases because the neutron cross section for manganese is 13.3 barns, compared to 0.33 barns for hydrogen, making it more likely a neutron is captured on manganese than hydrogen. Thus increasing the amount of manganese means more neutrons overall are captured.

135

The information in table 5.3 can be used along with equation 5.22 to produce figure 5.11. The solution starts to reach saturation after around 16 hours.
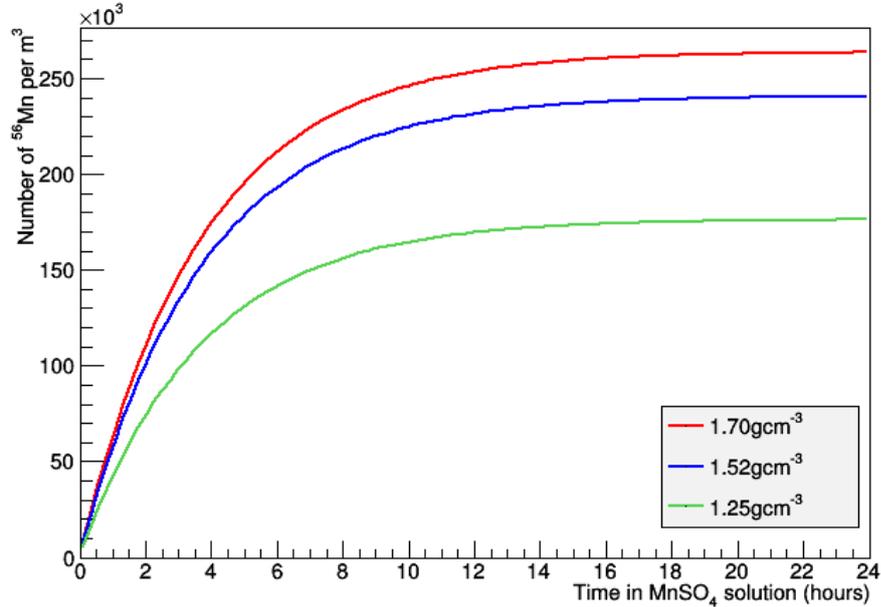
Figure 5.11: The number of ^{56}Mn atoms per cubic centimeter as time in the bath progresses for three different concentrations of the manganese sulphate solution. The curves level off as the solution reaches equilibrium between production and decay.

For the purpose of statistics, 10 million decays were simulated. The actual number of individual decays was increased to account for coincident decays. Figure 5.12 shows the spectrum recorded by the detector. At the 856.7 keV line, using the truth value for emitted gammas of this energy, the detection efficiency, ϵ_c, was 1.79%.

Table 5.4: Predicted numbers of ^{56}Mn in the bath and beaker, and the corresponding counts in the detector

Concentration (g/cm^3)	^{56}Mn after 20h atoms per m^3	^{56}Mn in Marinelli Beaker	Counts at 846.8 keV
0.25	175,790	1,717	29
0.52	240,004	2,345	37
0.70	262,870	2,568	47

The sample of activated solution should be allowed to decay away fully. Table 5.4 was calculated using equation 5.22, assuming the source is left in the bath for 20 hours. The counts at 846.8 keV were estimated from an average of simulation results with the appropriate number of decays. One day appears to be sufficient for all the ^{56}Mn to decay, see figure 5.13(a), which results in a detector spectrum like the one shown in figure 5.13(b).

The neutron source rate can be estimated with the formula [120]:

$$R = N_T \frac{\lambda}{1 - e^{-\lambda T}} \frac{1}{\epsilon_{Mn}} \qquad (5.24)$$

136
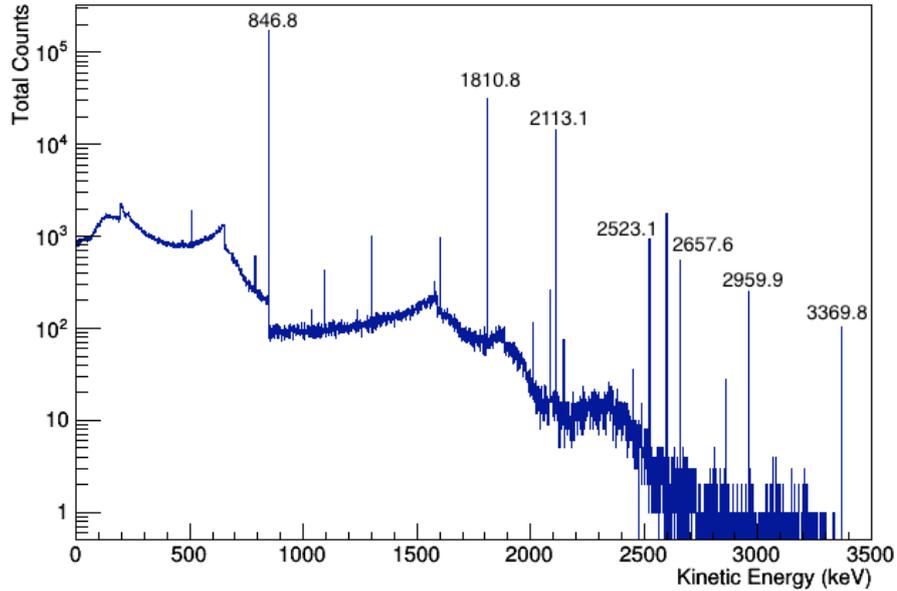
Figure 5.12: γ-ray spectrum measured by the detector. Several lines from the decay of ^{56}Mn can be identified and their energies are labeled.
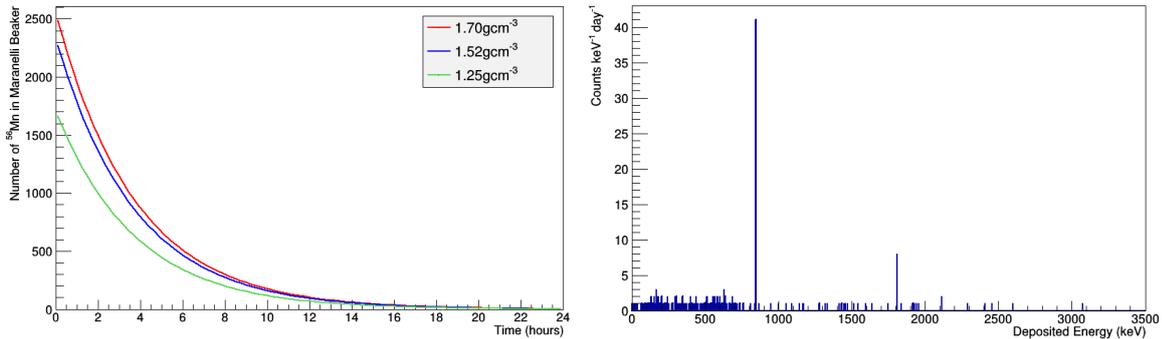
(a) Number of ^{56}Mn atoms inside the Marinelli beaker against time.

(b) Detector γ-ray spectrum from a 1.52 g/cm^3 solution assuming all ^{56}Mn decays in 24 hours.

Figure 5.13: The decay of ^{56}Mn inside the beaker is demonstrated on the left, and using this to assume it has all gone within 24 hours produces the detector spectrum on the right.

where T was the time in the bath, and ϵ_{Mn} is the probability of capture on Mn. N_T is the total number of ^{56}Mn in the bath after time T which can be calculated as:

$$N_T = \frac{C_\gamma}{\Gamma \epsilon_c} \frac{V_{beaker}}{V_{bath}} \tag{5.25}$$

where C_γ is the counts at the 846.7 keV line, Γ is the branching fraction of this line (0.9885) and V_{beaker} and V_{bath} are the Marinelli beaker volume and bath volume.

The uncertainties in this measurement come from the efficiencies for neutron capture by ^{55}Mn, the detector efficiency and the ^{56}Mn half life. Park et al [120] measure ϵ_c by using a sample of a known concentration of ^{56}Mn and the neutron capture efficiency by a Monte Carlo simulation

137

with inputs of nuclear cross sections. Here, the efficiencies are estimated by simulation and ideal background subtraction is assumed. Using Poissonian errors, the uncertainty on the capture probability is only 0.2%, and the uncertainty on the half life of ^{56}Mn is 0.008% [120]. This leaves the dominant source of uncertainty as the counting efficiency.

5.2.2.3   Conclusions and Feasibility

The experimental set up of method 1 could safely allow a measurement of the AmBe source activity to within 8%. In reality, the background of the detector would need to be subtracted, adding further uncertainty due to finite resolution. The geometry of the real-world setup, including details of the HPGe detector and shielding, will change the number of detected γ-rays. Additionally, systematic error is increased through uncertainty on the detector background rate, though this is a small amount for underground HPGe detectors such as at SURF or Boulby where backgrounds are highly reproducible. However, the indicative numbers reported here for a 2 kg crystal, suggest that with a 1 day exposure we may reasonably expect to constrain the source activity to within 20%. It should also be noted that this technique may be used to provide relative activity of the neutron calibration sources very easily. In this case, the results would be independent of the precise geometry.

For method 2, using the counts at the 846.7 keV line would hypothetically allow a measurement of the activity to $\frac{\sqrt{N}}{N} \approx 19\%$ for a 1.25 gcm^{-3} density, 16% for a 1.52 gcm^{-3} density and 15% for 1.70 gcm^{-3}. The effect of increasing the concentration isn't large but may be important depending on background. Therefore, it may be possible to constrain the neutron source activity to within 15% using this method, assuming perfect background subtraction, using the maximum solubility of MnSO$_4$ and only one extraction of solution into a Marinelli beaker. It would be possible to repeat the measurement with further extractions and reduce the uncertainty.

Method 1 exposes the HPGe detector to a small flux of neutrons, and method 2 appears to be the better choice in terms of the ease of repeating measurements and the safety of the HPGe detector. Both provide a novel way to measure the neutron source activity without removal from SURF and transportation, which can invoke health and safety concerns, and means the source is not available for calibration. This may be of future use at SURF for calibration sources.

5.2.3 ^{252}Cf Fission Neutrons

^{252}Cf undoes α-decay, with a branching ratio of 96.91%, and spontaneous fission with a branching ratio of 3.09%. Therefore, it is used in small neutron sources. Fission neutrons have an energy range of 0 to 13 MeV, with a mean of 2.3 MeV, and therefore cover a slightly larger energy range than AmBe. These neutrons can also be used to calibrate the NR band. Figure 5.14 shows the NR band from both AmBe and ^{252}Cf.
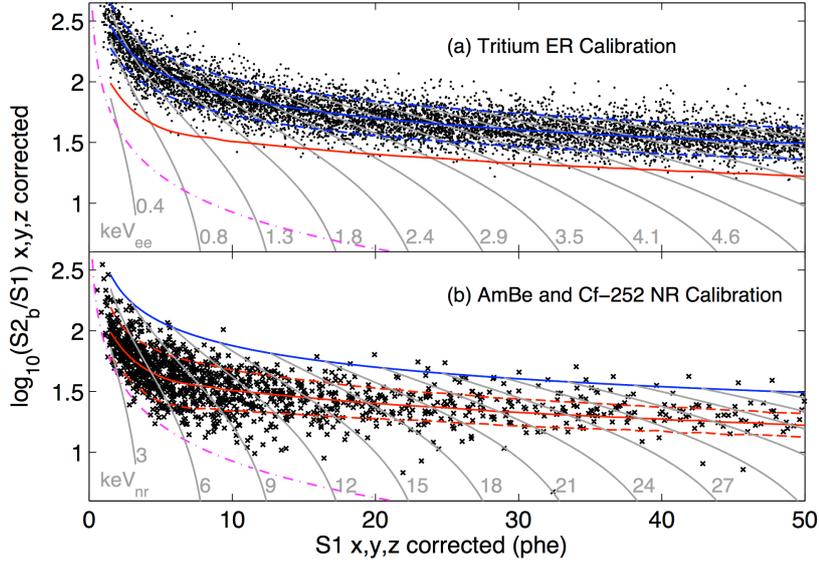
Figure 5.14: Top: LUX ER band from tritium data. Bottom: LUX NR band from AmBe and ^{252}Cf neutrons [106].

5.2.4 DD Neutrons

5.2.4.1 The Neutron Generator and Experimental Setup

A considerably more advanced technique than the AmBe and ^{252}Cf neutrons, which rely heavily on simulation and spectral matching with data to infer efficiencies, was developed for the run 3 reanalysis and run 4; a D-D neutron generator calibration. LUX have pioneered this technique for in-situ measurements. An Adelphi Technologies Inc. DD108 neuton generator produces a beam of mono-energetic neutrons of 2.45 MeV, using deuterium-deuterium fusion:

$$^{2}_{1}\text{H} + ^{2}_{1}\text{H} \rightarrow ^{3}_{2}\text{He} + ^{1}_{0}\text{n}; \quad E_n = 2.45\text{MeV}. \tag{5.26}$$

Deuterium gas is fed to the generator's plasma source. The generator can produce up to 1×10^8 neutrons/s if operated in a pulsed mode. For the calibration, it was operated using 100 μs pulses at a 500 Hz repetition rate.

The neutrons are fired through a 377 cm polyvinyl chloride conduit that passes through the LUX water tank and up to the edge of the cryostat. It is suspended with stainless steel rope from the top of the water tank. The tube is air filled, and has water-filled gaps at both ends totaling 6 cm. During normal WIMP search operation, the tube is not aligned with the TPC, to ensure maximum water shielding. For the calibration, the tube is raised to 16.1 cm below the liquid xenon surface and leveled to 1°. This position ensures more low multiplicity neutron scatters as there is a short distance to the liquid surface.

The neutron energy spectrum was measured prior to calibrations using a time of flight setup, measuring the time taken for neutrons to traverse 309±4 cm between a NaI(Tl) detector and a liquid scintillator detector [122]. Neutrons at 2.45 MeV are non-relativistic, meaning a simple kinetic energy calculation can be used. The mean neutron energy was measured as 2.40 ± 0.06 MeV, consistent with the expected 2.45 MeV. A Bonner sphere was used to measure the flux at the water tank as 78 ± 8 n cm^{-2}s^{-1}. The total livetime acquired for calibration was 107.2 hours.

5.2.4.2 Event Selection and Reconstruction

This technique does not rely on a scatter in LUX and the neutron being detected at some angle in an external detector, as is performed for small dedicated experiments used to measure LXe properties. Instead, multiple scatters within LUX are used, taking advantage of the precise position resolution of the detector. The basis of the technique is the measurement of the scattering angle of the neutron inside the LXe, which in turn allows a precise recoil energy calculation for each event, for energies within the WIMP search energy range. Using this alongside the size of the S2 signal for each event, the electron extraction efficiency and the mean size of an SE pulse, allows for a simple determination of Q_y. The L_y measurement is slightly more in depth, requiring single scatters, the previously determined Q_y and NEST simulations.

For the energy measurement, events containing 1 S1 and 2 S2s are selected for analysis. Only one S1 is observed despite the presence of two interaction sites because the maximum time between vertices is only ~ 30 ns, which is not resolvable due to the similar size of the S1 time constant. The S2 threshold applied is 36 phd (1.5 single electrons) raw area, i.e. prior to corrections. This is much lower than the threshold applied in WIMP search, as the accidental rate of two S2s is much lower than that for one. A cut was placed on the root-mean-square width of the S2 pulse at 775 ns, to ensure no overlapping S2s were selected. Events were only selected if they lay within the 4.9 cm diameter of the neutron beam projection into the active region, and also that they had traveled at least 15 cm into the active region along the beam path. Forward scatters were selected by ensuring the second scatter was deeper into the liquid xenon along the beam path than the first scatter. Furthermore, a cut of $\rho > 5$ was applied, where ρ is the 3D separation of scattering vertices, to remove a systematic bias from position reconstruction uncertainties. Finally, maximum signal cuts were applied on S1 and S2 to reject any electron recoil events in the data; S1 < 300 phd accepts more than 99% of nuclear recoils, and $S2 < 5000$ phd also accepts more than 99% of D-D S2s whilst rejecting all 39.6 keV$_{ee}$ gamma rays from inelastic scattering of neutrons on ^{129}Xe.

The angle between the two scatters, θ, can be measured after position reconstruction is applied for both S2s. The space between the S2s in the z co-ordinate is calculated from the time between the two S2s as usual. Knowing the initial energy of the neutrons allows for simple
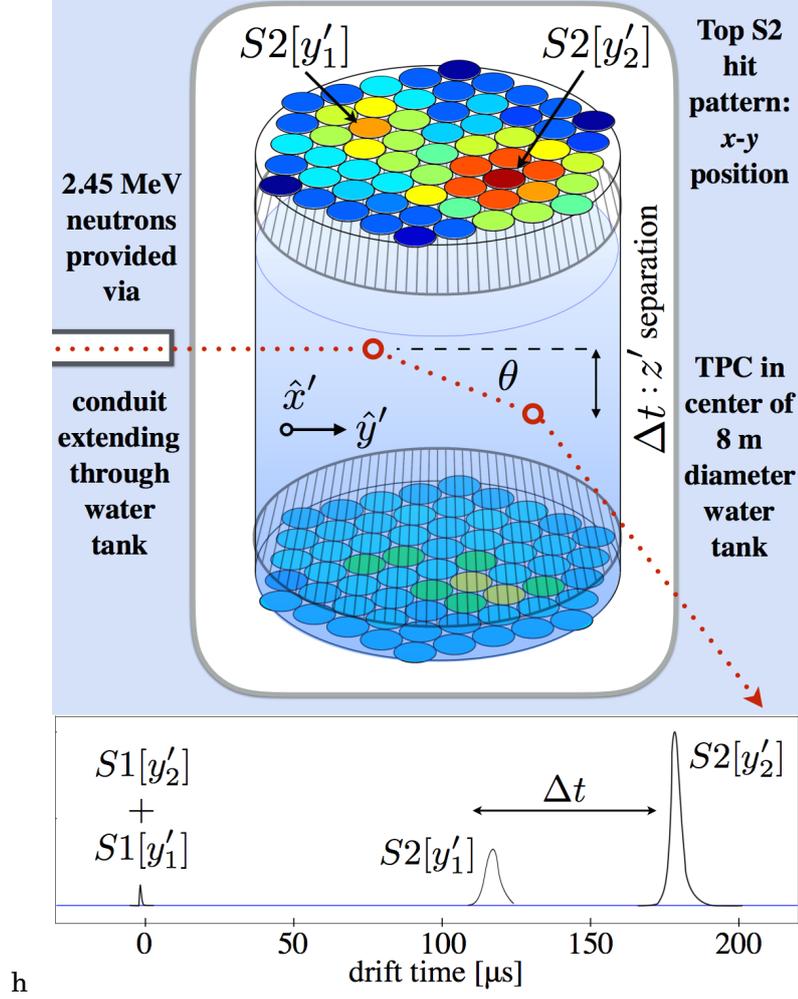
Figure 5.15: Schematic of the D-D calibration experimental setup. The monochromatic 2.45 MeV neutrons are collimated through an air-filled conduit passing through the water tank to the cryostat. This figure illustrates a D-D event that could be used for analysis: a neutron (red dotted line) enters the TPC, scatters twice, and then leaves. Bottom: the trace of such an event [123].

kinematical reconstruction of the recoil energy E_r:

$$E_r = E_n \frac{4 m_n m_{Xe}}{(m_n + m_{Xe})^2} \frac{1 - \cos\theta_{CM}}{2} \tag{5.27}$$

where θ_{CM} is the scattering angle in the centre of mass frame, related to the angle measured by reconstruction in the laboratory frame by:

$$tan\theta_{lab} = \frac{\sin\theta_{CM}}{\frac{m_n}{m_{Xe}} + \cos\theta_{CM}} \tag{5.28}$$

The assumption that $\theta_{lab}/\theta_{CM} \approx 1$ is valid to 1% for all scattering angles.

141

5.2.4.3 NR Band

The D-D data provides an excellent, high statistics measurement of the NR band, see figure 5.16. These results confirmed that the previous AmBe and ^{252}Cf results were a good measurement of the NR band. Note that compared to figure 5.14, the discrimination variable now uses the full S2 area, rather than just $S2_b$, which is the S2 light in the bottom PMT array only.
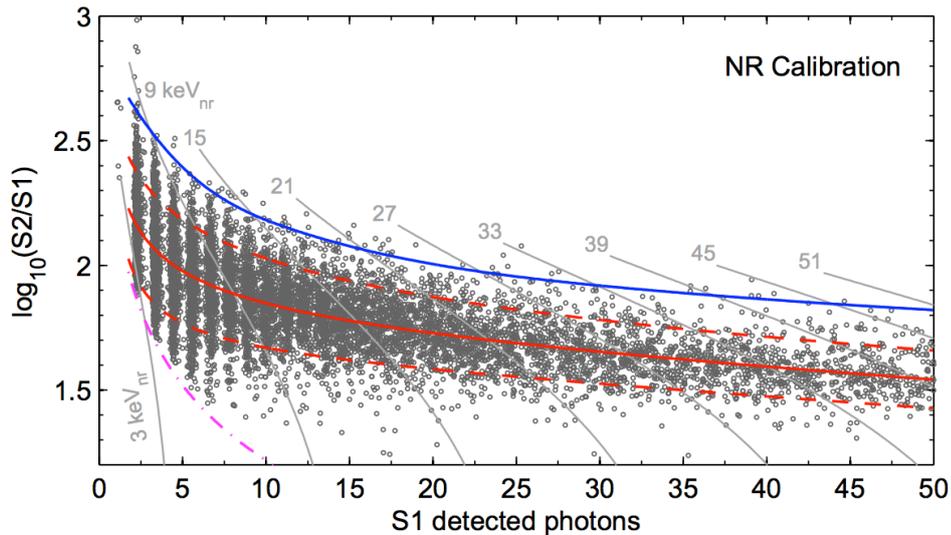
Figure 5.16: Newer LUX neutron calibration data shown for the WIMP search region of interest, 1 phd < S1 < 50 phd, against the discrimination variable $\log_{10}$(S2/S1). The NR band as measured with D-D calibration data is shown. The solid line represents the mean and the dashed lines represent the 10% and 90% contours. The S2 threshold of 165 phd is shown as a magenta dashed lines, and constant energy contours are shown in grey. The quantisation at low S1 is due to digital photon counting.

5.2.4.4 Q_y

The ionisation yield, Q_y, was measured from 0.7 to 24.2 keV$_{nr}$, using the size of the first S2 signal and E_R calculated with equation 5.27. Uncertainties in Q_y arise from the position reconstruction of the S2 (which gets worse as the S2 gets smaller). The statistical error on the x and y co-ordinates is maximum at ~ 2 cm at the S2 threshold, and is typically no larger than 1 cm for most signals, whilst for z it is ~ 0.1 cm. The systematic error on x and y was smaller with a best estimate of 0.35 cm and a maximum of 0.7 cm. Furthermore, an uncertainty of 0.6 cm on the beam entry point in x and z was also included.

A Monte Carlo simulation in LUXSim was used to produce corrections for the position reconstruction uncertainty and to also verify angular reconstruction. This took account of Eddington bias, the statistical fluctuation of the reconstructed energy of an event due to the non-zero resolution of the angle measurement. Furthermore, a Monte Carlo based model consisting of a Poisson convolved with a Gaussian was used to generate the expected probability of the number

of reconstructed electrons, given a mean number of electrons escaping recombination. This took account of statistical fluctuations around the S2 threshold. Using this along with the LXe purity and SE size, the distribution of the number of electrons in the gas can be determined. The energy resolution had a dependence of $a/\sqrt{E}$, where $a = 0.64 \pm 0.06$.

The ionisation signal, n_e, was determined by:

$$n_e = \frac{S2}{\epsilon_e \times \mu_{SE}} \tag{5.29}$$

where $S2$ represents the area of the S2 signal, μ_{SE} represents the single electron mean size and ϵ_e is the electron extraction efficiency.
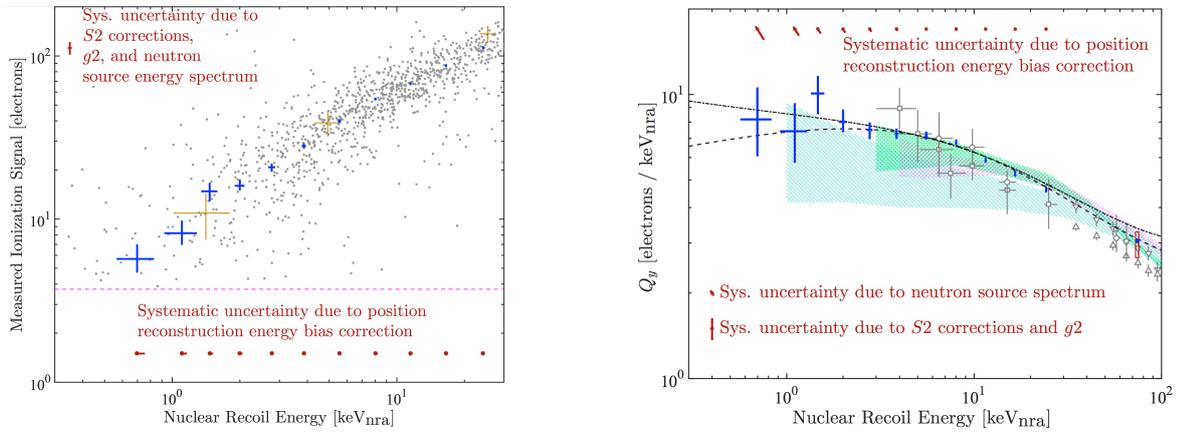
(a) Ionisation signals for events (grey points), showing the estimated error of the most precisely measured events as gold crosses, and the mean ionisation signal for each bin (blue crosses).

(b) Q_y (blue crosses). Other angle-based measurements shown in grey for comparison: squares (1 kV/cm), circles (4 kV/cm) [124], down triangles (0.3 kV/cm) up triangles (0.1 kV/cm) [125].

Figure 5.17: LUX 180 V/cm measurements of ionisation signals (left) and Q_y (right) [123]. Red error bars at the bottom show systematic uncertainties, and the Eddington bias uncertainties are shown at the top. On the right, hatched represent simulated-spectrum based measurements of 3.6 kV/cm (purple) [104], 730 V/cm (teal) [103] and 530 V/cm (green) [52], and the dashed and dot-dashed lines represent the LUX best-fit Lindhard and Bezrukov models respectively.

The final signal model was fit to the observed ionisation distribution for bins of energy using an extended unbinned maximum likelihood. The modelled resolution effects were included as a constrained nuisance parameter. Figure 5.17(a) shows the measured S2 signal in electrons as a function of nuclear recoil energy, and figure 5.17(b) shows the resulting measurements of Q_y, lower than anything previously measured for LXe.

5.2.4.5 L_y

The scintillation yield, L_y, is the more difficult measurement. For this, only single scatters were selected, with the S2 threshold returning to 55 phd to reduce accidental coincidence. The same beam purity cuts were used and a radial position cut of $r < 21$ cm was applied. Data quality cuts (such as a bad area cut of 219 phd) to remove events containing topologies such as e-trains from

large S2 pulses and photoionisation of impurities. The usual golden definition was altered so that there could not be any other unpaired S2s before the single scatter, ensuring the detector was quiet. These cuts resulted in a very clean population of single scatter nuclear recoils.

S1 and S2 distributions were modeled with a Lindhard-based NEST simulation that used the measured Q_y. The single scatters from data were placed into bins of S2 and the S1 distribution within each bin was compared to the model. The L_y measurement was made using an S2 range of 50 - 900 phd, corresponding to ~0-20 keV$_{nr}$. The data points ranged from 0.7-24.2 keV$_{nr}$. The best fit for each S2 bin was done by using a maximum-likelihood optimisation of the simulated S1 spectrum. Once the fit was complete, the parameter n_γ could be extracted.
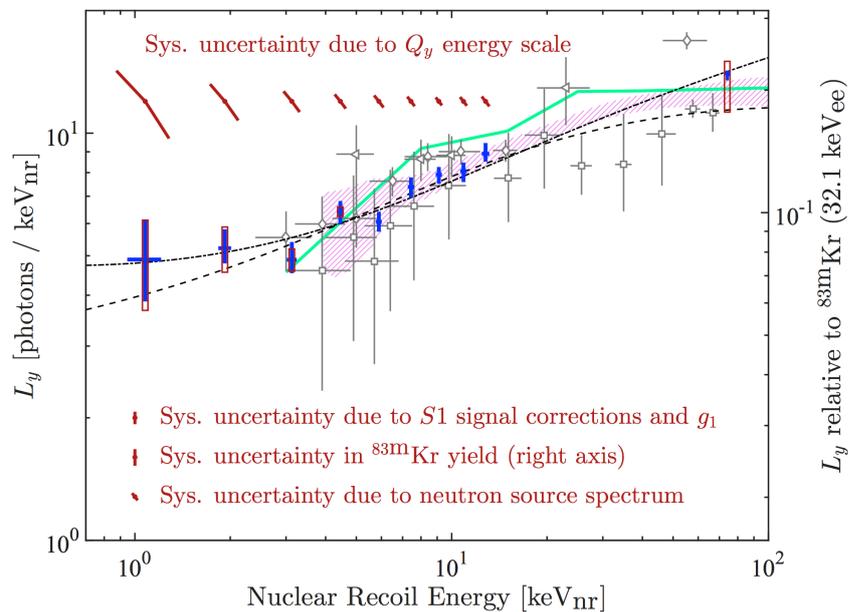
Figure 5.18: D-D measurements of L_y at 180 V/cm are shown as blue crosses. The red uncertainties at the top show the size of the 1σ systematic due to the Q_y measurement. The left y axis is absolute L_y and the right shows L_y relative to the LUX 32.1 keV$_{ee}$ ^{83m}Kr yield at 0 V/cm. The sizes of other systematics are shown at the bottom. Grey points (squares [126], left-facing triangles [127], diamonds [102]) show other angle based measurements at 0 V/cm. The purple band and cyan line are spectral fits corrected to 0 V/cm.

5.3 WIMP Search Results

LUX has completed two WIMP search runs over its lifetime, and three full analyses of the data. Table 4.1 at the beginning of chapter 4 lists the important differences in the analyses. Before discussing the results, it is useful to discuss the methods used by LUX and other experiments for setting limits on WIMP mass and interaction cross section. The following sections present limit setting, then the results from each of the analyses on spin-independent and spin-dependent (where available) cross sections.

5.3.1 Limit Setting

In the case of a null detection, dark matter experiments can place exclusion limits on the WIMP mass and interaction cross section to show which values of the phase space have been ruled out. The standard format of these plots has the WIMP mass in GeV along the x-axis and the cross section in either cm^2 or barns on the y-axis.

LUX used a double-sided profile likelihood ratio (PLR) test for limit setting. Cross sections are scanned over for each WIMP mass to construct a 90% confidence interval, using test statistic distributions evaluated by Monte Carlo using the RooStats package in ROOT. The profile likelihood ratio is defined as [128]:

$$\lambda(\mu) = \frac{L(\mu, \hat{\hat{\boldsymbol{\theta}}})}{L(\hat{\mu}, \hat{\boldsymbol{\theta}})} \tag{5.30}$$

where μ is the strength of the signal process (i.e. $\mu=1$ is the signal hypothesis), θ is a series of parameters characterizing probability density functions (PDFs) for signal and background, and L is a likelihood function. $\hat{\hat{\boldsymbol{\theta}}}$ represents the value of $\boldsymbol{\theta}$ that maximises the likelihood function for the specified hypothesis μ (for a set cross section), whilst $L(\hat{\mu}, \hat{\theta})$ is the maximised unconditional likelihood function, where everything is allowed to float.

To get the likelihood of a dataset, the probabilities of each event are calculated by evaluating the total signal and background PDF using measured quantities. For run 3 these were S1, S2, radius and depth, but for run 4, ϕ was additionally included because of field nonuniformities. These probabilities are multiplied together, along with a Poisson term that accounts for the fact we expect a total number of events N_{exp}, and observe N_{obs}. Gaussian constraints that multiply the likelihood are applied to the number of background events based on the measurements described in chapter 4. These backgrounds are included as nuisance parameters within $\boldsymbol{\theta}$. The PLR was used within the fiducial volume, but spatial background models were validated using data from the whole active region. The energy spectrum of the nuclear recoils was modeled with a standard isothermal Maxwellian velocity distribution with v_0 = 220 km/s, v_{esc} = 544 km/s and local dark matter density ρ = 0.3 GeV/cm^3 as described in section 1.2.3. The average Earth velocity is taken to be 245 km/s. Uncertainties in astrophysical parameters are considered beyond the scope of the experiment.

The test statistic is given by [128]:

$$t_\mu = -2\ln\lambda(\mu) \tag{5.31}$$

Due to the limited amount of data that is used within LUX's PLR analyses, the approximation of a χ^2 test statistic distribution is not valid. Instead, fake datasets (known as MC toys or pseudo-experiments) are generated, t_μ calculated for each one and a distribution $f(t_\mu|\mu)$ is built for each test value of the cross section. The p-value is then the integral from $t_{\mu,obs}$ (calculated with actual data) to infinity of the generated distribution. Therefore, the p-value is obtained by [128]:

$$p = \int_{t_{\mu,obs}}^{\infty} f(t_\mu|\mu) \quad dt_\mu \tag{5.32}$$

145

This is done for a range of test cross sections until $p = 0.1$ - meaning that only 10% of the MC datasets simulated yield more extreme disagreement than the observed data, i.e. a 90% confidence limit.

The shape of the limit curve can be investigated by looking at an "event contour". If no events are observed with zero expected background, the 90% upper limit on the number of events is 2.44 from Feldman-Cousins statistics. Therefore, plotting the cross section expected for each WIMP mass that would generate 2.44 events in an exposure gives an upper limit of sorts. This was done as a simple cross-check of the run 3 reanalysis limit curve after a mistake in the PLR code led to the limit being too low at low energies, but it is useful for understanding the limit curves. This contour can be calculated using:

$$\sigma_{M_\chi} = \frac{2.44}{\epsilon \int \frac{dN}{dE_R} \cdot \eta(E_R) \quad dE_R} \tag{5.33}$$

where ϵ has the same meaning as before; the detector exposure in kg-days. $\eta(E_R)$ is the overall golden efficiency after all cuts as shown in section 3.9, and the recoil spectrum is in events/kg/day/zb/keV, as in figure 5.19(a), and must be calculated for each M_χ.
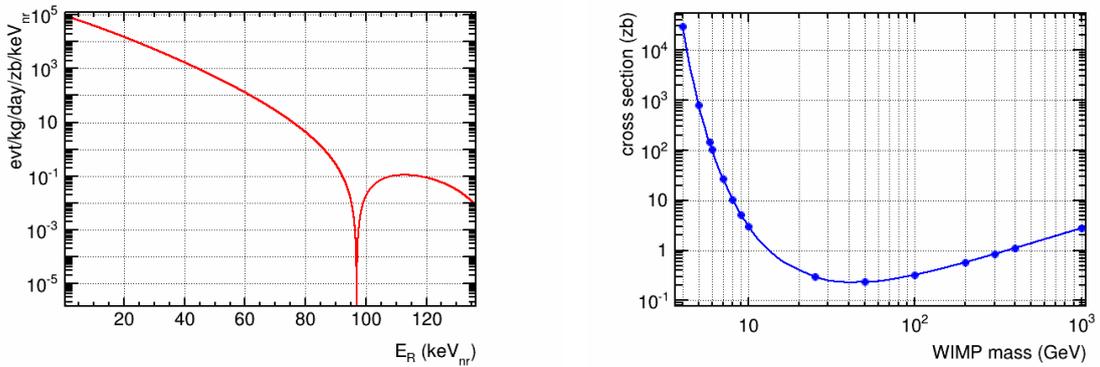
(a) Example WIMP-nucleus differential recoil spectrum.

(b) 2.4 event contour/ upper limit.

Figure 5.19: Left: WIMP-nucleus recoil spectrum for a 50 GeV WIMP (assuming $\sigma = 1 \times 10^{-36}$ cm^2). The feature just below 100 keV$_{nr}$ is where the spherical Bessel function (describing the density of the nucleus) in the nuclear form factor falls to zero for certain recoil momenta [129]. Right: the corresponding 2.4 event curve as a proxy for a limit. Blue circles show points where the cross section was calculated, and the blue line is an extrapolation between those points.

The rapid fall off at low energy can be accounted to the fall-off in efficiency at low energies, whilst the slow exponential increase as WIMP masses increase is seen to come from the differential rate. As the local dark matter density, (ρ in equation 1.31), is fixed, as the WIMP mass increases, the number density of WIMPs decreases. This also causes a loss in sensitivity, as the less WIMPs that are passing through the target, the smaller the chance of an interaction.

5.3.2 Run 3

Run 3 was the first underground WIMP search data-taking run and originally consisted of 85.3 livedays of data collected between April and August 2013. As LUX collects large amounts of data (70–80 GB a day) the first step of any analysis must be to select only those events of interest. Therefore, a series of cuts must be applied to the data, and these are listed in table 5.5 alongside the number of events remaining within the data collected for the run 3 WIMP search after each is applied sequentially. The detector stability cut was used to remove periods of data where either

Table 5.5: Table showing the number of remaining events after each analysis cut in the LUX run 3 WIMP search. A remarkable decrease from over 80 million events to just 160 is seen after all cuts are applied.

Cut	Events Remaining
all triggers	83,673,413
detector stability	82,918,902
single scatter	6,585,686
S1 energy (2-30 phe)	26,824
S2 energy (200-3300 phe)	20,989
SE background	19,796
fiducial volume	160

the liquid level, the gas pressure or the grid voltages were outside their nominal range. This removes spurious data as well as unusual events that affect data quality such as circulation outages. The next cut is the single scatter or golden event cut as defined in section 3.5. This removes events that do not contain a recoil signal as well as any multiple scatter neutron events. The energy cuts keep only those events within the WIMP search energy region of interest, and the SE cut removes data taken when single electron rates were high. Finally, the fiducial cut, defined as a reconstructed radius of less than 18 cm, and a drift time between 35 μs and 305 μs, selects only the inner region of the detector. The power of xenon self-shielding is evident from the reduction of ~20,000 events to only 160 for final analysis. These 160 events are shown in figure 5.21, plotted in S1 vs $\log_{10}$(S2/S1) space for discrimination. Almost all events lie within the ER band, and those that don't are consistent with leakage from the ER band, therefore no WIMP signal was observed.

The first results on the spin-independent (SI) WIMP-nucleon scattering cross section were published in October 2013, with an upper limit of 7.6×10^{-46} cm^2 at a WIMP mass of 33 GeV. 85.3 livedays of data in an 118 kg fiducial volume were used. A PLR analysis, as discussed in section 5.3.1, was performed to set limits on the WIMP-nucleon scattering cross section. L_y was assumed to fall to zero below 3.0 keV$_{nr}$ as a conservative cut off (at the time, this was the lowest energy for which a direct light yield measurement existed, since LUX had not yet performed D-D generator neutron calibrations). Uncertainties in the NR yield were not profiled as the assumed model was in good agreement with LUX data. The observed PLR for zero signal was consistent with its simulated distribution, giving a p-value for the background-only hypothesis of 0.35. The limit shown in figure 5.22 is at 90% confidence and rules out the phase-space above the line.
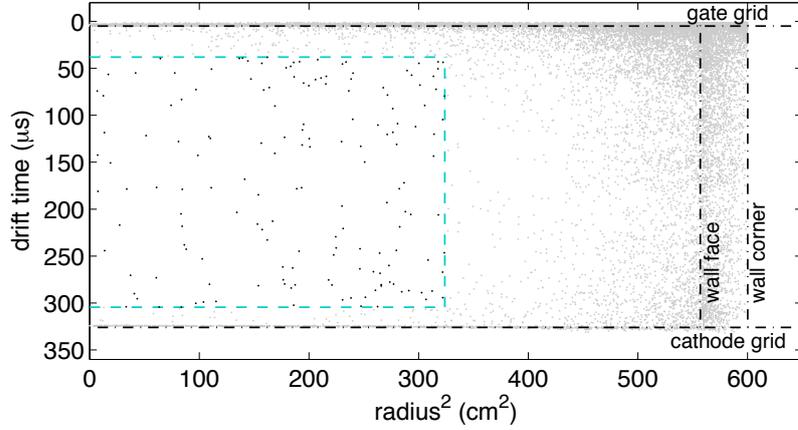
Figure 5.20: LUX run 3 WIMP search data in R^2 vs Z, showing all events with an S1 between 2 and 30 phd. The fiducial volume is shown by the cyan dashed line. The physical location of the grids and PTFE wall are also shown [106].
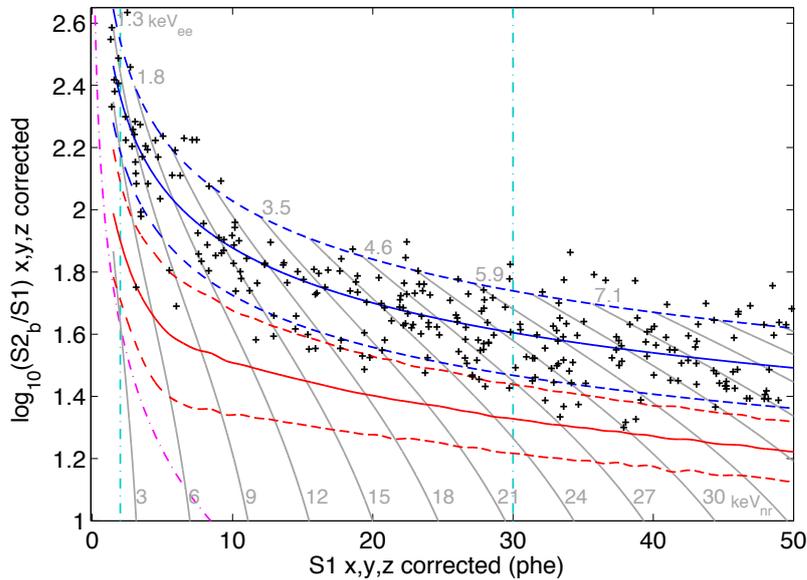
Figure 5.21: LUX run 3 WIMP search fiducial data in S1 vs $\log_{10}$(S2/S1), the discrimination phase-space [106]. Blue indicates the electron recoil band, and red the nuclear recoil band. The band means are shown by solid lines, whilst the dashed lines show 1.28σ contours (10% band tails). Vertical dashed cyan lines indicate the 2 - 30 phe search region used for S1. Grey contours show bands of energy in keV$_{nr}$.
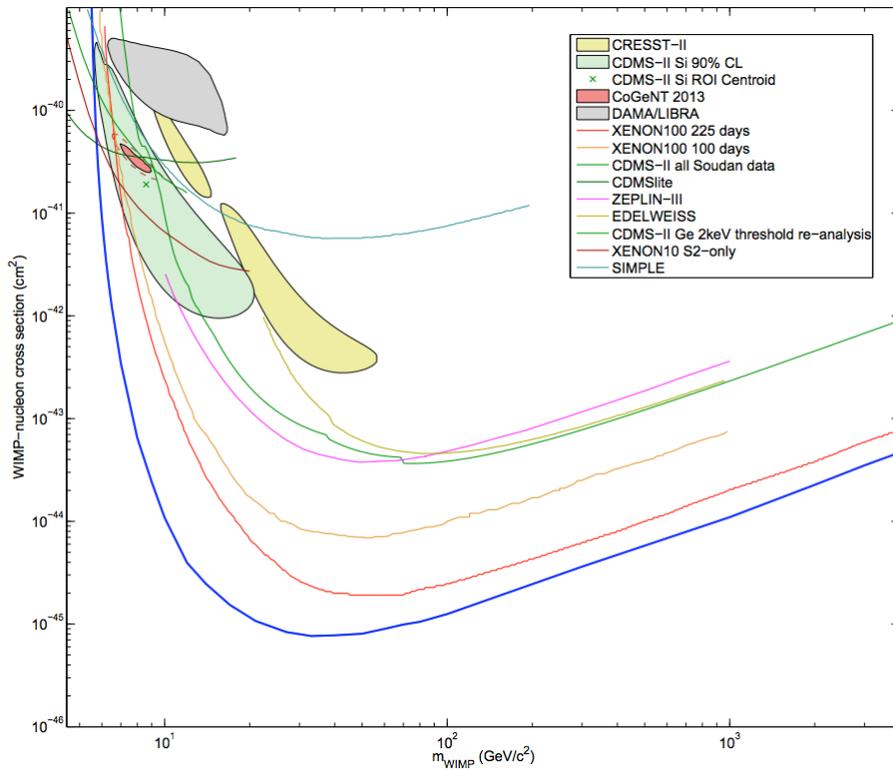
Figure 5.22: LUX Run 3 limit on WIMP mass and spin-independent WIMP-nucleon scattering cross section (blue). Previous results from other experiments are shown in various colours, and claimed discoveries from CRESST-II, CDMS, CoGent and DAMA/LIBRA are identified with 90% CL shaded contours. The LUX limit strongly disfavours the interpretation of these results as WIMP signals [106].

5.3.3 Run 3 Reanalysis

Due to a number of improvements in both LUX analysis and calibrations, it was decided that the collaboration would perform a reanalysis of the run 3 data and publish a new limit. The improvements key to this reanalysis were as follows:

- D-D calibration data measuring L_y and Q_y to lower than any previous measurements - see section 5.2.4

- Better selection of data; low rate ^{83m}Kr data was found to be suitable for WIMP search, so livetime was increased from 85 to 95 days.

- An improved background model allowing an increase in the fiducial volume from 118 kg to 145 kg - see section 4.4.

- Improvements to signal finding and classification - see chapter 3.

These improvements led to increased sensitivity at all WIMP masses, with most improvement below 20 GeV. The new measurements with D-D calibrations were dominant and showed the region 3.3 - 5.2 GeV/c^2 to be detectable in xenon for the first time.
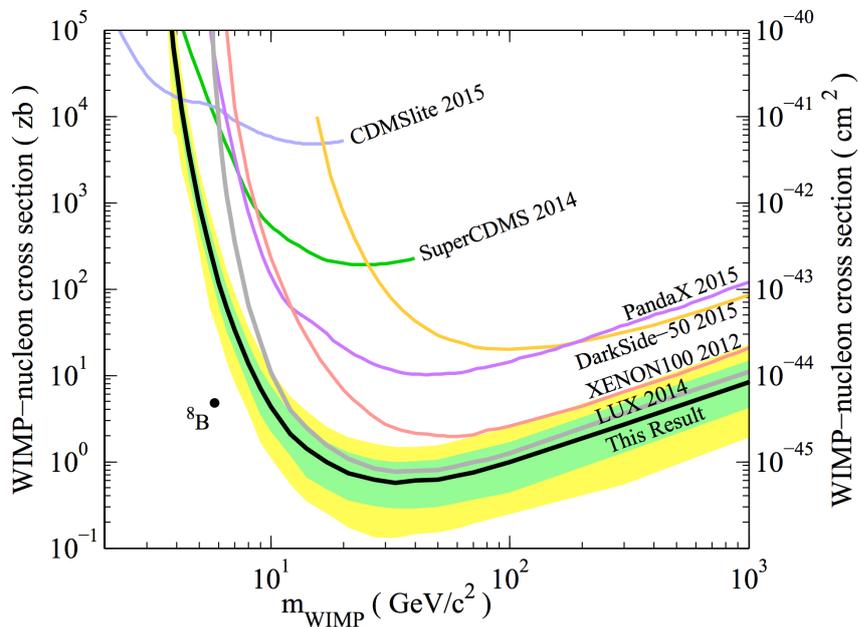
Figure 5.23: LUX run 3 reanalysis limit on WIMP mass and spin-independent WIMP-nucleon scattering cross section at 90% CL (black line). Green and yellow shaded bands show the 1 and 2 σ background-only trials. Limits from the first LUX analysis [106] are shown for comparison as well as those from SuperCDMS [130] (green), CDMSlite [131] (light blue), XENON100 [43] (red), DarkSide-50 [38] (orange), and PandaX [44] (purple). The black dot shows the location of expected events from neutrino-nucleus scattering of ^{8}B solar neutrinos [46].

5.3.4 Run 3 Spin-Dependent WIMP-Nucleon Scattering

The run 3 data was also used to set constraints on the WIMP-neutron and WIMP-proton cross-sections for spin-dependent scattering. Data selection was conducted in the same manner; single scatter events from within a volume defined by R< 20 cm and 38-205 μs drift time were selected.

Cancellation between the spins in nucleon pairs removes the A^2 enhancement (see section 1.2.4.1), and the recoil spectrum is suppressed compared to the SI case, therefore limits are naturally less constraining. Natural xenon contains ^{129}Xe at 29.5% and ^{131}Xe at 23.7% which both contain an unpaired neutron, allowing much higher sensitivity on the WIMP-neutron cross section than the WIMP-proton. The contributions to the differential event rate from each xenon isotope are added together in order to produce a signal PDF and as for the SI case a PLR using S1, S2, radius and depth is used to set upper limits. The limits are shown in figure 5.24 and reach minimums of σ_n=9.4×10^{-41} cm^2 and σ_p=2.9×10^{-39} cm^2. For WIMP-neutron scattering, this is lower than any previously published limit.

(a) WIMP-neutron cross section, σ_n.

(b) WIMP-proton cross section, σ_p.
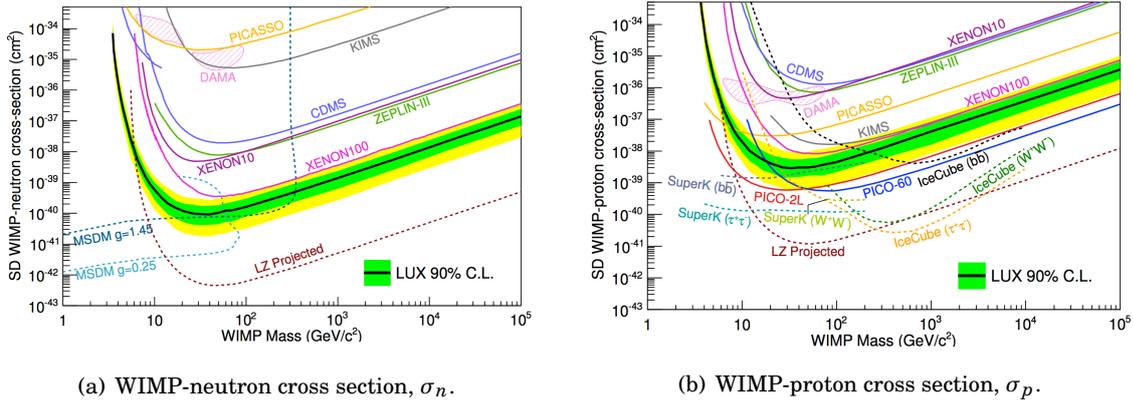
Figure 5.24: Spin-dependent limits at 90% CL placed by LUX using 1.4×10^4 kg/days of data (black line). $\pm 1\sigma$ and $\pm 2\sigma$ bands are shown in green and yellow respectively [132]. Also shown are limits from CDMS [133], KIMS [134], PICASSO [50], PICO-2L [48], PICO-60 [49], XENON10 [135], XENON100 [52] and ZEPLIN-III [41, 136], and the favoured region at 3σ from DAMA [137]. Indirect limits from Ice-Cube [33] and Super-K [53] are shown, and collider limits from CMS using the MSDM model [138]. Finally, the projected LZ limit is shown [139].

Furthermore, constraints on the couplings a_p and a_n can be determined using the formalism in [140]:

$$\sum_A \left(\frac{a_p}{\sqrt{\sigma_p^A}} \pm \frac{a_n}{\sqrt{\sigma_n^A}} \right)^2 > \frac{\pi}{24 G_F^2 \mu_p^2} \tag{5.34}$$

where $\sigma_{p,n}^A$ are the limits on the WIMP-proton and WIMP-neutron cross sections for a nucleon with mass A. The results for four different WIMP masses are shown in figure 5.25. LUX's limits on spin-dependent scattering are complementary to experiments with a greater WIMP-proton sensitivity, such as PICO, who utilize an unpaired proton in fluorine nuclei in their C_3F_8 target.
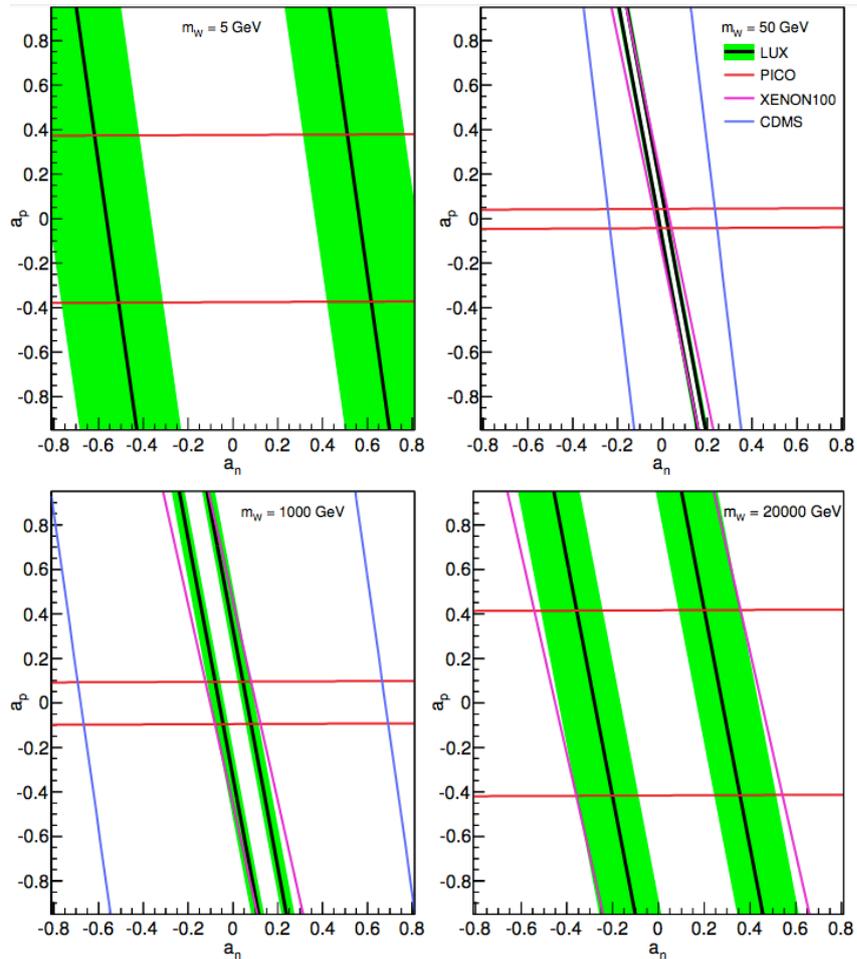
Figure 5.25: Constraints on the WIMP couplings to protons (a_p) and neutrons (a_n) at 90% CL for WIMP masses of 5, 50, 1,000 and 20,000 GeV [132]. Other limits shown are from CDMS [141], PICO-2L [48], PICO-60 [49] and XENON100 [52].

5.3.5 Run 4

Run 4 was the second and primary WIMP search data-taking period underground, which began with the intention of running for 300 livedays. In total, 332.0 livedays of data taken between September 2014 and May 2016 were used for analysis.

Run 4 presented a significant challenge as the data was taken with a time varying electric field. Between run 3 and run 4, LUX undertook 'grid conditioning', where cathode, gate and anode grid voltages were increased to just above the level of discharge and maintained there for extended periods of time in order to attempt to burn off any impurities on the grid wires. This was to investigate potential improvement through increased drift field strength in S2 identification, to extend thresholds to lower energies through increased field in the gas extraction region, and to provide valuable input for the design phase of LZ. The conditioning was successful in increasing

the extraction field from 2.9 kV/cm to 3.5 kV/cm and consequently improving the LUX extraction efficiency from $49 \pm 3\%$ to $73 \pm 4\%$ [36]. However, this had an unfortunate unintended side-effect; the radial component of the LUX drift field increased significantly, see figure 5.26. Not only this, but it was time varying and radially asymmetric. The wall radius at different depths in the detector was found to vary at different rates. It has been shown with electric field models that the effect is consistent with a build-up of negative charge on the PTFE walls; this charge is concentrated in the upper portion and appears to increase over time. This charge is thought to be the effect of exposure to coronal discharge during the grid conditioning, and will be studied in more detail when LUX is dismantled.

For analysis purposes, the time varying field could be mitigated due to the weekly calibrations performed with ^{83m}Kr. A 3D electrostatic model of LUX was built using the COMSOL multiphysics package [108], and this included a time-specific charge density in the PTFE panels. The ^{83m}Kr calibration data was used to fit the charge density. These models could be used to produce a map of co-ordinates in true space to the observed space, which is distorted due to the field variation. To generate corrections for S1 and S2, two tools were needed. Firstly, to account for field effects, the ratio of two ^{83m}Kr S1s can be used as a measurement of the field only. Secondly, the position of the peak of the tritium β spectrum (2.5 keV) varies almost only with geometrical effects, so tritium calibrations could be used for geometry-only corrections. Furthermore, the mass of any fiducial volume chosen could be checked by finding the ratio of uniformly distributed ^{83m}Kr events surviving the cut to the total number and applying this to the total mass of xenon, 250 kg. g_1 and g_2 were measured regularly using monoenergetic ER sources as described in 5.1. For run 4, g_2 was found to vary between 18.92 ± 0.82 and 19.72 ± 2.39 phd/e, and g_1 gradually fell from 0.100 ± 0.002 to 0.087 ± 0.001 phd/γ.
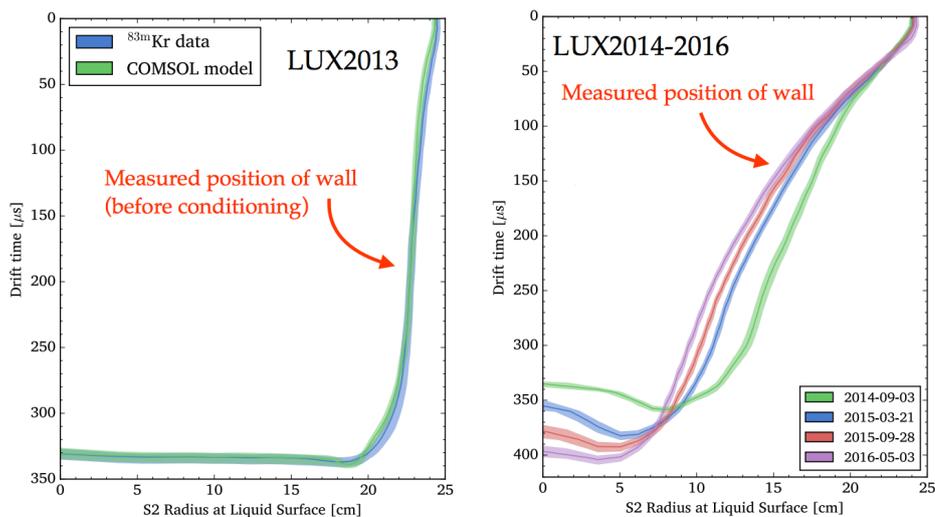
Figure 5.26: Position of the wall as measured with ^{83m}Kr data [142]. Left: 2013 data, before grid conditioning. Right: 2014-2015 data, showing the evolution in time using four datasets.

The run was broken into 16 total bins that effectively behave as separate experiments; 4 time bins to account for the changing field, and 4 volume voxels sliced in z that could each be treated as having a constant field and other uniformities in detector parameters. This was found to be optimal when considering how adequately the field variation is accounted for and the sparsity of the calibration data. Each of the 16 bins required its own response model from NEST, and the models were validated by comparison of their predicted ER band means those from calibration data. Calibration of the 16 bins can be seen in figure 5.27.
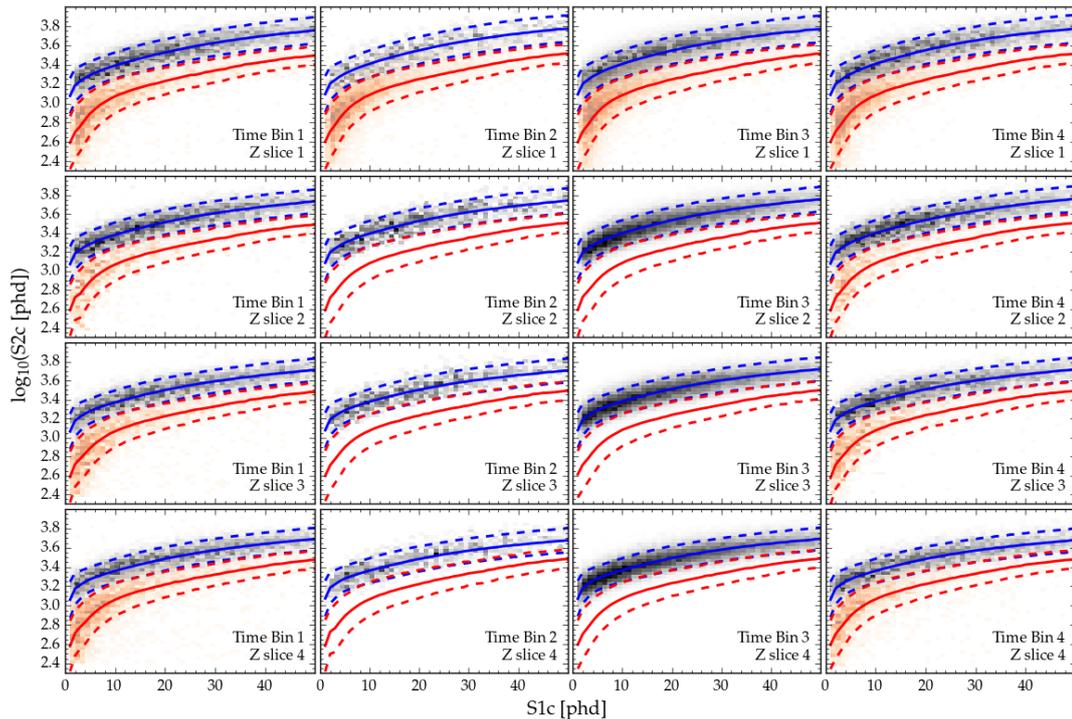
Figure 5.27: Calibration data and the ER (blue) and NR (red) bands in the 16 time and drift bins used for LUX run 4 analysis [142]. Solid lines show means and dashed lines show 10% and 90% contours, all from NEST models. S1c and S2c refer to corrected S1 and S2 areas.

Selection of data was similar to run 3, with a few exceptions. Each time bin had its own radial cut of 3 cm inwards from the measured position of the wall within that time bin (see the right of figure 5.26), and the drift cut applied was 40 - 300 μs, slightly smaller than that for run 3. The S2 threshold was 200 phd (raw), and the upper limit for the corrected S1 was 50 phd as in the run 3 reanalysis.

A new blinding technique, salting, was used for run 4 to check for bias, as run 3 was an unblinded analysis. Salting is the addition of fake NR events created by slicing together S1s and S2s from tritium ER calibration data to WIMP search data. LUX analysts did not know which events were salt and if all the salt survives all analysis cuts, this demonstrates that potential WIMP signal is not being accidentally cut away. Figure 5.28(a) contains the salted events shown in blue. Once salt was removed, there were three events within the fiducial volume that were

reconstructed just below the bottom contour of the NR band, see the red data points in figure 5.28(b). If put into the PLR, these events would give a detection of WIMPs at all masses. However, when handscanned, the events were found to be anomalous. Two contained S1s that had over 80% of their light in only one top array edge PMT, consistent with light from an energy deposit outside the TPC leaking into the detector through a gap near the PMT array. The third other had an unusual S1 shape characteristic of gas scintillation, which occurred less than 1 second after a high event rate. Therefore, these events did not correspond to interactions within the TPC. It was found these could be removed using loose cuts; a S1 area dependent cut on the maximum area in an individual PMT to remove the first two, and a cut on the prompt faction. These cuts were tested on tritium and DD calibration data and were found to have >99% acceptance, with a flat dependence on S1 area. Figure 5.28(c) shows the final results. The remaining events in the NR band were consistent with background predictions. A PLR was constructed using the background model described in 4.4, and spin-independent cross-sections were scanned over for each WIMP mass. The likelihood was constructed using the full signal + background PDF evaluated for each event, a Poissonian term for the number of events and a set of Gaussian constraints (parameters such as the Lindhard k, background counts, random coincidence counts). The 16 bins were taken as separate exposures with their own specific PDFs for S1, S2, r, ϕ and z. The p-value at 100 GeV was 0.39, consistent with the background-only hypothesis. The resulting limit has a minimum of 2.2×10^{-46} cm^2 at 50 GeV, and is seen as the grey curve marked 'LUX2014–16' in figure 5.29.

A PLR analysis was then performed combining both run 3 and run 4 data, by adding the run 3 exposure as separate detector number 17. Response, signal, and background models for the run 3 exposure were unchanged from [46], and nuisance parameters were treated separately. The resulting 90% confidence limit is also shown in figure 5.29, which has a minimum of 1.1×10^{-46} cm^2 at 50 GeV. A conservative power constraint was applied at the -−1σ extent of the projected sensitivity in order to avoid excluding cross sections where sensitivity is unreasonably enhanced through chance background fluctuation [36]. With this result, LUX remains (at the time of writing) the most sensitive WIMP search experiment to-date. The PandaX experiment recently published a very similar limit to the original run 4 only limit (see figure 5.29). PandaX is a dual-phase liquid xenon TPC experiment located at the China Jin-Ping Underground Laboratory (CJPL), containing 500 kg of xenon. Their results, announced shortly after the LUX run 4 results at IDM2016 [142], were from 98.7 livedays of data taking, giving a total exposure of 3.3×10^4 kg-days, and a minimum cross section of 2.5×10^{-45} cm^2 at 40 GeV [45]. Whilst the LUX run 4 limit was slightly more sensitive, and the run 3 and run 4 combined limit even more so, PandaX is still currently running and will soon surpass LUX's best sensitivity. At the time of writing, however, the limit shown in figure 5.29 is the most recent published by PandaX.

LUX has extended the search into electroweak parameter space by a factor 10 through its science runs, and with its exquisite low energy sensitivity, has conclusively ruled out all known claims of possible signal from other experiments at low masses.

(a) With salt in blue.

(b) Salt removed, bad events in red.

(c) Final data, after removal of salt and bad events.
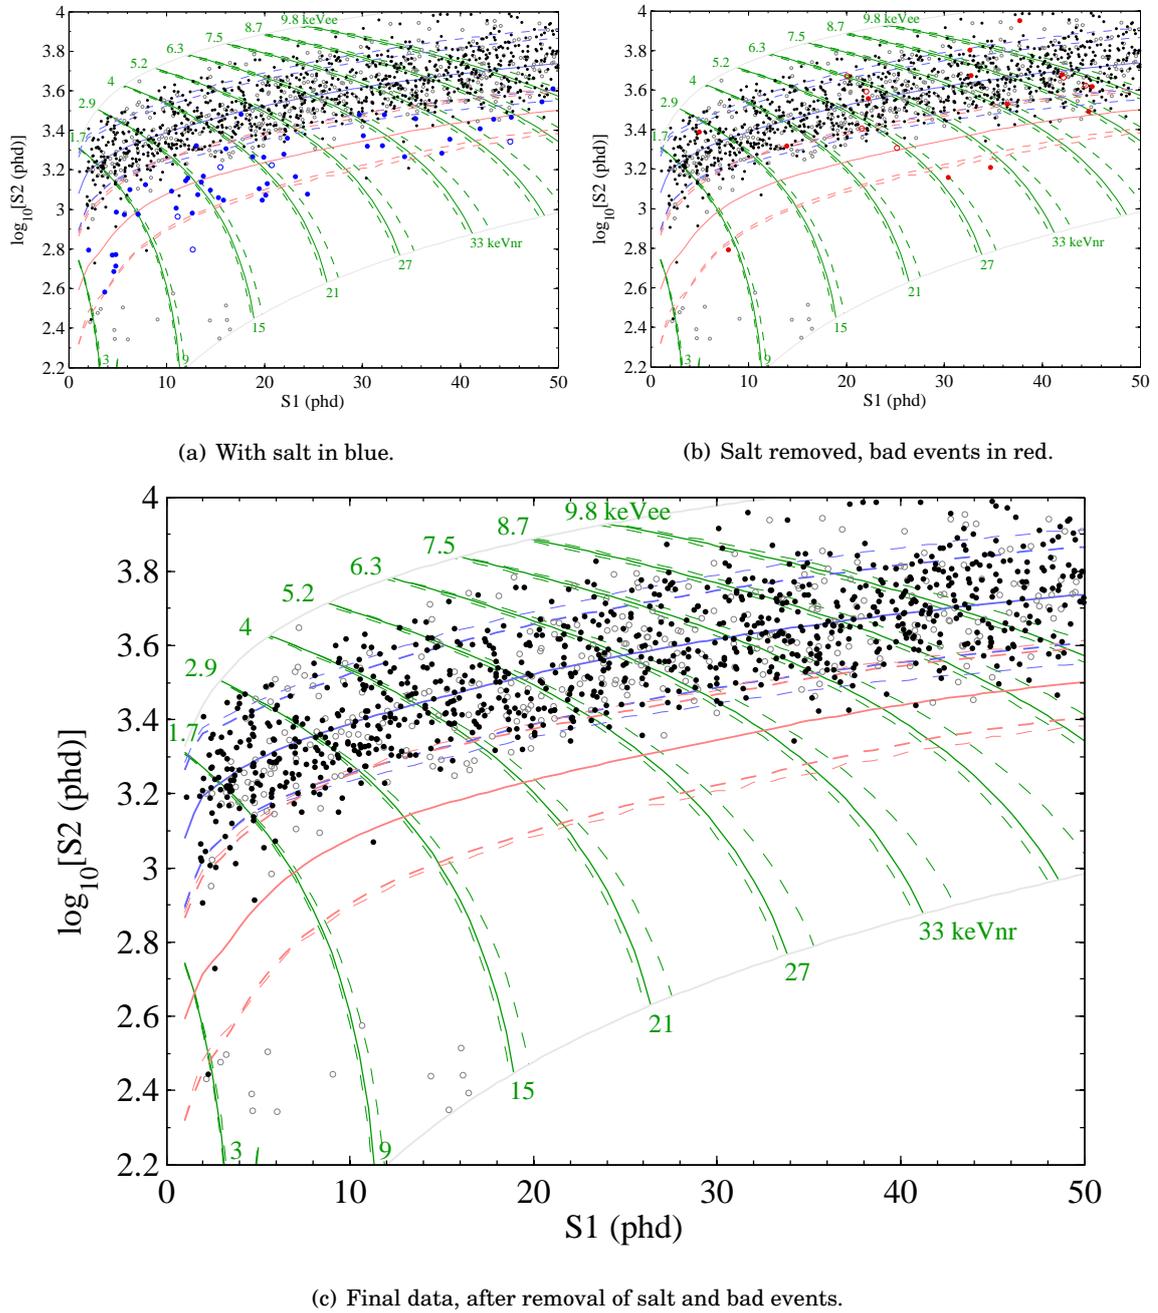
Figure 5.28: Run 4 discrimination plots [142]. Filled data points are within the fiducial volume, whist circled points are within 1 cm of the boundary. The ER and NR bands are exposure-weighted averages and the fainter dashed lines demonstrate the boundaries from lowest and highest S2 models to show the scale of variation of the 10% and 90% contours. Green lines show exposure-weighted mean energy contours labeled for ER (keV$_{ee}$) at the top and NR (keV$_{nr}$) at the bottom, and dashed line again demonstrate the extrema. They grey curves across the top and bottom show the data selection boundary applied before the PLR.
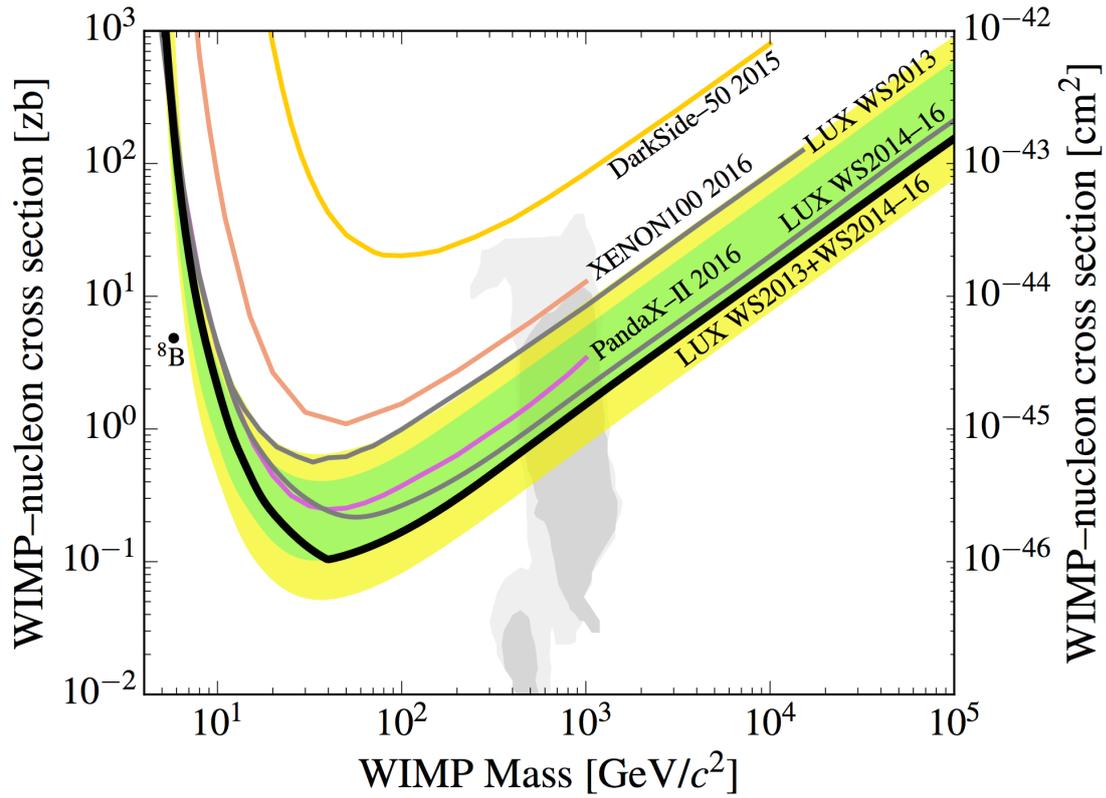
Figure 5.29: 90% confidence limit set on WIMP properties by LUX run 3 and run 4 combined (black) [36]. Green and yellow bands show the 1 and 2σ background-only trials. Limits from the LUX run 3 reanalysis [46] and run 4 only analysis are shown in grey for comparison as well as those from XENON100 [43] (red), DarkSide-50 [38] (orange), and PandaX [44, 45] (purple). The black dot shows the location of expected events from neutrino-nucleus scattering of ^{8}B solar neutrinos. Parameters favoured by SUSY CMSSM [143] are indicated as dark and light grey regions (1 and 2σ respectively).

Chapter 6

The LZ Experiment

L UX-ZEPLIN, or LZ, is LUX's multi-tonne successor, also named after the pioneering
ZEPLIN series of single- and two-phase xenon experiments that were conducted at
Boulby, UK. The collaboration consists of physicists and engineers from 31 institutions
in the USA, UK, Portugal, Russia and South Korea. Its sensitivity will be ideally matched to
explore the bulk of the remaining theoretically favoured electroweak phase space towards galactic
dark matter discovery. In order to reach the desired sensitivity LZ must have unprecedentedly
low backgrounds. This chapter will describe the design of LZ, discuss sources and mitigation of
backgrounds and present expected sensitivity.

6.1 LZ Assumptions and Goals

LZ has been designed to reach a sensitivity of 2×10^{-48} cm^2 or 2 yoctobarn for a 50 GeV WIMP
(spin independent scattering), after an exposure of 1000 tonne days. This estimate is made using
mostly conservative assumptions of important quantities, shown in table 6.1, alongside their
values in LUX. Note that the only estimate that is higher than achieved in LUX is the electron
extraction probability. This is expected to improve significantly as LZ will not have the same
electric field limitations as LUX did.

Many of the assumptions are based on the larger size of LZ, for example including a 3-fold
coincidence requirement in the PMTs for S1s, necessary because the larger quantity of xenon
and PMTs results in a higher single photon rate. The reduced light collection efficiency is also
due to the increased size of the TPC and the larger quantity of xenon.

The high level of expected sensitivity relies on several key differences to LUX. First of all,
fiducial volume will be ~40× bigger at 5,600 kg, and the aim is to collect 1000 days of data,

Table 6.1: Quantities assumed for the LZ sensitivity projections alongside real values from LUX.

Quantity	LZ	LUX
Threshold (50% efficiency)	6 keV$_{nr}$	3.3 keV$_{nr}$
S1 range	3 - 30 phd	2 - 50 phd
S2 range	> 450	> 165
S1 light collection efficiency	7.5%	14%
Photocathode efficiency	25%	30%
Electron extraction probability	95%	65%
ER discrimination	99.5%	99.8%

leading to a total nominal exposure of 5,600 kg days. Furthermore, LZ will take advantage of a vetoing system comprising of a scintillator outer detector and an instrumented skin region of liquid xenon which will significantly reduce neutron and γ-ray backgrounds by tagging them as clear background events. Finally, LZ will be built using extremely radio-pure materials that have undergone extensive screening and cleaning regimes, that far surpass what was done for LUX.

An overview of LZ is shown in figure 6.1, showing key components and their location within the water tank that originally housed LUX. The next sections will discuss the design of LZ in more detail.
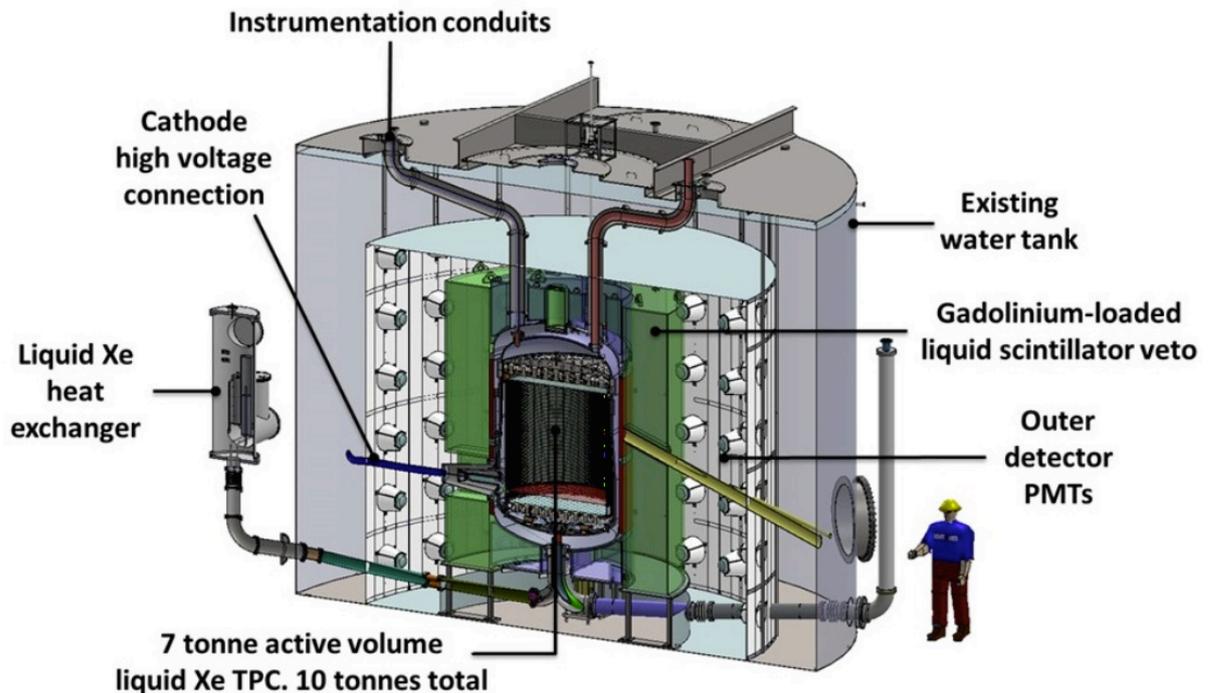
Figure 6.1: Overview of the LZ experiment showing the main components. Key differences to LUX are the size of the TPC, the liquid scintillator veto also placed inside the water tank and the cathode high voltage connection.

6.2 LZ Design

6.2.1 Xenon and TPC

LZ will use a total of 10 tonnes of xenon, with a WIMP target comprising 7 tonnes of active xenon. The TPC will contain three electrodes; a cathode grid at the bottom, a gate grid at the top below the liquid surface, and above the liquid surface an anode grid. As in LUX, the inside of the TPC will be coated with highly reflective PTFE panels to maximise light collection. Within the PTFE, an embedded field cage provides an electric field. Resistors are present in order to grade the field, keeping it in a vertical configuration for drifting electrons. Between the PTFE and the inner cryostat walls there will be a layer of xenon known as the 'skin'. This will be instrumented with PMTs to provide a veto region to be used in combination with the outer detector, and consists of 2 tonnes of xenon. This will be optically decoupled from the main xenon target. The xenon will be purified using a hot zirconium getter to ensure low levels of impurities that will restrict electron drift lifetime through the detector or shorten the attenuation length of scintillation photons in xenon, necessary as purification is essential to collecting maximum signal. Radioactive contaminants that act as an internal background also must be removed, namely ^{85}Kr and ^{39}Ar.

The delivery of the electric field to LZ is very different to LUX; a xenon filled 'umbilical' feedthrough will deliver high voltage to the cathode. The feedthrough cable has a 0.43 inch conductive core, a polyethylene sheath and enters the TPC horizontally. Outside of the TPC, the cable bends to vertical and travels up to the top of the water tank. Below the cathode, the reverse field region (RFR) is a small region of a high electric field designed to protect the bottom PMTs from the drift field. Thus the electric field is divided into three regions; the drift region, the RFR and the electroluminescence region (within the gas). This design will allow field strengths of ~3× greater than in LUX, which will allow a greater electron extraction efficiency. Key parameters of the main detector system are shown in table 6.2.

6.2.2 PMTs

The LZ active xenon will be viewed by two arrays of PMTs, with 241 contained in the bottom array within the liquid, and 253 in the top array within the gas layer. These will be Hammamatsu R11410-22 3-inch PMTs, and are ultra-low background, with ~1000 times less radioactivity than a standard PMT. These were shown to be suitable for LZ, with a quantum efficiency of >30% at 175 nm and <1 mBq/PMT of ^{238}U and ^{232}Th, and were even considered as a replacement in LUX after a screening campaign found them to have considerably less radioactivity than the LUX R8778 PMTs [84], see table 6.3.

Furthermore, LZ will have PMTs to view the outer layer of xenon outside of the TPC known as the skin; 93 1" PMTs at the top of the TPC, 20 2" PMTs at the bottom of the TPC and 18 2" PMTs in the bottom dome. The bottom and dome PMTs will be the same as the ones used for

Table 6.2: Parameters of the xenon and TPC detector system.

Category	Parameter	Value
LXe	TPC active mass	7,000 kg
	Skin mass	2,000 kg
	Total mass	9,600 kg
PMTs	TPC	253 (top) + 241 (bottom)
	Side Skin	93 (top) + 20 (bottom)
	Dome Skin	18
Dimensions	Electroluminesence (gate-anode)	13 mm
	Drift (cathode-gate)	1,456 mm
	RFR (sub-cathode)	137.5 mm
	TPC diameter	1,456 mm
	Field cage thicklness	15 mm
	Skin thickness	40 mm (surface), 80 mm (cathode)
Electric Fields	Electroluminesence field	10.2 kV/cm
	Drift field	0.31 kV/cm (baseline), 0.65 kV/cm (goal)
	Reverse field	2.9 kV/cm (baseline), 5.9 kV/cm (goal)
	Stages	57 (drift), 7 (RFR)
Operation	Pressure	1.8 bar(a), range of 1.6-2.2 bar(a)
	Equilibrium temperature	175.8 K

Table 6.3: Initial screening results of the R8778 PMTs used in LUX and the R11410 PMT candidate for LZ [84]. The LUX PMT shows detection for all radioisotopes, whilst the R11410 PMT only ^{60}Co is not an upper limit.

PMT	Activity (mBq/PMT)				
	^{238}U (^{234m}Pa)	^{238}U(^{226}Ra)	^{232}Th(^{228}Ra)	^{40}K	^{60}Co
R8778	<22	9.5 ± 0.6	2.7 ± 0.3	66 ± 6	2.6 ± 0.2
R11410	<6.0	<0.4	<0.4	<8.3	2.0 ± 0.2

LUX; Hammamatsu R8778. As will be see in chapter 7, the skin is important in rejecting γ-ray backgrounds.

6.2.3 Cryostat

The cryostat is a vessel designed to contain 10,000 kg of LXe at -100°C. There are three main parts - the inner cryostat vessel (ICV), the outer cryostat vessel (OCV) and the cryostat support (CS). The cryostat will be fabricated from commercially pure grade 1 titanium. This titanium was chosen after an extensive screening campaign identified it as the most radiopure of not only the series of samples taken from different manufacturers, but of all reported screened titanium worldwide. Details on the procurement, screening, and selection of this titanium, and its simulated background contribution (which can also be found in chapter 7), will be described in an upcoming publication [144].

The ICV is split into two sections joined by a flange near the top. It comprises of a cylindrical body with one end tapered in order to minimise the LXe passive volume between the TPC and inner cryostat whilst keeping the electric field at the cathode below the requirement of 50 V/cm,

(a) The outer cryostat vessel and cryostat support.   (b) The inner cryostat vessel.

Figure 6.2: CAD model of the cryostat vessels and support. The large port on the front will be for the HV umbilical feedthrough. The two ports at the front and back of the OVC and ICV top domes are for PMT and thermosyphon cabling, whilst the three tubes that appear only on the OCV will be calibration source tubes. There is also a bottom port visible for the bottom array PMT cabling and heat exchange.

and two ellipsoidal heads with a 2:1 (top) and 3:1 (bottom) aspect ratio. The differing ratios mean the bottom needs additional thickness, making it 12 mm thick compared to the rest of the vessel at 8 mm. The top head contains two ports for cabling and heat exchange conduits. A high voltage conduit port sits on the front of the vessel at the largest radius after the tapering, and the bottom head contains a port for heat exchange. The ICV is suspended within the OVC on tie bar assemblies.

The OCV is designed in three sections, a straight cylinder body and 2:1 ellipsoidal head and base, and a thickness of 8 mm. The top head consists of the dome, three ports for calibration source tubes, two conduit ports for thermosyphon and PMT cabling and a large indent for a cylindrical tungsten alloy pig that will contain photoneutron calibration sources. The bottom section has a straight section starting with a flange and a port at the front for the high voltage conduit, a heat exchange and cabling port on the bottom of the ellipsoidal section and three support stands. The cylindrical body has two matching flanges to connect it to the top and bottom sections. The flanges are joined with stainless steel nuts and bolts. The CS will connect to three base plates mounted on studs at the water tank base. The support has been designed to not only support the cryostat but to also resist the expected horizontal force applied in a seismic event.

6.2.4 Outer Detector

The need for an outer detector is driven by the low background rate required for LZ to reach its design sensitivity. A WIMP scatter within the LXe within the region of interest (~ 5 keV$_{nr}$ to 50 keV$_{nr}$) will never be accompanied by an energy deposit in any of the surrounding detector components. A background γ-ray or neutron, however, is likely to deposit energy elsewhere on its way into or out of the detector if it has scattered within the xenon. The outer detector serves two purposes; it will veto primarily neutrons but also γ backgrounds with high efficiency, allowing the enlargement of the fiducial volume, and the second is to measure and characterise the background [145]. This is important in the case of a WIMP signal, which will need exceptional supporting proof of a thorough understanding of backgrounds.

<div align="center">(a) OD tanks assembled   (b) OD tanks exploded</div>
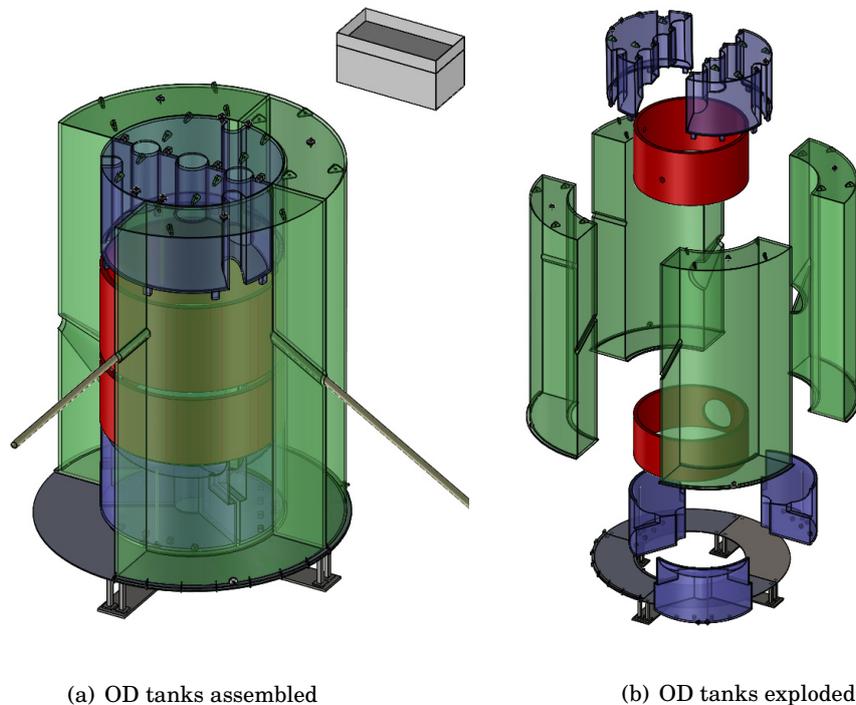
Figure 6.3: Liquid scintillator acrylic tanks. The side tanks are shown in green and the top and bottom tanks in blue. Displacer cylinders are shown in red. At the top in grey is a reservoir for the liquid scintillator.

The LZ outer detector system consists of 10 segmented acrylic tanks containing 21.5 tonnes of gadolinium-loaded linear alkylbenzene (LAB). The acrylic tanks are designed as four side tanks, three bottom tanks and two top tanks, see figure 6.3. The tenth tank is removable in order to allow the use of the YBe calibration source that is placed atop the cryostat.

LAB is a liquid scintillator with good optical transparency (20 m), high light yields, low contamination with radioactive impurities and a high flash point that makes it safe to have underground. It is employed by the SNO+, Daya Bay and RENO neutrino detectors, and its

qualities make it also ideal to use as a veto for dark matter experiments. The LAB is loaded with gadolinium at 0.1% by mass in order to increase the probability of neutron capture. Gadolinium captures a neutron by the processes:

$$^{157}_{64}\text{Gd} +^1_0 n \rightarrow^{158}_{64}\text{Gd} + \gamma(7.9\text{MeV}) \tag{6.1}$$

$$^{155}_{64}\text{Gd} +^1_0 n \rightarrow^{156}_{64}\text{Gd} + \gamma(8.5\text{MeV}) \tag{6.2}$$
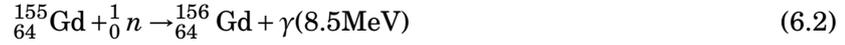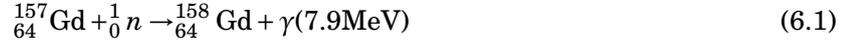
where the number of γ-rays is usually 3 or 4, totaling the given energy together. For thermal neutrons, the capture cross section on ^{157}Gd is 259,000 barns. This is the highest of all stable isotopes; only the unstable isotope ^{135}Xe has a higher cross section at 2×10^6 barns.
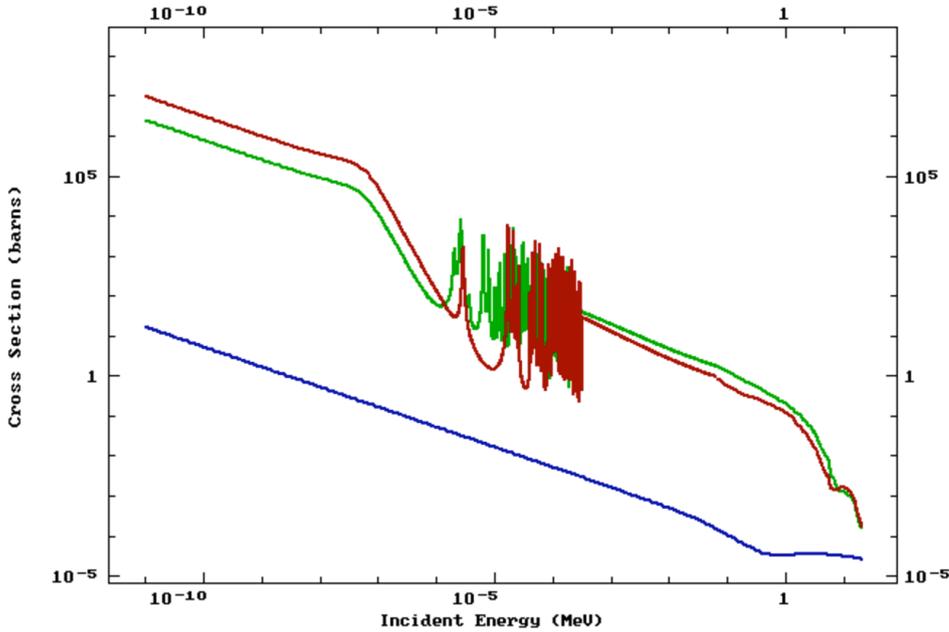
Figure 6.4: Capture cross sections for ^{157}Gd (red), ^{155}Gd (green) and hydrogen (blue) [146].

Most neutrons (90%) will be captured on either ^{157}Gd or ^{155}Gd, with the remaining 10% captured on hydrogen. The capture cross sections and their dependency on neutron energy of the three atoms is shown in figure 6.4. Including gadolinium decreases the capture time of the scintillator from $\sim 200~\mu$s to $\sim 30~\mu$s. The veto efficiency for neutrons must be greater than 90% to meet LZ sensitivity requirements.

γ-rays of a few MeV are able to scatter at small angles in the outer region of the LXe, leave 0.5 - 10 keV of energy and then exit the TPC, making them a dangerous background. The requirement for the outer detector is to veto > 70% of these backgrounds.

120 8-inch Hamamatsu R5912 PMTs instrument the water tank in order to detect the γ-rays released by neutron capture. This means the acrylic must be transparent to photons with wavelengths greater than 300 nm. The PMTs are arranged in a cylindrical array of 20 ladders

with 6 PMTs each on the walls of the water tank. Radioactivity from the water PMTs is prevented from reaching any active components by the water.

6.3 Calibration Systems

LZ will build on the calibration systems that were developed and refined for LUX. However, due to the large size difference between LUX and LZ, there are difficulties with using some of the techniques. Original LZ calibration plans were the use of a dispersed ^{83m}Kr source as in LUX, described in section 5.1.1. However, simulations of xenon flow suggest there will not be sufficient mixing of the much larger xenon volume within the 1.86 hour half life of ^{83m}Kr to ensure a uniform calibration source. Therefore, it is now proposed that LZ will use ^{131m}Xe as a calibration source, which will be present at a rate of 1-10 Hz naturally within the xenon, and decays as:

$$^{131m}\text{Xe} \rightarrow ^{131}\text{Xe} + \gamma(163.9\text{keV}) \tag{6.3}$$

with a half life of 12 days. The isotope is produced using a ^{131}I parent source, which due to its use in the medical industry is known to be common, cheap and radio-clean. There is likely to be some saturation in a single PMT from this relatively high energy γ-ray, but it is expected that with further study of PMT light response the position reconstruction algorithm will still work effectively.

Several neutron sources will be used in LZ. The NR band will be calibrated with an americium-lithium, or AmLi source. AmLi is favoured over the AmBe previously used in LUX as its lower maximum neutron energy of 1.5 MeV results in an enhanced fraction of events at low recoil energy (<10 keV). There will also be photoneutron calibrations; a YBe source is planned, and further R&D is underway on ^{205}Bi and ^{206}Bi sources, which emit γ-rays at 1,764 and 1,719 keV respectively. These can also be combined with beryllium and cause the ejection of 88.5 and 47.5 keV neutrons, which result in low energy NR endpoints of 1.4 and 2.7 keV in xenon.

A D-D generator calibration will be performed as in section 5.2.4 to measure NR light and charge yields. There will be the addition of two operational techniques that were tested in LUX after run 4. The first technique uses a pulsed operation mode, where event timing information is taken from the generator so that an S2-only analysis can be done, allowing event selection to reach lower energies. The second involves directing neutrons away from the detector and at a D_2O reflector, which then backscatters at 135°into the detector. The neutrons have a reduced energy of 355 keV, resulting in a NR spectrum endpoint of 10.5 keV.

A further suggestion currently undergoing research is the idea of evaporating ^{210}Po onto the cathode at fixed radial positions outside of the fiducial volume, which would provide a tool to constrain position reconstruction algorithms. It would need to be ensured that this did not cause increased photon or electron emission from those spots. In LUX, CH_3T's uniform distribution in LUX's TPC allowed tight constraints on the fiducial volume definition; due to the mixing

problem mentioned previously this may be difficult in LZ. The cathode ^{210}Po point-sources may be a solution to this. However, mobility studies of ^{210}Po and daughters must be validated first.

LZ will have a series of low-deadtime calibrations before and during the WIMP search: ^{131m}Xe will be used to measure the electron lifetime and PMT gains and the NR band will be calibrated with AmLi. Furthermore, to calibrate the skin, ^{220}Rn will be used to check for low light collection corners, ^{131m}Xe will be used as a uniform source throughout the skin and ^{57}Co in the calibration tubes allows a scan in the z direction. Finally, the outer detector will be calibrated with a ^{228}Th source, which is an α-decay, and this can be performed early as there will be no need to wait for high liquid purity levels to be reached.

Calibrations that require significant deadtime are intended to take place after the 1000-day WIMP search run, and include the D-D calibrations, the photoneutron sources, and tritiated methane to measure the ER light and charge yield, and band tails. This calibration plan is likely to change and evolve over time as LZ is constructed and the timeline becomes clearer.

6.4 LZ Backgrounds

6.4.1 Sources of Background

LZ will face similar backgrounds to LUX, described in chapter 4. Electron recoils are by far the dominant background in any liquid noble dark matter search, due to the penetration ability of γ-rays and the prevalence of isotopes that emit γ-rays or β-particles. Whilst ER events are obviously not usually candidate WIMP events, there will be some fraction of leakage of ER events into the nuclear recoil band due to fluctuations, light loss, reconstruction etc. The goal for LZ is for a leakage fraction of just 0.5%, but for LUX 0.2% was obtained. In order for LZ to reach its goal sensitivity of $\sim 2 \times 10^{-48}$ cm^2 within three years of data taking, it can only tolerate a maximum of 37 μDRU$_{ee}$ (see the definition of DRU at the beginning of chapter 4).

Neutrino physics plays a large role in LZ backgrounds; the solar neutrino rate is even used to place constraints on other backgrounds. A few hundred ER events are expected from pp, ^{7}Be and ^{13}N solar neutrinos scattering from electrons. Additionally, LZ is expected to detect 7 ± 3 neutrino-nucleon coherent scattering events, a standard model process described in section 1.4 that has never been previously observed. However, these events are expected at very low energies and have little impact on the WIMP search across the vast range of WIMP masses. These events are produced by ^{8}B neutrinos from the sun, which have the highest scattering cross section of astrophysical neutrinos in xenon. A further (sub-dominant) ER background below 20 keV$_{ee}$ is presented by $2\nu\beta\beta$-decay of ^{136}Xe, and this is expected to generate tens of ER counts throughout the LZ lifetime.

The key contributors to ER backgrounds from material radioactivity are the isotopes within the ^{238}U and ^{232}Th decay chains, specifically from the radon sub-chains, and the single decays of ^{40}K and ^{60}Co. As was detailed in section 4.2.3, radon is a significant and dangerous source

of background because of its emanation into the active volume. Naked and semi-naked β-decay of radon daughters is a concern due to vetoing ability, and as in LUX, ^{214}Pb presents the most difficulty. Mobility of ^{210}Bi from the walls is also a concern (as discussed in section 4.5). For NR backgrounds from material radioactivity, neutrons are produced in (α,n) reactions in various detector materials. These are predicted to be emitted at rates that vary by component between 0–124 neutrons/year, with the exception of the outer detector PMTs (3,308 n/year, but these are so far from the TPC the high rate is allowable - this has been shown in simulation, see chapter 7). Dominant internal components for neutron emission are the PMTs and the cryostat vessel, the latter simply because of its large mass. Material radioactivity has been constrained so that it must only contribute an event rate that is <10% of the solar pp neutrino scattering rate, and a maximum of 0.2 nuclear recoils (assuming 50% signal acceptance).

Intrinsic backgrounds within the LXe are expected from ^{85}Kr and ^{39}Ar, which decay as:

$$^{85}_{36}\text{Kr} \rightarrow {}^{85}_{37}\text{Rb} + e^- + \bar{\nu}_e \qquad (Q = 687 \text{ keV}, t_{1/2} = 10.8 \text{ years}) \qquad (6.4)$$

$$^{39}_{18}\text{Ar} \rightarrow {}^{38}_{19}\text{K} + e^- + \bar{\nu}_e \qquad (Q = 565 \text{ keV}, t_{1/2} = 269 \text{ years}) \qquad (6.5)$$
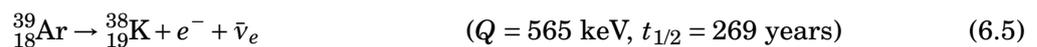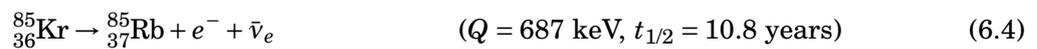
For ^{85}Kr, the decay spectrum is well understood and events can be measured using $\beta - \gamma$ coincidences, whilst ^{39}Ar is a naked β-decay and so more dangerous, albeit purification of the xenon should remove both isotopes. The contribution from krypton and ^{222}Rn combined is also constrained to 10% of the solar pp rate, limiting ^{222}Rn to <0.67 mBq [1] and the krypton concentration to <0.002 ppt (g/g). For ^{39}Ar, the constraint is <10% of the rate of ^{85}Kr, or <2.6 μBq.

LZ will also have similar cosmogenic backgrounds to LUX; ^{127}Xe ($t_{1/2}$=36.4 days), ^{129m}Xe ($t_{1/2}$=8.9 days) and ^{131m}Xe ($t_{1/2}$=11.9 days) will be present in the xenon when it is first taken underground, but will start to decay away once there. The xenon stockpile is expected to be taken underground over a year before operations begin and so it is expected that there will be not be measurable rates of these isotopes once LZ is taking data. However, as ^{127}Xe in particular can generate energy deposits within the WIMP search region of interest (5.2 keV, ≤1.2 keV), it is desirable for the xenon to be shielded from thermal neutrons even whilst stored underground to mitigate any risk of generating this isotope. ^{60}Co can be produced in copper components and will produce γ-rays, and ^{46}Sc is known to be present in the titanium chosen for fabrication of the cryostat and PMT array structures at a level of 2.0 mBq/kg. Cosmic muons are easily tagged by their Cherenkov light within the water tank, but muon-induced neutrons through spallation or photonuclear interacts are a concern, and the flux is calculated to be 0.54×10^{-9} n/cm^2/s, with half above 10 MeV, and 10% above 100 MeV, a rate much lower than the material radioactivity neutron flux. Neutrons are also produced at a rate of the same order of magnitude within the water, but the water attenuates these to a flux of only 0.2 n/day.

---

[1]This is the goal, but for backgrounds and sensitivity projections, the baseline requirement of <2μBq/kg is used.

6.4.2 Screening and Cleanliness

The background rates in LZ must be unprecedentedly low and accurately modelled in order for it to reach its design sensitivity goal, whether for limits or discovery. This will be achieved by an extensive screening campaign for all detector materials, Monte Carlo simulations to estimate the background rate, thus providing feedback on the suitability of screened components, and finally cleanliness protocols for manufacture, transportation and commissioning.

Material screening is the principle method of controlling radioactive backgrounds in LZ. Candidate materials are screened for their contamination with radioisotopes, and only selected for use if the results are sufficiently low (which often requires feedback from simulations, detailed in the next chapter). With so many different components being manufactured in many different places, there must be clear protocols set out for all sub-systems on the proper handling, cleaning and shipping of components, especially those close to the TPC. There must be control measures in place during the construction of components, especially those close to the TPC, to keep contamination to a minimum. Examples include working in a cleanroom and cleaning materials only with products that are known to be suitable for a low-background experiment such as deionised water, cleaning surfaces of dust and limiting exposure to radon plate-out. The LZ cleanliness team has determined which materials are suitable for packaging and has provided this information as well as information on cleaning protocols to the rest of the collaboration.

There is also a process in place to use 'witness plates' and 'coupons' which are kept alongside materials as they are produced and transported. These can then be assayed once the items reach their destination to check for cleanliness.

The requirement for dust contamination on surfaces is <500 ng/cm^2, although we work to a goal of <5 ng/cm^2 and for radon plate-out <0.5 mBq/m^2. Predictions can be made on dust and plate-out during construction in a cleanroom using analytical calculators. These use models that predict surface fallout and plate-out using airborne contaminant concentration (ASTM-DELFT for dust, Jacobi for radon), and require inputs of cleanroom class (particles per volume), the size of the cleanroom, the filters in use, the air exchange rate, etc. Dust contributes half of the predicted amount of radon emanation in LZ, so keeping it to a minimum is very important.

6.4.2.1 γ Spectroscopy

γ-ray spectroscopy involves the use of High Purity Germanium detectors (HPGe) in measuring γ-ray spectra of candidate construction material. Such spectra can be used to infer the concentrations of radio-isotopes based on which energy peaks appear. LZ has the use of 11 HPGe detectors located both above and below ground, with a range of crystals: n-type, p-type and broad energy (BEGe). Using several different types allows precise measurements over a wider range of energies than could be achieved with one type alone. Material samples are placed close to the Ge crystal inside a castle made of low radioactivity lead and copper, which helps to shield the detectors from background radiation. The chambers are also continuously flushed with nitrogen to remove

any radon. Placing the detectors underground reduces background even more, so most of LZ's Ge detectors are operated at SURF in the Black Hills State University Underground Campus (BHUC) and at the U.K. Boulby Underground Laboratory within the Boulby Germanium Suite (BUGS). The detectors that remain on the surface are generally used for pre-screening to identify anything unsuitably hot, so only viable candidates are taken underground for more sensitive screening.

Table 6.4: HPGe detectors in use for LZ material screening, showing the depth of their location in meters water equivalent (mwe), the crystal type and mass and their sensitivities to the uranium and thorium late chains. Relative efficiencies are measured in comparison to the detection efficiency of a 3×3-inch NaI crystal for 1.33 meV γ-rays from a ^{60}Co source 25cm from the detector face.

Detector	Site	Depth (mwe)	Crystal	Crystal Mass (kg)	Relative Efficiency	Sensitivity	
						U(mBq/kg)	Th(mBq/kg)
Chaloner	Boulby	2805	BEGe	0.8	48%	0.6	0.2
Ge-II	Alabama	0	p-type	1.4	60%	4.0	1.2
Ge-III	Alabama	0	p-type	2.2	100%	4.0	1.2
Lunehead	Boulby	2805	p-type	2.0	92%	0.7	0.2
Lumpsey	Boulby	2805	Well	1.5	80%	0.4	0.3
Maeve	SURF	4300	p-type	2.1	85%	0.1	0.1
Merlin	LBNL	180	n-type	2.3	115%	6.0	8.0
Mordred	SURF	4300	n-type	1.2	60%	0.7	0.7
Morgan	SURF	4300	p-type	2.1	85%	0.2	0.2
SOLO	SURF	2200	p-type	0.6	30%	0.5	0.2
Wilton	Boulby	2805	BEGe	0.4	18%	7.0	4.0

6.4.2.2 ICP-MS

Inductively Coupled Plasma Mass Spectroscopy, or ICP-MS, is a technique that uses inductively coupled plasma to ionise a sample for screening; the ionised sample is then passed into a mass spectrometer which separates them by their mass to charge ratio using magnets. ICP-MS measures individual elements at sensitivities up to 1 in 10^{15} , or 1 part per quadrillion (ppq).

UCL operates an ICP-MS facility in an ISO class 6 cleanroom dedicated to LZ. The instrument is an Agilent 7900 and has sensitivity to U and Th below 10^{-12} g/g. ICP-MS has the advantage of being able to directly measure ^{238}U and ^{232}Th; without it, we can only extrapolate from γ-ray spectroscopy measurements of the late chain activity, and make assumptions about secular equilibrium. The results from another ICP-MS facility at the University of Alabama have been used to validate γ-screening results for titanium samples for the LZ cryostat, see section 6.4.3

6.4.2.3 Radon Emanation

Any components that may contribute to radon entering the active xenon volume must be screened in a radon emanation chamber. This involves placing a sample within a sealed chamber, waiting for it to outgas, and measuring the resulting radon activity using a silicon pin diode detector. The positively charged α-particles produced by radon decay are attracted electrostatically to the PIN

diode detectors and appear as a clear signal. LZ has such systems at Maryland and SDSM&T, and also utilises the SuperNEMO radon emanation system at MSSL [147]. Another method used by the University of Alambama is a liquid scintillator viewed by a PMT; radon-bearing gas is passed through and the Bi-Po coincidence is detected. ^{214}Bi decays with a half life of 20 minutes into ^{214}Po, which decays rapidly with a half life of 124 ms, so the two can be detected together (for example in LUX, they would appear within one event window).

Radon emanation from materials accounts for half of the total radon counts, with the other half coming from dust. We expect less radon in LZ than indicated by assays, since they are performed at room temperature. At LXe temperatures, some suppression of Rn diffusion is expected.

6.4.2.4 NAA

Neutron Activation Analysis (NAA) has also been used for LZ, by utilisation of the 5.5 MW$_{th}$ MIT Reactor II to activate samples with neutrons and measure the subsequent γ-rays with HPGe detectors. Samples can be irradiated for minutes - days, allowing high analysis sensitivity. Activation products of interest are elements such as ^{42}K, ^{233}Pa and ^{239}Np. A sensitivity for uranium and thorium of 10^{-12} g/g has been achieved with this technique; in particular this has been used to meet requirements for assaying PTFE for LZ.

6.4.3 Radiopure Titanium

For certain key materials, screening campaigns were undertaken to identify the most suitable materials. One example, important due to the mass and the proximity of the material to the LXe, is that for the cryostat; this was a highly successful campaign where the world's most radiopure titanium was selected to construct the LZ cryostat and an assessment of the expected background was performed.

Titanium is an attractive alternative to stainless steel and copper for cryostat vessels. It has a high strength-to-weight ratio, is resistant to corrosion, and is already popular in industry with aerospace, metal finishing, oil refining and medicine. It has a lower atomic mass and therefore lower attenuation of radiation than stainless steel, should in principle contain very little ^{60}Co, and also benefits from little cosmogenic activation. LUX successfully manufactured and deployed a cryostat constructed from low-activity titanium to contain 350 kg of LXe, and so the goal was to procure similarly low-radioactivity titanium for LZ. Many other experiments have failed to do so, however.

To produce metallic titanium, TiCl$_4$ (produced from mined ore via the Kroll process) is reduced by liquid magnesium or sodium. The enclosing vessel is filled only with argon to avoid oxygen or nitrogen contamination. This produces a porous titanium sponge, which is purified by leaching or heated vacuum distillation. There are typically no controls in place at this stage to limit contamination from surfaces or fluids in contact with the sponge, and thus variability in

radioactivity of sponge products is expected. After crushing and pressing, the sponge is melted in an arc furnace, and mixed with any alloy additions required for the grade of titanium being produced, as well as scrap metal, before being melted in the arc furnace again, and cooled to form an ingot. The use of alloys and scrap metal, often without traceability of source material, constitutes the highest risk for introduction of radioactivity into the bulk material. At this stage the radiopurity may be greatly affected even if the purity has been preserved from the $TiCl_4$ through to the sponge [144].

Refinement of titanium can be undertaken using either vacuum arc remelting (VAR) or electron beam cold hearth (EBCH) techniques. VAR is used for most commercial titanium and involves melting the metal under a vacuum in a crucible. A current is passed through a titanium ingot electrode and across a gap to another piece of the metal, creating an arc and beginning a continuous melting process. As the electrode melts, it must be lowered towards the bottom of the vessel, allowing control over the solidification rate, which affects the micro-structure of the titanium. The vacuum conditions allow any remaining gas contamination (e.g. N_2, O_2, H_2) to escape the titanium into the vacuum chamber. Other impurities such as carbon, sulphur and magnesium have a high vapour pressure and will also be lowered in concentration. The disadvantages of VAR are that it is a slow process and nitrogen-rich particles have a tendency to sink to the bottom where the temperature is too low for dissolution to take place effectively.

Electron beam remelting uses impinging electron beams to heat and melt metals in a water-cooled copper crucible under a vacuum. In EBCH, cold hearth refers to a watercooled copper hearth in which the feedstock is drip melted, overflowing into a withdrawal mold. EBCH holds some advantages for manufacturers: it is a faster process as it involves less steps than traditional refinement implying fewer opportunities for further contamination of radiopurity, its high yield, its flexibility in melting feedstock of any geometry and finally the possibility of producing ingots and slabs of a wide range of cross-sections. A final and key advantage for applications such as low-background experiments is that EBCH is expected to exceed other refinement methods in removing high density contaminants [148].

For the LZ campaign, 22 samples of titanium of varying grades was sampled in partnership with several manufacturers and at various stages of the production cycle beyond radiopure $TiCl_4$ (including sponge, slabs, and sheet), using either VAR or EBCH remelting, and with varying scrap concentrations. 13 stainless steel samples were also screened as an alternative. Following the assays, the highest and most reproducible radiopurity was found in ASTM Grade 1 titanium that contained 0% scrap and had been refined using EBCH technology. Of these, the lowest radioactivity was observed in material from TIMET Heat Number (HN) 3469, a single 15,000 kg slab of titanium, produced by TIMET at its Morgantown (Pennsylvania) mill.

A sample from TIMET HN3469, denoted HN3469-T, was screened at the Berkeley Low Background Facility using MAEVE in May of 2015. This consisted of 10.1 kg of plates selected from the top portion of the single slab. A second sample taken from the middle (HN3469-M), was

acquired and assayed in September of 2015 to confirm the uniform distribution of contamination. The samples were analyzed using a GEANT4 [89] model of the MAEVE detector (as was used in section 5.2.2) with an exact sample geometry of the titanium placed around the detector on four sides and the top to simulate detection efficiency of γ-rays emitted from the titanium. The radioactivities of both samples of this titanium stock were found to be consistent [144].

In both of the samples, each counted for approximately three weeks, the early uranium chain was non-detectable or barely detectable within the limited abilities of HPGe detectors to assay this portion of the decay chain via γ-ray spectroscopy. The late portion of the chain (at ^{226}Ra and below) however, is quite accessible via γ-ray spectroscopy due to both the branching ratio and detection efficiency for its γ-rays, and registered no detectable activity above background down to the few ppt level. The late uranium value is based upon the 609 keV peak from ^{214}Bi and is consistent when compared to upper limits from other useful peaks in the late uranium chain, such as 295, 352 (^{214}Pb) and 1764 keV (^{214}Bi). Both samples had detectable levels of the thorium series in both the early and late portions of the chain, consistent with secular equilibrium in both samples. The late series measurement is based upon the 238 keV gamma-ray from ^{212}Pb, which is the strongest peak given the product of the detection efficiency of that gamma-ray line and its branching ratio in the thorium chain. The assays are summarized in Table 6.5.

Table 6.5: Results from gamma-ray spectroscopy of TIMET titanium sample HN3469 for both top and middle samples. All limits are 1σ upper limits and uncertainties are statistical only. A systematic uncertainty of up to 10% should also be assumed. The results are shown as activities in the ^{238}U, ^{232}Th, and ^{40}K chains—all in units of mBq/kg. Results were obtained with the MAEVE detector and confirmed with the CHALONER detector at Boulby and ICP-MS assay [144].

	Top	Middle
Date	5-2015	9-2015
Sample mass	10.07 kg	8.58 kg
Livetime	23.9 days	20.8 days
^{238}**U**$_e$	*<1.6*	2.90 ± 0.15
^{238}**U**$_l$	*<0.09*	*<0.10*
^{232}**Th**$_e$	0.28 ± 0.03	*<0.20*
^{232}**Th**$_l$	0.23 ± 0.02	0.25 ± 0.02
^{40}**K**	*<0.54*	*<0.68*
^{60}**Co**	*<0.02*	*<0.03*
^{46}**Sc**	2.0 ± 0.1	2.7 ± 0.1

In terms of cosmogenic activation there were several isotopes of scandium present, most of which are the result of cosmic ray-induced reactions with the five stable isotopes of titanium. Detected in the sample was ^{46}Sc (889, 1121 keV, $t_{1/2}$~84 days); as well as small amounts of ^{47}Sc (159 keV, $t_{1/2}$~3 days), ^{48}Sc (984, 1038, 1312 keV, $t_{1/2}$~44 hours), and ^{44,44m}Sc (271, 1157 keV, $t_{1/2}$~59 hours and ~4 hours, the metastable state being the longer-lived). Screening results for ^{46}Sc are also listed in Table 6.5, whilst the ^{47}Sc, ^{48}Sc, and ^{44,44m}Sc activities were not listed as their short half lives mean they essentially disappear over the course of the measurement. The reported value for ^{46}Sc was decay corrected to the start of counting for each of the samples. All

limits are 1σ upper limits and uncertainties are statistical only, and systematic uncertainties are estimated to be up to 10%. The cosmogenic production of ^{46}Sc can be mitigated by moving components underground as soon as possible after manufacture. This campaign was considered a success, and the titanium from the HN3469 slab has been procured to not only manufacture the 2,292 kg cryostat but also the field-shaping rings (260 kg) and PMT support structures (104 kg) within the inner detector. Results from Monte Carlo simulations of the background from this titanium cryostat can be found in the next chapter.

Similar screening campaigns are being performed for every single component in LZ, to ensure sufficiently low backgrounds and full understanding of levels of radioactivity within all detector systems.

S imulations are an essential tool in physics; not only do they allow the testing of theories and models by comparison to data, for an experiment like LZ they enable the study of a physical system in high detail before its components have begun production in real life. This allows for design optimisation and a greater understanding of the challenges we may face when building and running LZ. In this chapter, I will detail my role in the building of an accurate and full model of LZ for Monte Carlo simulations, developing and validating an analysis code for simulation output and producing a reliable background model using inputs from radioassay techniques. Furthermore, I will describe my identification of a major flaw in a uranium chain generator used for background simulations, my work on simulations of the outer detector, which led to a rethink of its design, and finally my implementation of a spontaneous fission generator, the results of which led to the removal of spontaneous fission neutrons from the background model, again significantly impacting choice of construction materials and acceptance criteria.

7.1 Monte Carlo Simulations

7.1.1 LZSim

The pre-existing simulation framework LUXSim (see section 2.6) was used to create LZSim. A new detector model for LZ was built using GEANT4 geometry classes with dimensions informed by the CAD model. The model contains all components relevant to backgrounds and realistic particle propagation, including the cryostat vessels, the PMT arrays, the grids, PTFE, field rings, conduits, the outer detector scintillator tanks and the liquid and gaseous xenon volumes. There has been extensive work done on ensuring the physical accuracy and completeness of the LZSim geometry. Several validation tests were developed in order to check for any bugs, incorrect geometry or

volume overlaps within the model, and these are described in section 7.1.4.

Geometry with GEANT4 involves coding simple individual geometrical shapes such as cuboids, spheres and cylinders, and augmenting them to more realistic shapes by adding and subtracting new volumes. Once the shape is complete, it is assigned a material and a position and rotation within a 'mother volume', thus becoming a daughter volume. For example, the outer detector acrylic tanks are daughter to the water in the water tank, and the LAB scintillator itself is daughter to the acrylic. With a complex detector such as LZ, geometry modifications are not trivial, as there are thousands of different volumes which may be affected by the placement of a new one. Several contributions were made to the geometry of the LZ detector model; the outer and inner cryostat volumes were updated to match the latest CAD models. Previous to these modifications, both vessels were simple cylinders with domed ends; there were no flanges or ports and the dimensions were outdated. CAD designs were used to modify the geometry to match the correct dimensions, then new volumes were created to add in the HV umbilical port, the calibration tube ports, the cabling ports and the flanges. In addition, the field rings within the PTFE were placed; a process that involved creating two different ring shapes for the reverse field region and the drift region, and duplicating these rings with the correct spacing to cover the whole TPC.

LZSim is under version control with GitLab. Users given developer access can create a new branch of the code to edit, but they must submit a merge request to be approved by an administrator in order to change the master code. There are tagged releases of the master branch after major updates, with the ability to switch instantly back to an earlier tag if older results need checking. This system has been found to be much more effective for collaborative work than the Subversion system used for LUXSim. The output files of LZSim are converted to ROOT files using a converter developed for LUXSim. The information saved depends on the commands issued before running the simulation, see section 2.6.2 for more details. If a particle passes through a volume that is given a record level greater than 0, a new entry is created in the output file. Figure 7.1 shows an example of the information contained in the ROOT file of an LZSim event.

7.1.2   Data format for Background Simulations

In addition to assessing detector design, one of the main uses of LZSim is for analysing the potential backgrounds that LZ will face. The need for running high statistics simulations for many different backfround sources from many different components poses a problem for file storage. The .root files produced by LZSim can become very large, especially when higher record levels and more volumes are used, therefore we developed a reduced file format for background simulation analysis. The first of these formats is known as the "Reduced Tree" format and only contains information for four key volumes important for backgrounds; the main liquid xenon (within the drift field), the reduced field region (as events here will only create an S1), the skin veto volume, and the outer detector scintillator volume.

```
======> EVENT:1
 iPrimaryParNum  = 1
 cPrimaryParName = neutron&
 fPrimaryParEnergy_keV = 3250
 fPrimaryParTime_ns = 4.97827e+07
 fPrimaryParPosX_mm = -270.737
 fPrimaryParPosY_mm = -830.957
 fPrimaryParPosZ_mm = 1990.69
 fPrimaryParDirX = -0.467151
 fPrimaryParDirY = 0.296089
 fPrimaryParDirZ = -0.833128
 iRecordSize     = 6
 EvtN            = 1
 V_name          = 367
 cVolumeName     = OuterTitaniumVessel
 iStepNum        = 1,
                   9, 10, 1, 2, 3
 iParticleID     = 2112,
                   22, 22, 11, 11, 2112
 cParticleName   = neutron&gamma&gamma&e-&e-&neutron&
 iTrackID        = 1,
                   4, 4, 6, 6, 2
 iParentID       = 0,
                   1, 1, 4, 4, 1
 cCreatorProcess = primary&NeutronInelastic&NeutronInelastic&compt&compt&NeutronInelastic&
 cStepProcess    = Transportation&compt&Transportation&eIoni&S1&Transportation&
 fKEnergy_keV    = 3250,
                   752.684, 624.098, 128.517, 0, 2265.8
 fOptPhoWaveLength_nm = 0,
                   0, 0, 0, 0, 0
 fDirectionX     = -0.467151,
                   0.049647, 0.529334, -0.764064, -0.764064, -0.75995
 fDirectionY     = 0.296089,
                   0.709241, 0.48959, 0.594122, 0.594122, -0.646459
 fDirectionZ     = -0.833128,
                   0.703216, 0.692898, 0.251447, 0.251447, -0.0675783
 fEnergyDep_keV  = 0,
                   0.06861, 0, 128.517, 0, 0
 fPositionX_cm   = -27.4747,
                   -35.6018, -35.0163, -35.603, -35.603, -50.9704
 fPositionY_cm   = -82.8416,
                   67.2056, 67.7471, 67.2065, 67.2065, -76.8301
 fPositionZ_cm   = 198.353,
                   210.55, 211.316, 210.55, 210.55, 166.588
 fStepTime       = 4.97827e+07,
                   4.97827e+07, 4.97827e+07, 4.97827e+07, 4.97827e+07, 4.97827e+07
 TotEDep         = 128.585
 TotOptNum       = 0
 TotThermElecNum = 0
```

Figure 7.1: Example output of the information contained in one event from LZSim. This includes primary particle type, position, energy and direction, volume name (the outer cryostat vessel, called OuterTitaniumVessel), the names of each particle created in the event and their respective creator processes, step processes, positions, directions, energies and energy depositions. Here, a neutron creates two γ-rays from inelastic scattering, one of which creates an electron through Compton scattering, which goes on to create another via ionisation. The other γ-ray and the original neutron leave the volume with the step process transportation.

A later decision was made to produce a further reduced tree known as the 'Background Analysis Tree', used to produce final background counts and histograms. Table 7.1 demonstrates the differences in data contained in each of the three types of output file. The key difference between the Analysis Tree and the Background Analysis Tree is the application of the background selection cuts (detailed in section 7.1.3), this allows for quick plotting of results and calculations of cut efficiencies.

Table 7.1: Information stored in the three types of ROOT files used in LZ background simulations.

Type	.root files	Analysis Tree	Background Analysis Tree
Volumes	volumes with record level > 0	LXe, RFR, OD, Skin	LXe, RFR, OD, Skin
Energy deposits	total if record level 1, per step if record level 2	total ER (keV$_{ee}$) and total NR (keV$_{nr}$)	combined ER and NR total (keV$_{ee}$ and keV$_{nr}$)
S1 & S2	-	calculated by NEST (phd)	calculated by NEST (phd)
Position	x,y,z - LZSim coordinates (mm)	x,y,z,r - energy weighted mean (cm)	x,y,z,r - energy weighted mean (cm)
$\sigma_{r,z}$	-	energy weighted variance (cm)	energy weighted variance (cm)
Time	LZSim time (ns)	time of first hit in volume (ns)	time of first hit in volume (ns)
Cut Booleans	-	-	**ROI**, **SS**, **Skin**, **OD**, **FV**

7.1.3 Background Selection Cuts

Analysis of simulation results involve determining how many events constitute a background. Such events must satisfy the following, which imitate the cuts that will be applied to real WIMP search data in LZ:

1. The total energy deposited in the LXe is within the WIMP search region of interest (**ROI**)

2. the particle was reconstructed to have scattered once within the LXe, within the expected detector resolution (**SS**)

3. the event was not vetoed by the xenon skin veto (**Skin**)

4. the event was not vetoed by the liquid scintillator outer detector (**OD**)

5. the scatter location lies within the fiducial volume (**FV**)

Abbreviations given in brackets will be used to refer to these cuts from now on, and will be stated in bold to distinguish them from other abbreviations, for example OD refers to the outer detector in general, **OD** refers to the analysis cut. Note that 'passing' a veto cut means the event remains unvetoed, so that passing all of the cuts renders the event as a final background count.

In the energy deposit-only simulations initially performed, the region of interest for WIMP search is considered to be 1–6 keV$_{ee}$ or 6–30 keV$_{nr}$, and so passing **ROI** required the total energy to be within those bounds. The ER energy is calculated by summing the deposited energy of any photons or electrons in the event, whilst for the NR energy neutrons, protons and ions are used. The total energy is calculated in both keV$_{ee}$ and keV$_{nr}$ by scaling the energy of the other type, i.e.:

$$E_R(\text{keV}_{ee}) = E_{ER} + \frac{1.5}{6.0}E_{NR} \tag{7.1}$$

$$E_R(\text{keV}_{nr}) = E_{NR} + \frac{6.0}{1.5}E_{ER} \tag{7.2}$$

The energy scale used depends on the type of background being simulated; for γ and β simulations keV$_{ee}$ is used and for neutron simulations keV$_{nr}$ is used. However, in order to more closely replicate real life cuts, NEST (see section 2.6.5) was implemented into the LZSim framework so it could be used during analysis. NEST takes the position and energy of individual events and returns S1 and S2 sizes in detected photons. NEST has been well validated using the LUX data, where it replicated S1 and S2 signals very well at energies in the WIMP search ROI. The cut applied in S1 is 0–20 phd, and there is also a S2 lower threshold of 450 phd.

Furthermore, in order to boost statistics for electron recoils, where very few of the initially generated γ-rays make it to the fiducial volume, the energy window is expanded to 1–100 keV$_{ee}$, a flat spectrum is assumed and the number of events is scaled down by the ratio of the two energy windows. This greatly reduces the occurrence of simulations where no events remain after all cuts. This was trivial for energy deposit simulations, but for S1/S2 analysis simulations, the energy range corresponding to 0–20 phd must be determined. This was chosen as 1–7.81 keV$_{ee}$; the upper limit was chosen as the mean of the distribution of total energy in keV$_{nr}$ for S1s from 19.9 to 20.1 phd as calculated by NEST, and the lower 1 keV$_{ee}$ cut off is based on the cut-off in the light yield in NEST. Thus, the ER scaling factor for the S1/S2 analysis is $(100 - 1)/(7.81 - 1) =$ 14.537.

For the single scatter (**SS**) cut, we must define the energy weighted mean position:

$$\langle r_E \rangle = \frac{\sum_i (E_i r_i)}{\sum_i E_i} \quad \langle z_E \rangle = \frac{\sum_i (E_i z_i)}{\sum_i E_i} \tag{7.3}$$

where r_i and z_i are the positions hits in either the radial and the z direction and E_i are the energy deposits for each hit. The energy weighted variance is then calculated as:

$$\sigma_r = \sqrt{\frac{\sum_i E_i (r_i - \langle r_E \rangle)^2 \times \sum_i E_i}{(\sum_i E_i)^2 - \sum_i (E_i^2)}} \quad \sigma_z = \sqrt{\frac{\sum_i E_i (z_i - \langle z_E \rangle)^2 \times \sum_i E_i}{(\sum_i E_i)^2 - \sum_i (E_i^2)}} \tag{7.4}$$

The event is then determined to have passed the single scatter cut if $\sigma_r < 3.0$ cm and $\sigma_z < 0.2$ cm. This is intended to imitate a position reconstruction resolution of 3 cm in r and 0.2 cm in z. Position resolution in LZ is expected to be < 1 cm radially (but varies with energy and the layout of the PMT array) so this is conservative.

The event is vetoed by the skin (**Skin**) if the total energy deposited within the skin is greater than 100 keV or the S1 size is greater than 3 phd if S1 cuts are being used, and is deposited within 800 μs of the first hit in the LXe. For the outer detector (**OD**), the energy threshold is 200 keV, and it must be deposited within 500 μs. For future use, we define:

$$\Delta t_{OD} = |t_{LXe}[0] - t_{OD}[0]| \tag{7.5}$$

where [0] indicates that these values are the first energy deposits in those volumes. See section 7.3 for how these values affect the vetoing efficiency.

Finally, in current simulations, the fiducial volume (**FV**) defined as $r < 68.8$ cm and $1.5 < z < 132.1$ cm.

7.1.4 Testing and Validation of LZSim

The LZSim detector geometry had originally been updated only sporadically and was missing a lot of important elements. In October and November of 2015 the LZ Simulations group conducted a geometry overhaul to update the geometry sufficiently to produce reliable new numbers for the Technical Design Report (TDR). This involved several people working individually on different aspects of the detector, with code merged via Github. In order to test any new code for obivous failure modes, every change was required to pass a test macro, LZQuickCheckNoVis in table 7.2. The other macros listed are used for more detailed checks before any new tagged version of LZSim is released.

Table 7.2: Validation macros used to test LZSim for problems. These are usually used before any change is fully merged into the master branch.

Name	Purpose	Details
LZQuickCheckNoVis	Checks simulation runs successfully	Releases 1000 30 keV electrons from the centre of the detector
LZDetectorDimensionsCheck	Tests correctness of key dimensions	Fires geantinos along several important detector axes/components with tracking output
LZRecursiveOverlaps	Checks for overlapping volumes	Issues the GEANT4 command /geometry/test/recursive_test
LXeSkin_Photons	Tests optical physics and properties	Releases 100,000 LXe scintillation photons within the skin

The GEANT4 geometry tests are able to detect overlap of physical volumes, and running a recursive test allows GEANT4 to extend the test to all daughter volumes; this is necessary as overlaps can cause errors in particle propagation or cause a simulation to exit with a segmentation fault. Firing 'geantinos' with tracking output on allows dimensions of each volume to be checked, as these are a special particle that undergoes no interactions and will travel continuously in whatever direction they are fired.

A series of 'sanity checks' were performed using simulations that recorded particle tracking for every single volume, for both neutrons and γ-rays. Particles were fired radially through the detector, and then their survival properties analysed to compare with predictions to check for incorrect materials, densities and thicknesses within the LZSim model. Neutron cross sections were obtained from the Group Neutron Data Library (GNDL) [149] and γ-attenuation coefficients from the National Institute of Standards and Technology (NIST) [121]. Simulations were run with both the beam of neutrons starting outside the outer cryostat vessel and traveling inwards, and starting in the centre of the TPC and traveling outwards. Several key layers of the detector were selected: the outer vessel, inner vessel, skin, PTFE, and several slices through the xenon,and the cumulative probability of a particle in each direction reaching that layer after passing through all the preceding ones, P(R), was calculated using the cross section and attenuation data. The number of neutrons/γ-rays that had survived that far in the simulation (N_{sim}) were counted, and compared to the number predicted using P(R), N_{exp}. This was calculated from multiplying the

number of generated particles by P(R), which with the cross sections/attenuation coefficiencts (N_{sim}). The layers, their radius in the simulation, P(R) and the ratio of simulation to data is shown in table 7.3. The prediction relied on assumptions of no energy loss, a linear path, and also ignored inelastic scatters as this changes a neutron's ID in the simulation. 84% of the ratios are within 30% of the prediction, which for such a simple approximation is good agreement. The worst agreement was seen for neutrons traveling the longest distance, where N_{sim} is less than 10% of N_{exp}, but this is expected as most neutrons would have scattered and lost energy by this point.

Table 7.3: Table showing cumulative survival probabilities P(R) for several layers within the detector model in LZSim, and the ratio of the number of surviving particles in simulation, (N_{sim}) to the number expected using P(R) (N_{exp}). Numbers in bold are within 10% of the prediction. "In" refers to beams of neutrons fired radially into the detector from outside the outer vessel and "out" is the opposite; a beam of neutrons fired from the center of the TPC radially outwards. *For these, no particles were predicted to survive, and none were seen in simulation.

| Layer | R (cm) | Neutrons | | | | γ-rays | | | |
| | | P(R) | | N_{sim}/N_{exp} | | P(R) | | N_{sim}/N_{exp} | |
		In	Out	In	Out	In	Out	In	Out
Outer vessel	91.6	0.8807	0.0010	**1.10**	0.03	0.8086	2×10^{-6}	**1.00**	**1.00***
Inner vessel	80.7	0.7756	0.0011	1.17	0.62	0.6539	3×10^{-6}	**0.98**	1.11
Xenon Skin	79.7	0.4711	0.0012	1.62	**0.96**	0.2579	4×10^{-6}	**1.04**	**0.90**
PTFE	72.8	0.4499	0.0020	1.38	0.89	0.2457	9×10^{-6}	0.80	0.71
Fiducial edge	68.8	0.3145	0.0021	1.24	1.11	0.1259	1×10^{-5}	0.81	**0.98**
5 cm from wall	67.8	0.2876	0.0023	1.21	**1.08**	0.1065	1.2×10^{-5}	0.81	1.94
10cm from wall	62.8	0.1838	0.0036	1.13	**1.08**	0.0462	2.8×10^{-5}	0.81	0.84
15 cm from wall	57.8	0.1175	0.0057	**1.06**	**1.07**	0.0200	6.4×10^{-5}	0.82	0.89
50% to centre	36.4	0.0173	0.0384	**0.91**	**1.04**	0.0006	0.0023	0.88	**0.98**
75% to centre	18.2	0.0037	0.1961	0.59	**1.01**	3×10^{-5}	0.0478	0.74	**1.00**
1 cm from centre	1.0	0.00073	0.9144	0.08	**1.00**	1×10^{-5}	0.8461	**1.00***	**1.00**

A rough validation of the **SS** cut was attempted by predicting the probability of a neutron only scattering once. The probability of a neutron scattering within a cylinder of radius r, when a narrow width neutron beam is directed at its centre is:

$$P(S) = \frac{\Sigma_S}{\Sigma_T}(1 - e^{-2\Sigma_T R}) \tag{7.6}$$

where Σ_T is the total macroscopic cross section and Σ_S is the scattering macroscopic cross section:

$$\Sigma_T = \Sigma_e + \Sigma_i + \Sigma_c \quad \Sigma_S = \Sigma_e + \Sigma_i \tag{7.7}$$

where e, i and c refer to elastic, inelastic and capture respectively. Macroscopic cross sections are the cross section multiplied by the density of the material:

$$\Sigma = \sigma\rho \tag{7.8}$$

As the single scatter cut will allow events with a radial energy weighted variance (σ_r defined in equation 7.4) of less than 3 cm to be called a single scatter, the formula can be used with R=3 cm
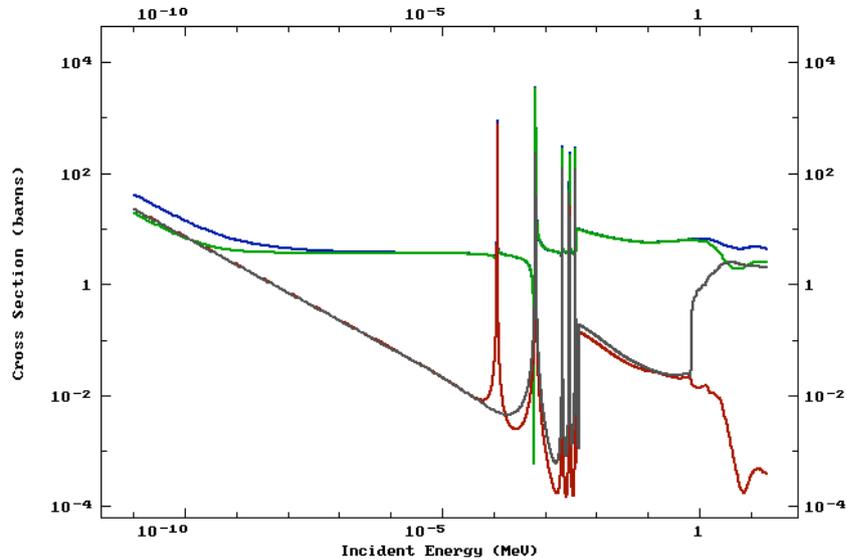
Figure 7.2: Neutron cross sections on ^{132}Xe as a function of incident neutron energy. The total cross section is shown in blue, the elastic in green, the capture cross section in red and inelastic in grey [146].

to estimate the probability of this happening for a 1 MeV neutron as 0.235. The probability of a scatter within the whole xenon volume of a 1 MeV neutron is found to be $P(S_{xe}) = 0.996$. These can be combined to obtain the probability of the event being tagged as a single scatter as:

$$P(SS) = P(S_{xe})[1 - P(S_{xe})] + P(S)P(S_{3cm})[1 - P(S_{xe})] = 0.005 \tag{7.9}$$

where i.e. the sum of the probability that a neutron will scatter and then escape the xenon and the probability the neutron will scatter, scatter again within 3 cm and then escape the TPC. This is expected to be a large overestimate as it does not take into account energy loss when scattering, and the cross sections grow as the energy decreases, see figure 7.2. The number obtained from simulation for the probability of neutrons to pass the **SS** cut is 0.0007, so 14% of the predicted value, but as stated the prediction was an overestimate, so this seems reasonable.

7.1.5 Generator Validation

ER simulations were initially performed using several generators; *DecayChain*, *SingleDecay* and *U238/Th232*. As their names suggest, *DecayChain* was used for simulating ^{238}U and ^{232}Th chains, and *SingleDecay* was used for ^{60}Co, ^{40}K, ^{46}Sc and any other single isotope decays of interest. *U238* and *Th232* were simpler models of the decay chains, as explained below.

The *DecayChain* generator was set using a command of the form:

/LUXSim/source/set volume DecayChain_xxx activity units age of source units

so as an example:

/LUXSim/source/set InnerTitaniumVessel DecayChain_U238 1 mBq/kg 4.5e7 y

The requirement for the age of the source was so the generator could calculate the population ratios of each isotope in the chain. Furthermore, when beginning the simulation with the command:

/LUXSim/beamOn nEvents

the number of events would be used along with the set activity and component mass to calculate a livetime. This would be used along with the population ratios and half lives of the isotopes to determine which decays to generate. This meant that simulation results were actually highly dependent on the number of event simulated.

The *U238* and *Th232* generators used a simpler method, setting the chains in secular equilibrium; each isotope in the chain was assigned equal probabilities of decay, except for where there are two possible decays, where it used the relevant branching ratios. The generator used the GEANT4 commands:

/grdm/nucleusLimits Amin Amax Zmin Zmax and */gps/ion Z A Q E*

to generate each isotope. *grdm* is the G4RadioactiveDecay module, and nucleusLimits sets the limits on the atomic weight and number for the generated decays. *gps* is GEANT4's general particle source, and the ion command will generate an ion with atomic number *Z*, mass *A*, charge *Q* and excitation energy *E*.

A validation of the *U238* and *Th232* generators was performed using 300,000,000 events generated within the inner and outer cryostat vessels. A measure of how well the generator was modelling the decay chains was to count the total number of γ-rays emitted on average throughout the decays in one full chain, which are known to be 2.235 and 2.629 on average for ^{238}U and ^{232}Th respectively.

Initial counting results showed an excess in γ-rays in the uranium chain, by approximately a factor of 2. In order to investigate, each decay was generated individually using the *grdm* and *gps* commands that the generator used. A problem was quickly identified in the decay of ^{234}Th, which decays to a metastable state of palladium, which then decays to ^{234}U:

$$^{234}_{90}\text{Th} \rightarrow \, ^{234m}_{91}\text{Pa} + e^- + \bar{\nu}_e \tag{7.10}$$

$$^{234m}_{99}\text{Pa} \rightarrow \, ^{234}_{92}\text{U} + e^- + \bar{\nu}_e \tag{7.11}$$

^{234m}Pa has a half life of 1.159 minutes. It was found that the following decay to ^{234}U was automatically simulated within the same event, presumably either intentionally because of the metastable state, or because of a bug within GEANT4. Then, when individually generating ^{234}Pa, with the *grdm* and *gps* commands, it would not generate the metastable state but a ground state, which decays to an excited state of ^{234}U. This excited uranium isotope emitted on average 2 γ-rays as it decayed to the ground state. This explained the excess in γ-rays generated by the generator. A fix was applied to the *U238* generator by removing the generation of ^{234}Pa as it was already taken care of within the generation of ^{234}Th.

A similar check was then done with the *DecayChain* generator, and the same issue was found to be present. This was the generator in use for LZ background simulations, so this was problematic. Table 7.4 shows the results of tests done on the original *U238* generator, the *DecayChain* generator and the fixed *U238* generator, which counted the proportion of each isotope of the generated decays, and the γ-rays emitted per full chain. Note that the modified version generated on average 2.237 γ-rays, almost the expected number of 2.235.

Table 7.4: Chain splitting and γ-ray counting for three generators; the original *U238* generator, the DecayChain, and a modified version of *U238*. The percentages show the proportion of all generated decays that were of each isotope.

Parent Isotope	*U238*	*DecayChain_U238*	Modified *U238*
^{238}U	7.289%	7.681%	7.601%
^{234}Th	7.207%	7.767%	7.912%
^{234}Pa	7.054%	7.761%	0 (in ^{234}Th)
^{234}U	7.227%	7.847%	7.797%
^{230}Th	7.187%	7.706%	7.614%
^{226}Ra	7.195%	7.732%	7.662%
^{222}Rn	7.059%	7.779%	7.799%
^{218}Po	7.033%	7.781%	7.553%
^{218}At	0.000%	0.002%	0.000%
^{214}Pb	7.204%	7.779%	7.714%
^{214}Bi	7.007%	7.780%	7.74%
^{214}Po	7.231%	7.780%	7.486%
^{210}Pb	7.072%	7.603%	7.637%
^{210}Tl	0.000%	0.000%	0.001%
^{210}Bi	7.126%	6.384%	7.931%
^{210}Po	7.109%	0.618%	7.553%
γ **Per Chain**	4.199	4.584	2.237

This finding led to the development of an entirely new generator, called *G4Decay*. This could be used to generate any single decay as well as the full uranium and thorium chains. *G4Decay* generates a full decay chain for every number given to the *beamOn* command but saves each decay as a separate event for easier analysis. It also grouped together very close in time decays such as the ^{214}Bi - ^{214}Po decays. The results from the *DecayChain_238_92* uranium chain generator were found to match the modified *U238* results to within 1%, and the correct number of γ-rays was generated. Background ER simulations from then on were performed using *G4Decay*.

7.2 LZ Background Model

7.2.1 Electron Recoil Simulations

ER simulations all invoke the use of the *G4Decay* generator, which is used with the command:

/LUXSim/source/set volume G4Decay_A_Z activity units

where A and Z are the atomic mass and number of the required isotope. The activity can be set as a total for example in mBq, or by mass, in mBq/kg. LZSim calculates the mass of the component (which can also be set with a command) to use to scale the activity. In general, the activities are set as 1 mBq/kg and scaled to the screening activities in analysis.

Table 7.5 lists the simulations performed for LZ electron recoil backgrounds for the Technical Design Review.

Table 7.5: Electron recoil background simulations performed for the LZ TDR. *These components were not present in the geometry and so were simulated as point source

	Component	Sources
PMTs	Xe	^{238}U, ^{232}Th, ^{60}Co, ^{40}K
	Skin	^{238}U, ^{232}Th, ^{60}Co, ^{40}K
	Dome	^{238}U, ^{232}Th, ^{60}Co, ^{40}K
	Water	^{238}U, ^{232}Th
TPC	PTFE	^{238}U, ^{232}Th, ^{60}Co, ^{40}K
	Loop antennae*	^{238}U, ^{232}Th, ^{60}Co, ^{40}K
	Position sensors*	^{238}U, ^{232}Th, ^{60}Co, ^{40}K
	Long level sensors*	^{238}U, ^{232}Th, ^{60}Co, ^{40}K
	Weir precision sensors*	^{238}U, ^{232}Th, ^{60}Co, ^{40}K
	Acoustic sensors*	^{238}U, ^{232}Th, ^{60}Co, ^{40}K
	Thermometers*	^{238}U, ^{232}Th, ^{60}Co, ^{40}K
Cryostat	Inner vessel	^{238}U, ^{232}Th, ^{46}Sc, ^{40}K
	Outer vessel	^{238}U, ^{232}Th, ^{46}Sc, ^{40}K
	Seals	^{238}U, ^{232}Th, ^{60}Co, ^{40}K
	Flange bolts	^{238}U, ^{232}Th, ^{60}Co, ^{40}K
Benchmarks	Bottom fluid conduit	^{238}U, ^{232}Th, ^{60}Co, ^{40}K
	Up conduit	^{238}U, ^{232}Th, ^{60}Co, ^{40}K
	Top dome	^{238}U, ^{232}Th, ^{60}Co, ^{40}K

The benchmark points in table 7.5 were point sources at important locations in the geometry. These were used to calculate background from components such as cabling.

7.2.2 Nuclear Recoil Simulations

For neutron sources, the neutron (α,n) energy spectrum for each material was obtained from SOURCES4A [150], and single neutrons emitted isotropically with energies sampled from the spectrum. The uranium chain was split into early and late, as many components showed evidence of equilibrium breaking in uranium chain screening results. For thorium this was deemed unnecessary, and the full chain spectra was generated. Table 7.6 shows the components that were

185

simulated for nuclear recoil backgrounds, along with the materials used to generate the (α,n) spectra.

Uranium early chain spontaneous fission can contribute to neutron yields significantly in some materials where (α,n) rates are low and U_e radioactivity is high. Initial background estimations for LZ included spontaneous fission neutrons, emitted singularly. However, when spontaneous fission neutron was correctly modeled as near-simultaneous emission of up to 6 neutrons and 20 γ-rays, the mulitplicity of the event allowed vetoing of a much higher efficiency (100$\times$ better when the spontaneous fission events occurred in the cryostat) than with single neutrons. Therefore, spontaneous fission neutrons are excluded from the background simulations. The study that led to this decision can be found in the section 7.4.

Table 7.6: Nuclear recoil background simulations performed for the LZ TDR. For each component, (α,n) spectra for uranium early chain, uranium late chain and the full thorium chain were generated using the material shown. *These components were not present in the geometry and so were simulated as point source.

	Component	Material
PMTs	Xe - bodies	Kovar
	Xe - photocathodes	SiO_2
	Xe - ceramics	Al_2O_3
	Xe bases - capacitors	$BaTiO_3$
	Xe bases - connectors	BeCuNi
	Xe bases - resistors	Al_2O_3
	Skin	Al_2O_3
	Dome	Al_2O_3
	Water	Si_2
TPC	PTFE	PTFE
	Loop antennae*	Peek
	Position sensors*	Cirlex
	Long level sensors*	Al
	Weir precision sensors*	Al
	Acoustic sensors*	PVDF
	Thermometers*	Cirlex
Cryostat	Inner vessel	Ti
	Outer vessel	Ti
	ICV Flange Seals	Al
	Other Seals	Elastomer
	Flange bolts	SS316l
Benchmarks	Bottom fluid conduit	Cu
	Up conduit	Cu
	Top dome	Cu

7.2.3 Analysis and Background Counts

The reduced tree previously described was analysed for each component-source pair to determine how many events constitute a background. Such events must satisfy the **ROI**, **SS**, **Skin** and **OD**

cuts. After application of these cuts, a survival factor is calculated as:

$$P(ER/NR) = \frac{N_{surviving}}{N_{generated}} \tag{7.12}$$

where $N_{surviving}$ is the number of events surviving all cuts, and $N_{generated}$ is the total number of events generated.

Each simulation was done with a number of events equivalent to at least a single LZ 1000-day exposure, with most being several orders of magnitude longer than this. Components expected to be dominant sources of radioactivity such as the PMTs and the cryostat were simulated with more events to increase statistics. NR simulations always required less events than ER as neutron survival factors are several orders of magnitude higher than for ER sources. The number of events generated for each component-source pair ranged from 1,000,000 for lower priority neutron simulations to 1,000,000,000 for high priority ER such as ^{238}U in the cryostat.

Table 7.7 demonstrates survival factors for events simulated for each component-source pair. It is important to distinguish between P(NR) which is the survival probability for a single neutron and P(ER) which is the survival probability per decay.

Table 7.7: Survival factors P(ER) (per decay) and P(NR) (per neutron) for each component-source pair used in the LZ TDR background simulations. A '0' indicates no surviving events. All probabilities are given to three significant figures to demonstrate the size of the difference between similar simulations. *These components were not present in the geometry and so were simulated as point sources. †This was a ^{46}Sc simulation.

Component	P(ER)				P(NR)		
	^{238}U	^{232}Th	^{60}Co	^{40}K	^{238}U$_e$	^{238}U$_l$	^{232}Th
Xe PMTs - bodies					6.90E-5	7.30E-5	6.35E-5
Xe PMTs- photocathodes	1.58E-8	1.65E-8	1.36E-8	4.13E-9	5.35E-5	5.05E-5	4.85E-5
Xe PMTs- ceramics					5.35E-5	5.40E-5	4.80E-5
Xe PMT bases - capacitors					3.05E-5	2.80E-5	3.40E-5
Xe PMT bases - connectors	1.07E-8	1.58E-8	1.24E-8	2.75E-9	2.60E-5	3.65E-5	2.40E-5
Xe PMT bases - resistors					5.40E-5	4.15E-5	4.50E-5
Skin PMTs	6.88E-10	1.03E-9	3.47E-10	0	3.90E-5	3.50E-5	3.55E-5
Dome PMTs	1.38E-9	1.38E-9	0	0	5.55E-5	4.85E-5	4.50E-5
Water PMTs	0	0	0	0	0	0	0
PTFE	3.85E-8	3.47E-8	6.02E-9	9.36E-9	1.45E-4	1.10E-4	1.25E-4
Loop antennae*	0	5.52E-9	0	0	2.70E-5	1.80E-5	1.20E-5
Position sensors*	0	0	6.88E-9	0	1.40E-5	1.70E-5	1.20E-5
Long level sensors*	0	0	1.38E-8	0	5.90E-5	4.20E-5	5.20E-5
Weir precision sensors*	1.38E-10	1.38E-10	0	0	2.20E-5	2.70E-5	2.50E-5
Acoustic sensors*	5.50E-9	1.10E-8	6.88E-19	6.88E-9	5.00E-5	4.20E-5	2.80E-5
Thermometers*	9.63E-9	1.53E-9	1.89E-8	2.29E-9	3.40E-5	4.57E-5	4.33E-5
Vessels	4.68E-9	4.47E-9	2.58E-9†	1.03E-9	3.55E-5	3.20E-5	3.25E-5
ICV Flange Seals	1.03E-9	6.88E-10	3.53E-9	0	1.85E-5	1.95E-5	1.85E-5
Other Seals	2.75E-9	2.06E-9	3.53E-9	0	2.30E-5	1.55E-5	1.55E-5
Flange bolts	6.88E-10	2.75E-9	0	0	2.25E-6	1.85E-5	1.85E-5
Bottom fluid conduit	0	0	0	0	5.10E-5	7.80E-5	6.90E-5
Up conduit	0	0	0	0	2.10E-5	1.00E-5	2.00E-5
Top dome	0	0	0	0	2.00E-5	3.30E-5	2.90E-5

Because of the chain breaking mentioned previously, ^{238}U chain ER simulation results contained additional information in the form of a boolean that flagged whether an event came from the early chain. Due to the low branching ratios of γ-rays in the early chain, most of the survival factors in table 7.7 are calculated from entirely late chain events. The exceptions are shown in table 7.8

Table 7.8: Components for which some surviving ^{238}U events were from the early ^{238}U chain. Survival probabilities are split into early and late chain for combination with screening results.

Component	Fraction U_e	P(ER)$_e$	P(ER)$_l$
Xe PMTs	0.029	4.58E-10	1.53E-8
Xe PMT Bases	0.016	1.71E-10	1.05E-8
PTFE	0.018	6.93E-10	3.78E-8
Cryostat Vessels	0.029	1.36E-10	4.54E-9

The fraction of early chain events was calculated and separate survival factors used with early and late chain activities.

Component masses, activities and neutron yields are used to calculate emissions or decays per year for each component i as:

$$E^i_{yr} = N_{sec,yr} \cdot A^i \cdot M^i \cdot Y^i \qquad (7.13)$$

where $N_{sec,yr}$ is the number of seconds in a year, A is the component activity in Bq/kg, M is the mass of the component in kg and Y is the yield (1 for γ-rays, and in units of neutrons/decay for neutrons). P(ER)/P(NR) can be combined directly with each E_{yr} to determine an expected background rate. For components that were not simulated, the results of the simulation most well matched to their physical location are used. For neutrons simulations the material for (α,n) reactions is also considered to choose the closest match.

The total background from all components calculated as described above for the nominal LZ exposure is summarised in table 7.9. This is the result of a combination of the aforementioned screening procedures and background simulations. The backgrounds table also includes backgrounds from internal ^{222}Rn, ^{220}Rn, krypton and argon within the LXe. Furthermore, an estimation of the laboratory and rock backgrounds (simulated as in section 7.3) and fixed surface contamination are included. The ^{210}Bi is included due to reasons discussed in section 4.5. Finally, physics backgrounds from double β-decay of ^{136}Xe and neutrinos are shown. The final two lines show the backgrounds before and after an ER discrimination of 99.5% and a NR acceptance of 50% are applied. The 50% for NR acceptance is based on the signal region defined as below the mean of the NR band.

A total of 1,240 ER counts is expected (before discrimination), although if radon goals are met, this is reduced significantly to 438. Moreover, this would result in astrophysical neutrinos being the major background, meaning that most analyses would encounter only irreducible backgrounds.

Table 7.9: Estimated background counts in the WIMP search region of interest from all background sources in the LZ 1000 day exposure. Mass-weighted average activities are shown for composite materials. Solar ^{8}B neutrinos are expected to contribute 7 ± 3 NR but only at very low energies and are excluded from the table.

	Mass	^{238}U$_e$	^{238}U$_l$	^{232}Th$_e$	^{232}Th$_l$	^{60}Co	^{40}K	n/yr	ER	NR
	(kg)	mBq/kg							(cts)	(cts)
Upper PMT Structure	40.5	3.90	0.23	0.49	0.38	0.00	1.46	2.53	0.05	0.000
Lower PMT Structure	69.9	2.40	0.13	0.30	0.24	0.00	0.91	6.06	0.05	0.001
R11410 3" PMTs	91.9	71.6	3.20	3.12	2.99	2.82	15.4	81.8	1.46	0.013
R11410 PMT Bases	2.8	288	75.8	28.4	27.9	1.43	69.4	34.7	0.36	0.004
R8778 2" Skin PMTs	6.1	138	59.4	16.9	16.9	16.3	413	52.8	0.13	0.008
R8520 Skin 1" PMTs	2.2	60.5	5.19	4.75	4.75	24.2	333	4.60	0.02	0.001
R8520 PMT Bases	0.2	213	108	42.2	37.6	2.23	124	3.62	0.00	0.000
PMT Cabling	104	29.8	1.47	3.31	3.15	0.65	33.1	2.65	1.43	0.000
TPC PTFE	184	0.02	0.02	0.03	0.03	0.00	0.12	22.5	0.06	0.008
Grid Wires	0.8	1.20	0.27	0.33	0.49	1.60	0.40	0.02	0.00	0.000
Grid Holders	62.2	1.20	0.27	0.33	0.49	1.60	0.40	6.33	0.27	0.002
Field Shaping Rings	91.6	5.41	0.09	0.28	0.23	0.00	0.54	10.8	0.23	0.004
TPC Sensors	0.90	21.1	13.5	22.9	14.2	0.50	26.3	24.8	0.01	0.002
TPC Thermometers	0.06	336	90.5	38.5	25.0	7.26	3360	1.49	0.05	0.000
Xe Tubing	15.1	0.79	0.18	0.23	0.33	1.05	0.30	0.64	0.00	0.000
HV Components	138	1.90	2.00	0.50	0.60	1.40	1.20	4.90	0.04	0.001
Conduits	200	1.25	0.40	2.59	0.66	1.24	1.47	5.33	0.06	0.001
Cryostat Vessel	2410	1.59	0.11	0.29	0.25	0.07	0.56	124	0.63	0.013
Cryostat Seals	33.7	73.9	26.2	3.22	4.24	10.0	69.1	38.8	0.45	0.002
Cryostat Insulation	23.8	18.9	18.9	3.45	3.45	1.97	51.7	69.8	0.43	0.007
Cryostat Teflon Liner	26	0.02	0.02	0.03	0.03	0.00	0.12	3.18	0.00	0.000
Outer Detector Tanks	3200	0.16	0.39	0.02	0.06	0.04	5.36	78.0	0.45	0.001
Liquid Scintillator	17600	0.01	0.01	0.01	0.01	0.00	0.00	14.3	0.03	0.000
Outer Detector PMTs	205	570	470	395	388	0.00	534	7590	0.01	0.000
OD PMT Supports	770	1.20	0.27	0.33	0.49	1.60	0.40	14.3	0.00	0.000
Subtotal (Detector Components)									6.20	0.070
^{222}Rn (2.0 μBq/kg)									722	-
^{220}Rn (0.1 μBq/kg)									122	-
natKr (0.015 ppt g/g)									24.5	-
natAr (0.45 ppb g/g)									2.47	-
(continued on next page)										

189

	Mass	$^{238}\text{U}_e$	$^{238}\text{U}_l$	$^{232}\text{Th}_e$	$^{232}\text{Th}_l$	^{60}Co	^{40}K	n/yr	ER	NR
	(kg)			mBq/kg					(cts)	(cts)
^{210}Bi (0.1 μBq/kg)									40	-
Laboratory and Cosmogenics									4.3	0.06
Fixed Surface Contamination									0.19	0.37
Subtotal (Non-ν counts)									921	0.50
^{136}Xe $2\nu\beta\beta$									67	0.00
Astrophysical ν counts (pp+^{7}Be+^{13}N)									255	0.00
Astrophysical ν counts (^{8}B)									0.00	0.00
Astrophysical ν counts (Hep)									0.00	0.21
Astrophysical ν counts (diffuse supernova)									0.00	0.05
Astrophysical ν counts (atmospheric)									0.00	0.46
Subtotal (Physics backgrounds)									322	0.72
Total									1,240	1.22
Total (with 99.5% ER discrimination, 50% NR acceptance)									6.22	0.61
Sum of NR-like ER and NR in LZ for 1000 day, 5.6 T FV, with all analysis cuts									**6.83**	

7.3 Outer Detector Simulations

7.3.1 Neutron capture in the OD

Given the importance of the OD, the simulation must be extensively validated. The capture cross sections for thermal neutrons can be used along with the abundances of the isotopes to calculate the probability of capture. A validation of the isotopic content of the OD was done by comparing these probabilities with the spectrum of γ-rays released from the OD. Figure 7.3 shows the energy of any particle with a GEANT4 creator process of *nCapture*.

Table 7.10: Table of cross sections, interaction lengths and energy released as γ-rays for the three isotopes that capture neutrons in the OD.

Isotope	Thermal σ (barns)	Interaction length (cm)	% Captured	E Released (MeV)
H	0.33	44	12.3	2.23
^{155}Gd	61,000	34	16.2	8.54
^{157}Gd	250,000	7.7	71.5	7.94

Capture on hydrogen results in the release of a 2.2 MeV γ-ray, whilst gadolinium captures are higher energy, see table 7.10. By integrating the energy spectrum shown in figure 7.3 above and below 2.23 MeV, a lower limit was placed on captures on gadolinium of >86.2% and an upper limit on captures on hydrogen of <13.8%. These are limits as some of the low energy continuum will also be from gadolinium. This agrees with the percentages shown in table 7.10.
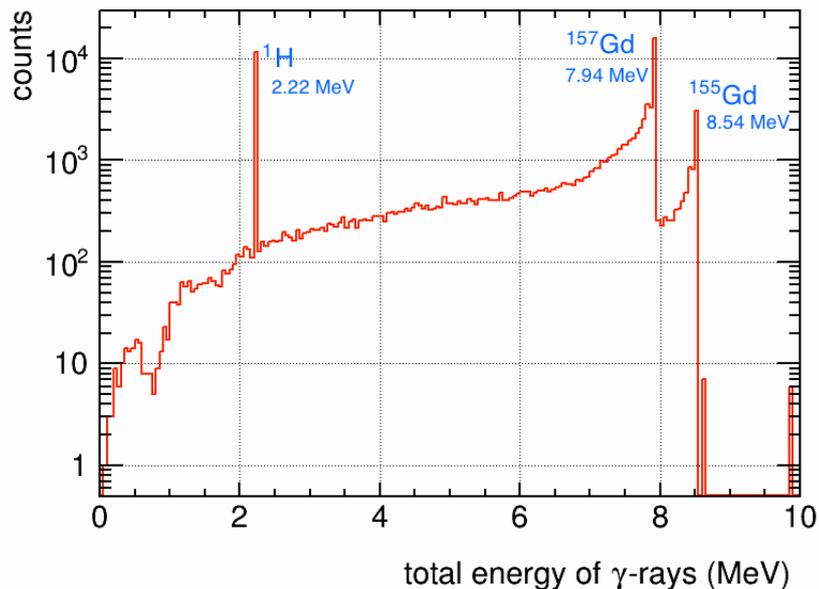
Figure 7.3: Summed energy of all neutron capture products within the outer detector, with the lines from ^{1}H, ^{155}Gd and ^{157}Gd identified.

Furthermore, the capture time of neutrons within the scintillator is an important parameter, as the veto window time must be large enough to allow >95% of neutrons to be captured, but low enough to not adversely increase deadtime to over 5%. The initial decision for the length of the vetoing window, Δt_{OD}, was 800 μs, but this would result in unacceptable deadtime. With 125 μs, deadtime is < 2%. Early simulations had shown the capture time constant in the Gd-loaded LAB to be ~30 μs. However, when we shortened the vetoing window in the TDR background simulations to 125 μs, twice as many neutron events survived the **OD** cut. This seemed unusual, as 125 μs is still greater than 4 capture time constants and so the number of events with capture times greater than this should have been small, if not negligible, but as figure 7.4 demonstrates, there are many events with capture times between 125 and 800 μs. In order to investigate this, thorium chain neutron simulations were performed whilst recording information at record level 3 in all volumes, in order to allow full tracking of neutrons throughout the simulation. As events in LZSim output are stored by volume, code was written to loop through records (a new record is created each time a particle enters a new volume) until the actual event number changed. The code would then loop backwards through each record in the event and create a text output containing time ordered lists showing particle propagation. Background events that had deposited energy within the LXe target, and had also passed $\Delta t_{OD} <$ 800 μs were selected. Looking at events with windows between 125 μs and 800 μs may explain why so many more events are unvetoed when the window was shortened. An example of this output is shown in table 7.11. Here, $\Delta t_{OD} =$ 158 μs between the hit in LXe and the neutron capture and subsequent γ-ray production in the OD. Firstly, many of these events were looked at by eye to attempt to determine
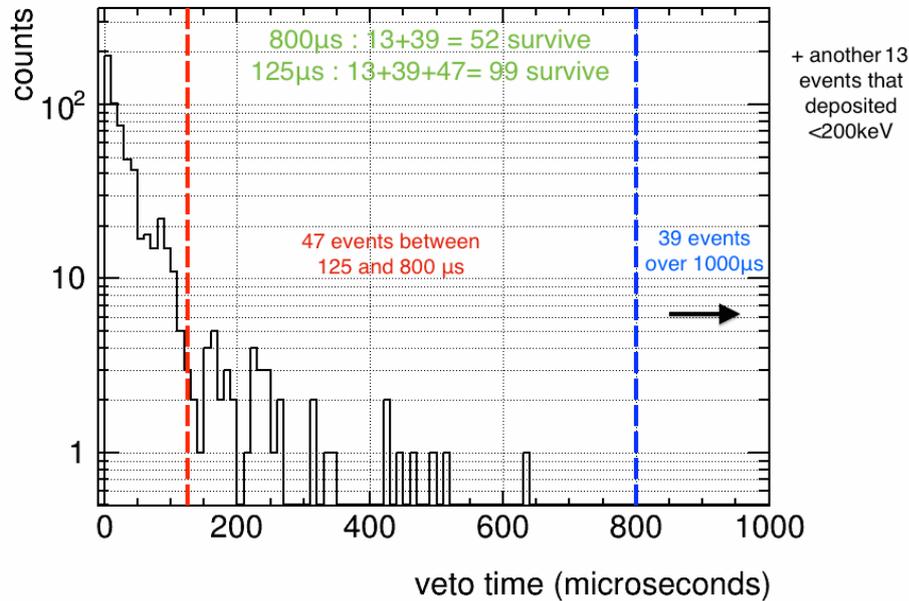
Figure 7.4: Distribution of Δt_{OD} of background events, demonstrating why shortening the veto window caused a factor ~2 difference in the number of background NR counts. Events to the right of the cut lines will pass the cuts and not be vetoed.

if there was some bug or problem, or if the neutrons were just spending a large amount of time outside the scintillator before getting captured. The example in table 7.11 shows the primary neutron starting in the outer vessel, depositing energy in the TPC, then passing through several other volumes including the skin, the vessels again, foam and reflectors and the scintillator tanks and volumes. It was not immediately obvious from the tracking output what the problem was; however, it was noticed that the acrylic tanks and the water of the water tank were appearing in many of the delayed neutron capture events. The tracking output text files were used to calculate the total time spent by each neutron in each volume, and plotted the fraction of Δt_{OD} spent in each volume against Δt_{OD}.

From figure 7.5 it is clear that neutrons do not frequently spend large fractions of time in the xenon volumes and the inner and outer vessels. The water and scintillator itself show longer times, but the surprising result is the acrylic tank. The top of the plot, where the fraction of the veto window is close to 1, is densely populated. This signifies that the problem is the acrylic tank; neutron thermalisation and capture time in acrylic is much longer (O~ 200 μs) than it is in the gadolinium loaded LAB. This study prompted a redesign of the geometry of the acrylic tanks holding the scintillator and an increase of the veto window to 500 μs.

Table 7.11: Tracking of a neutron event. The top row (*italics*) is the γ-ray that triggered the veto. Subsequent rows track the neutron, beginning with its creation in the outer vessel. Times are relative to the hit in LXe (**bold**). For multiple hits with the same process and volume, only the first line is shown.

| Time (μs)	Particle ID	Energy (keV)	Step Process	Creator Process	Volume
158.011	*22*	*383.501*	*Transportation*	*nCapture*	*ScintillatorCenter*
-0.0801411	2112	2250	Transportation	primary	OuterTitaniumVessel
-0.0712823	2112	2250	Transportation	primary	FoamDisplacer
-0.0712173	2112	2250	hadElastic	primary	CenterReflector
-0.0711852	2112	2092.21	Transportation	primary	CenterReflector
-0.0583495	2112	2092.21	hadElastic	primary	FoamDisplacer
-0.0561235	2112	665.95	Transportation	primary	FoamDisplacer
-0.0552142	2112	665.95	Transportation	primary	OuterTitaniumVessel
-0.0414269	2112	665.95	Transportation	primary	VacuumSpace
-0.0404417	2112	665.95	Transportation	primary	InnerTitaniumVessel
-0.0401942	2112	665.95	Transportation	primary	PTFELinerLiquid
-0.0368494	2112	665.95	hadElastic	primary	LiquidSkinXenon
-0.0357026	2112	662.448	Transportation	primary	LiquidSkinXenon
-0.0353713	2112	662.448	Transportation	primary	PTFEWallsInLiquid
-0.0351006	2112	662.448	Transportation	primary	ForwardFieldRing_51
-0.0348378	2112	662.448	hadElastic	primary	PTFEWallsInLiquid
-0.0336084	2112	634.474	Transportation	primary	PTFEWallsInLiquid
0	**2112**	**634.474**	**hadElastic**	**primary**	**LiquidXenonTarget**
0.0202759	2112	616.29	Transportation	primary	LiquidXenonTarget
0.0363975	2112	616.29	Transportation	primary	Top_PMT_Window_034
0.040984	2112	616.29	Transportation	primary	GaseousSkinXenonBank
0.0419073	2112	616.29	Transportation	primary	GaseousSkinXenon
0.0421376	2112	616.29	Transportation	primary	PTFELinerGas
0.0430564	2112	616.29	Transportation	primary	InnerTitaniumVessel
0.0561605	2112	616.29	Transportation	primary	VacuumSpace
0.057038	2112	616.29	Transportation	primary	OuterTitaniumVessel
0.0626967	2112	616.29	Transportation	primary	FoamDisplacer
0.0627615	2112	616.29	Transportation	primary	CenterReflector
0.0642411	2112	616.29	hadElastic	primary	ScintillatorTank
0.0661305	2112	338.803	Transportation	primary	ScintillatorTank
0.0680112	2112	338.803	hadElastic	primary	ScintillatorCenter
0.0839189	2112	20.0534	hadElastic	primary	ScintillatorTank
0.122164	2112	0.511666	Transportation	primary	CenterReflector
0.174913	2112	0.511666	hadElastic	primary	FoamDisplacer
1.67146	2112	0.02338	Transportation	primary	CenterReflector
1.70864	2112	0.02338	hadElastic	primary	ScintillatorTank
3.51268	2112	0.000290067	Transportation	primary	CenterReflector
19.6411	2112	0.000290067	Transportation	primary	FoamDisplacer
22.5439	2112	0.000290067	Transportation	primary	OuterTitaniumVessel
111.095	2112	0.000290067	Transportation	primary	VacuumSpace
113.043	2112	0.000290067	hadElastic	primary	OuterTitaniumVessel
113.598	2112	0.000275645	Transportation	primary	OuterTitaniumVessel
124.46	2112	0.000275645	Transportation	primary	FoamDisplacer
124.583	2112	0.000275645	Transportation	primary	CenterReflector
124.885	2112	0.000275645	hadElastic	primary	ScintillatorTank
127.725	2112	6.2463e-05	Transportation	primary	ScintillatorTank
128.153	2112	6.2463e-05	Transportation	primary	CenterReflector
129.86	2112	6.2463e-05	hadElastic	primary	FoamDisplacer
133.539	2112	6.89598e-05	Transportation	primary	FoamDisplacer
133.769	2112	6.89598e-05	Transportation	primary	CenterReflector
138.736	2112	6.89598e-05	hadElastic	primary	ScintillatorTank
147.276	2112	2.27942e-05	Transportation	primary	ScintillatorTank
148.822	2112	2.27942e-05	hadElastic	primary	ScintillatorCenter
158.011	2112	4.45006e-05	nCapture	primary	ScintillatorCenter
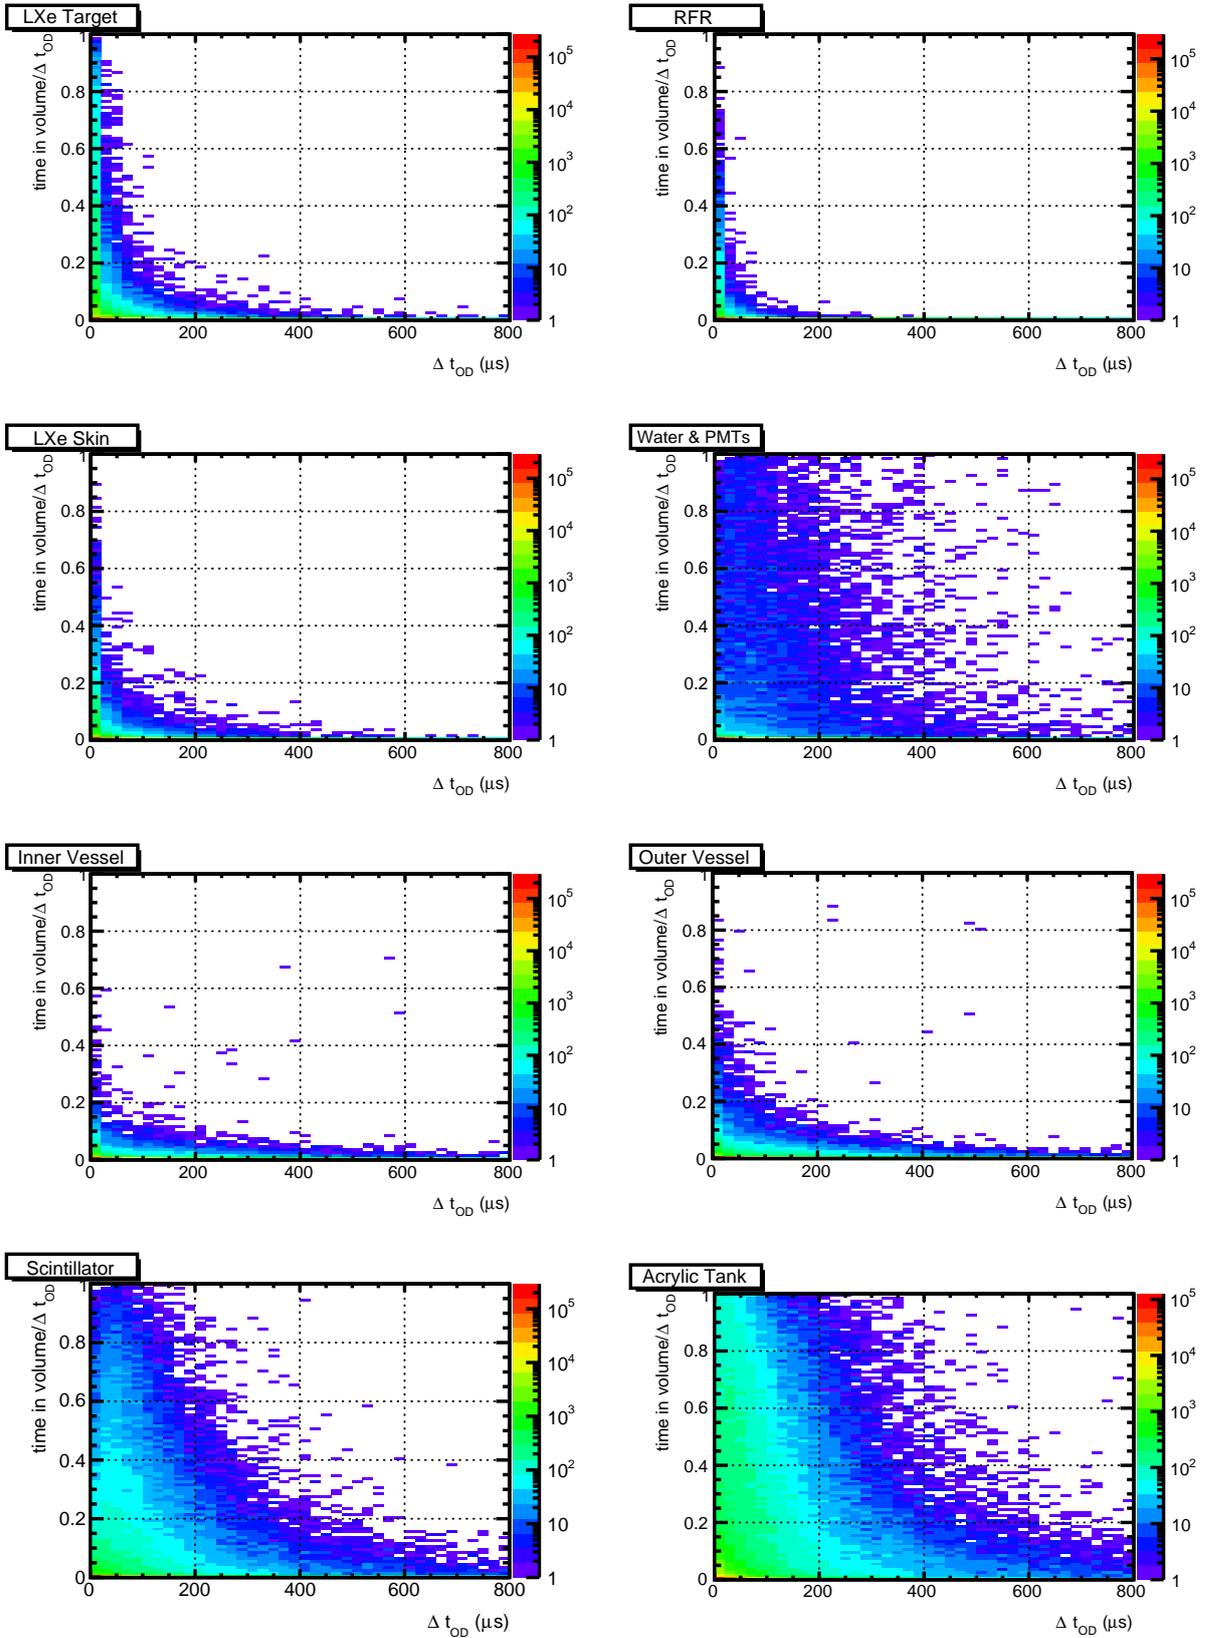
Figure 7.5: Fraction of Δt_{OD} against Δt_{OD} for different LZSim volumes.

7.3.2 Veto Inefficiency

The LZ requirements for the outer detector demand at least a 95% vetoing efficiency for neutrons, and > 70% for γ-rays of a few MeV (an energy range where it is possible for them to deposit energy in the TPC and then escape it). After the issue in the previous section was identified, a new value of Δt_{OD} was required. Background simulations were analysed in order to determine the efficiency of the OD as a function of the vetoing window, for no threshold (i.e. any energy deposit), >100 keV and >200 keV. The original threshold was 100 keV, but there was concern over intrinsic ^{14}C contamination in the LAB: ^{14}C decays by β-decay with an endpoint of 156 keV. Raising the threshold to 200 keV would remove almost all of this background. Neutron simulations that were produced for the TDR were used to produce figure 7.6. The inefficiency is defined as:

$$\epsilon(E_{OD}, \Delta t_{OD}) = \frac{\text{events passing } \textbf{ROI}, \textbf{SS}, \textbf{FV}, \textbf{Skin} \text{ and } \textbf{OD}(E_{OD}, \Delta t_{OD})}{\text{events passing } \textbf{ROI}, \textbf{SS} \text{ and } \textbf{FV}} \quad (7.14)$$

where E_{OD} and Δt_{OD} refer to the OD energy threshold and time window respectively. This gives a measure of the percentage of ROI single scatter events that went completely unvetoed.
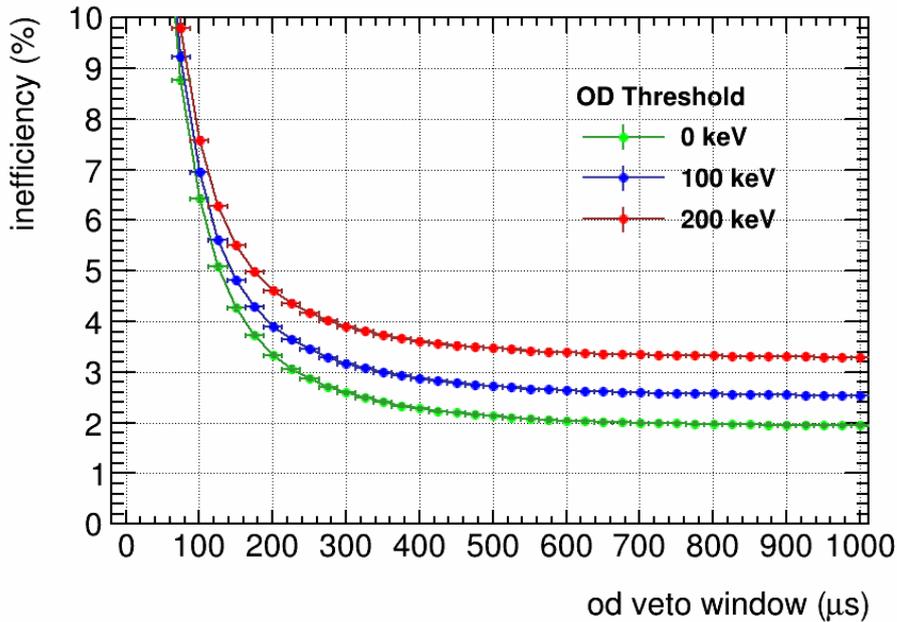

Figure 7.6: The OD inefficiency (i.e. % of unvetoed neutrons) as a function of the veto window, for three energy thresholds.

7.3.3 OD Rate

The outer detector must be able to successfully veto as many background events in the LXe as possible; a high overall rate in the OD from background sources that do not necessarily have to deposit energy in the TPC would have a negative effect on the OD's vetoing ability. In order to assess the rate of energy deposits within the OD from all background sources, the set of simulations performed for the TDR were analysed whilst only looking at hits in the OD volume. The initial results are shown in table 7.12. The same methodology using survival factors from the simulations with emission rates from radioassays was used to calculate a rate for each component. Where a component had not been simulated, for example for conduits and cables, the closest matching physically located source volume was chosen for the survival factor instead.

Table 7.12: OD background rates from detector component radioactivity, calculated using radioassay and simulation results, for an OD threshold of 200 keV. Numbers in brackets show which simulation results were used when there was not one available for that component. The maximum rates are calculated assuming every single emission causes an event in the OD. *This refers to the benchmark points.

No.	Detector Component	Rate (Hz)	Max Rate (Hz)
1	LXe PMTs	0.32	21.98
2	Skin PMTs	0.09	0.89
3	PMT Bases	0.10	3.82
4	Skin PMT Bases	0.02 (2)	1.09
5	Upper PMT Structure	0.01 (1)	0.41
6	Lower PMT Structure	0.01 (1)	0.45
7	PMT Cabling	0.12 (BP*)	6.97
8	PTFE	0.00	0.05
9	Grid Wires	0.00 (8)	0.00
10	Grid Holders	0.03 (8)	0.37
11	Field shaping rings	0.02 (12)	1.34
12	TPC Sensors	0.00	0.12
13	TPC Thermometers	0.00	0.04
14	Xe Recirculation Tubing	0.00 (BP*)	0.06
15	HV Conduits and Cables	0.03 (BP*)	1.58
16	HX and PMT Conduits	0.13 (BP*)	1.88
17	Cryostat Vessels	0.78	11.00
18	Cryostat Seals	0.36	3.77
19	Cryostat Insulation	0.41 (17)	2.38
20	Cryostat Teflon Liner	0.07 (17)	0.15
21	Outer Detector Tanks	1.93 (17)	7.02
22	Liquid Scintillator	0.26 (17)	1.37
23	Outer Detector PMTs	0.57	285.19
24	Outer Detector PMT Supports	0.01 (23)	4.61
	Total	**5.28**	**356.54**

The total rate for detector components was found to be 5.28 Hz, which is suitably low. A modified version of LZSim was used that included a shell of rock to mimic the rock in the Davis cavern, in order to investigate the event rate that γ-rays from the rock would produce within the OD. The simulation had originally been designed to investigate the rate within the liquid xenon

from rock γ-rays, and so included a feature that saved γ-rays on a surface a certain distance from the detector, then re-propagated these γ-rays with a boost in statistics in a second, separate simulation. This could be done as many times as necessary, but for the outer detector it was found that no boosting was needed.

Activities in the rock of the Davis cavern have been measured with the HPGe detector MAEVE in the East Counting Room, but more recently an assay has been done of gravel taken from under the LUX water tank, also performed by MAEVE but in its new location in the Black Hills Underground Campus at SURF. As the gravel is unprocessed natural rock, the values should be valid for the entire decay chains in secular equilibrium. Results from both are shown in table 7.13. This study used the higher, γ-spectroscopy results as the gravel measurement was unavailable at the time. Currently, because of the discrepancy it is unclear what the actual flux of γ-rays is inside the cavern, but there are plans to place a small liquid scintillator screener inside the water tank which will be able to measure the flux that the OD will actually receive.

Table 7.13: Radioassay measurements for rock performed in the East Counting Room and gravel from under the water tank. γ-fluxes are taken from the ^{214}Bi 1.74 MeV high energy γ-line for ^{238}U, the ^{208}Tl 2.62 MeV line for ^{232}Th and the ^{40}K 1.46 MeV line, and are converted to ppm and Bq/kg.

		Rock		Gravel	
Isotope	**γ-flux (cm^{-2}s^{-1})**	**Concentration (ppm)**	**Activity (Bq/kg)**	**Concentration (ppm)**	**Activity (Bq/kg)**
^{238}U	0.059	5.95	73.4	1.65	20.3
^{232}Th	0.056	6.42	26.1	0.302	1.21
^{40}K	0.36	2.31% (^{nat}K)	716	0.0662% (^{nat}K)	20.4

Initial results showed a very high rate of 3,699 Hz in the OD from rock γ-rays for a 200 keV threshold, especially near the bottom of the tanks, see figure 7.7. 94% of this is in the bottom half of the tank. With no threshold on the energy deposit, the rate was 6,137 Hz. This was very concerning as it would severely compromise OD performance.

Upon further investigation, it was found that the steel pyramid under the water tank had its material accidentally set as water within the GEANT4 geometry file. Changing this to steel reduced the rate to 1,520 Hz for a threshold of 200 keV, and 2,963 Hz with no threshold. This is better, but still considerably higher than needed for the $< 5\%$ deadtime requirement, which is 130 Hz.

Possible design alterations to mitigate this and reduce the rate are to shorten the tanks, moving them further into the water tank vertically to increase the γ-ray shielding from the water. A second, more difficult suggestion is to place lead on the bottom and top of the water tank. Whether or not this is feasible depends on the mechanical strength of the water tank. However, because of the different results obtained from the γ-ray flux and the radioactivity in the gravel, the full scope of the problem is not yet known. These results were obtained using the higher activities measured in the East Counting Room, and so the rate could already be a factor of 4–25 times smaller depending on the isotope. The screener that will be deployed in November 2016
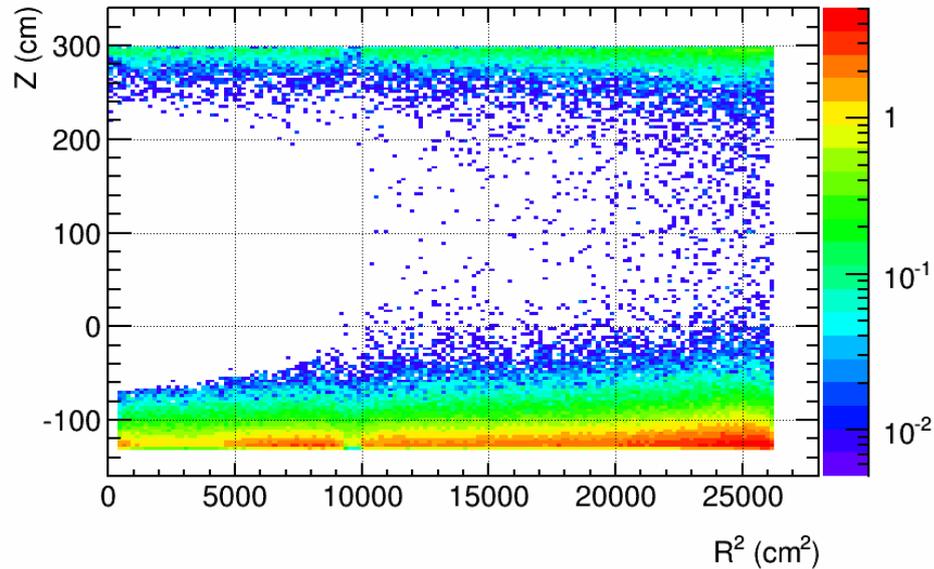
Figure 7.7: Temperature plot showing the location and rate (z axis) of energy deposits above a threshold of 200 keV within the outer detector from γ-rays emitted by radioisotopes in the cavern rock. 94% are within the bottom half, Z < 85 cm.

will give a definitive flux and further changes to the design will be actioned then. It should be noted design changes to tank geometry that mitigate this have recently been shown to meet requirements.

7.4 Spontaneous Fission in LZ

7.4.1 Motivation

Some heavy nuclei are energetically unstable and can undergo a process known as spontaneous fission, where the nuclei splits into two fragments. This process occurs with the emission of several neutrons and γ-rays. Fission neutrons can be a major component in the neutron background in several detector materials, for example contributing ~35% of the neutron flux from the titanium cryostat and 62% of the neutron flux from the PMTs, therefore it is important to understand detector response to these events. A LUXSim generator designed to imitate spontaneous fission was written and implemented for use within the LZ detector simulation, with the intention of determining what contribution fission products could make to the LZ background model.
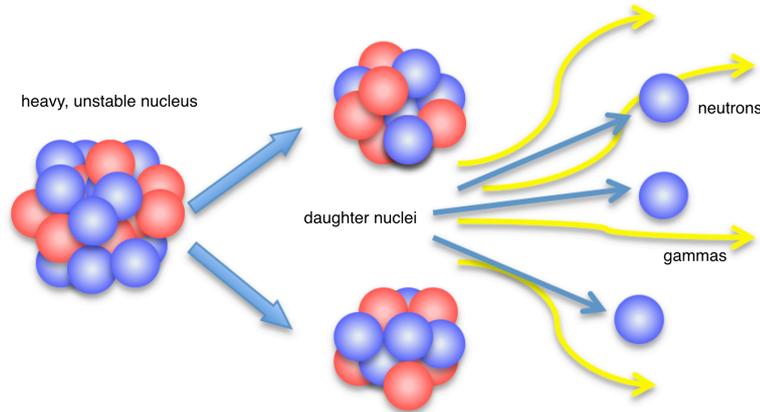
Figure 7.8: Spontaneous fission of a heavy nuclei results in emission of multiple γ-rays and neutrons as well as two daughter nuclei with masses approximately half the parent mass. Usually it is an uneven process, leaving one heavy daughter and one light daughter.

7.4.2 The Fission Process

Spontaneous fission is one of two types of nuclear fission—the other being induced, the type used in nuclear reactors. A spontaneous fission requires no incident particle; a very heavy, energetically unstable nuclei can fission into two daughter fragments, usually unevenly, producing one heavy and one lighter fragment. The spontaneous fission process needs activation through quantum mechanical tunnelling of the fission fragments through a potential barrier that arises from both surface tension and Coulomb repulsion, and advocates use of the liquid drop model of the nucleus. This models the nucleus as a drop of uniform liquid that can be deformed. Enough deformation and a saddle point can occur in the nucleus. Once a critical point is reached, Coulomb forces overcome the surface tension of the 'liquid' and eventually scission occurs, leaving two fragments. Several γ-rays and neutrons may be emitted both at scission and from the fragments.

7.4.2.1 Fission Rates

Spontaneous fission occurs in early uranium and thorium chain elements, see table 7.14 below. This contains data on abundance from SOURCES [150] (atoms per gram assuming 10 ppb concentration, gram per gram), the radioactive decay half life (not used, only for comparison), the spontaneous fission half life, the S.F. activity per cubic centimetre (using $A = \lambda N$) and finally the activity weighted probability of spontaneous fission. Therefore, for each fission event, there is a 99.969% chance it will be a ^{238}U fission, a 0.03% chance it is ^{234}U, etc. At this stage it was

decided to ignore ^{231}Pa and ^{230}Th fissions. They are extremely rare, and no FREYA data was available for them. It is likely it would have been sufficient for the purposes of this study to just simulate ^{238}U fission but all three uranium isotopes were included regardless for completeness.

Isotope	Atoms g^{-1}*	Abundance (by atoms)	Decay t$_{1/2}$(yr)	S.F. t$_{1/2}$ (yr)	Activity (s^{-1}g^{-1})	Probability per fission
^{238}U	2.51×10^{13}	99.30%	4.468×10^9	8.2×10^{15}	6.73×10^{-11}	99.969%
^{235}U	1.85×10^{11}	0.00546%	7.04×10^8	1.0×10^{19}	4.06×10^{-16}	6.02×10^{-4}%
^{234}U	1.38×10^9	0.730%	2.455×10^5	1.5×10^{15}	2.02×10^{-14}	0.030%
^{231}Pa	8.59×10^6	3.34×10^{-5}%	3.276×10^4	2.0×10^{17}	9.44×10^{-19}	1.40×10^{-6}%
^{230}Th	4.24×10^8	0.00168%	7.54×10^4	2.0×10^{18}	4.66×10^{-18}	6.92×10^{-6}%
^{232}Th	2.14×10^{13}	100%	1.40×10^{10}	1.2×10^{21}	4.75×10^{-16}	100%

Table 7.14: Table listing the abundances of fissionable isotopes in the uranium (above the line) and thorium (only ^{232}Th) decay chains, their half lives, activities and the relative probability of fission—i.e. the percentage of total fissions attributable to each isotope. *Assuming 10 ppb natural abundance uranium and thorium.

7.4.2.2 Emission of Neutrons and γ-rays

There are several waves of particle emission during a fission event. Prompt neutrons are emitted from accelerated fragments, within the window 10^{-18} to 10^{-14} seconds after scission. There is some experimental evidence for neutron emission at the point of scission, but the bulk of the neutrons are prompt neutrons from the fragments [151]. Prompt γ emission follows, in the window 10^{-14} to 10^{-3} seconds after fission. 1 ms is quite a significant delay but as this is the size of an LZ event window, it is acceptable to model the γ emission as totally instantaneous with the neutron emission within a simulation, as any energy deposits from the delayed γ-rays would be very likely to lie within the same event/veto window, and so can be used for vetoing.
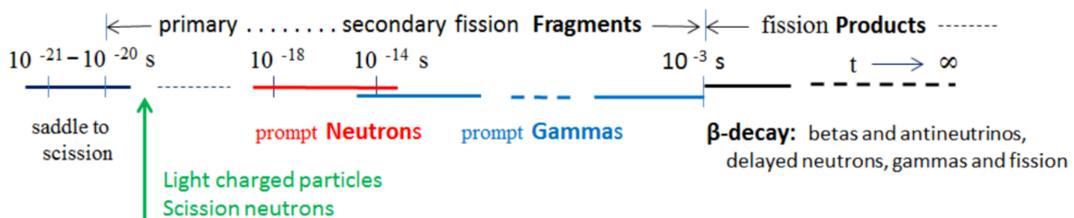

Figure 7.9: Emission timescale of various fission products [151].

The emission of fission products is isotropic in the rest frame of each fission fragment. There is in reality some relativistic boosting in the direction of travel of the fragments, but as they travel

in opposite directions from each other there will, on average, be an equal number of particles boosted in all directions. Running the angular correlation code provided by FREYA for ^{238}U shows for a large number of fissions the distribution of cosines between emitted neutrons is flat. For the purpose of this study, an isotropic distribution will be appropriate for each individual event, as any effects of boosting will be washed out quickly by the attenuation of neutron energies within the detector materials.
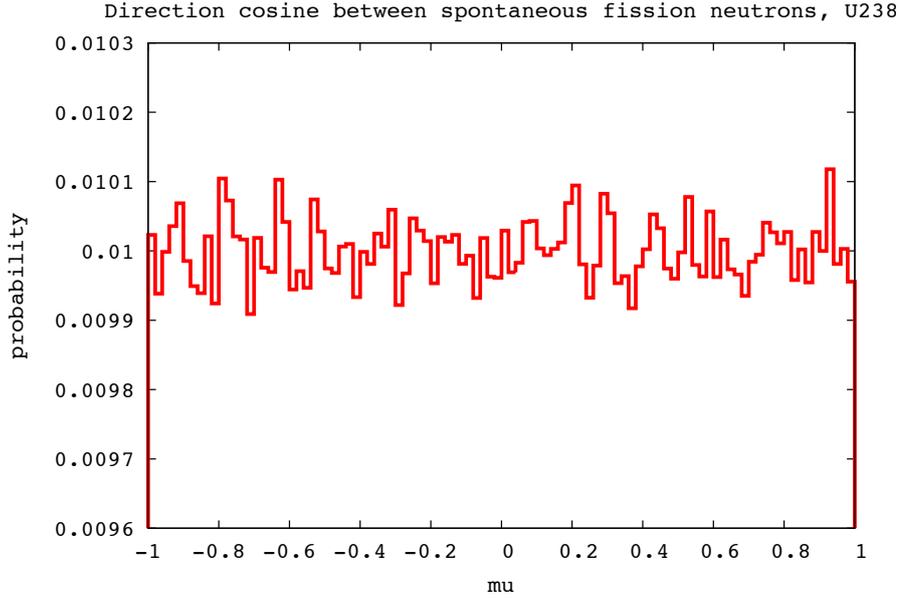

Figure 7.10: Cosine of the angle between spontaneous fission neutrons emitted from ^{238}U spontaneous fission. The angular distribution appears flat.

7.4.2.3 Particle Multiplicities

The fraction of fissions producing each multiplicity of neutrons has been of interest to physicists for decades due to the importance of keeping nuclear reactors in a steady state. In a reactor, fission neutrons can go on to cause further nuclear fissions and are essential for maintaining a reaction, but too many could cause a reactor to go critical, hence the need for neutron absorbing materials such as boron or cadmium rods. Experimental measurements and theoretical calculations have resulted in an accurate distribution function, which varies between isotopes. The probability P_ν of observing ν neutrons from a fission is approximated by a Gaussian-like distribution [152]:

$$\sum_{n=0}^{\nu} P_n = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{(\nu - \bar{\nu} + \frac{1}{2} + b)/\sigma} e^{-\frac{t^2}{2}} dt = \frac{1}{2} + \frac{1}{2} f\left[(\nu - \bar{\nu} + \frac{1}{2} + b)/\sigma\right] \tag{7.15}$$

where $\bar{\nu}$ is the average number of neutrons per fission, σ is the width of the distribution (set to 1.079, which is good for all nuclei except ^{252}Cf), $f(x)$ is a normal probability integral given by:

$$f(x) = (2\pi)^{-\frac{1}{2}} \int_{-x}^{x} e^{\frac{-t}{2}} dt \tag{7.16}$$

201

and b is a small correction factor [153]:

$$b \cong \frac{1}{2} - \frac{1}{2} \oint \left[(\bar{\nu} + \frac{1}{2})/\sigma \right] \tag{7.17}$$

which for all experimental data is < 0.01. An explicit neutron distribution is available for ^{238}U [152], but for the other nuclei this distribution will have been used to calculate the multiplicities within the FREYA code. Figure 7.11 shows the neutron number distribution for ^{238}U, ^{235}U, ^{234}U
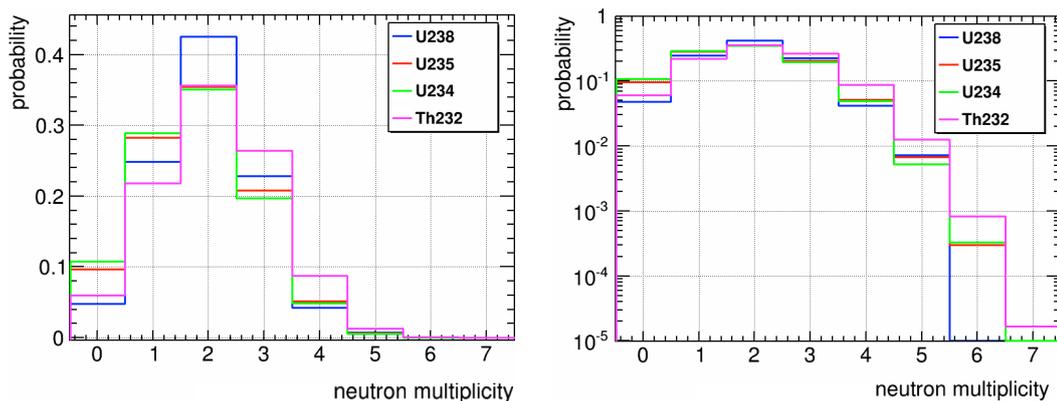

Figure 7.11: Neutron multiplicities for different fissionable isotopes, shown in a log scale on the right.

and ^{232}Th, on both a normal and a log scale to enhance the smaller probabilities at high ν. The probability of emitting G γ-rays follows a negative binomial distribution [152]:

$$\Pi(G) = \binom{\alpha + G - 1}{G} p^G (1-p)^G \tag{7.18}$$

where $p = \frac{\alpha}{\alpha + \bar{G}}$, $\alpha \sim 26$ and $\bar{G}$ is average number of γ-rays per fission. The average number of γ-rays $\bar{G}$ is related to the average number of neutrons $\bar{\nu}$ by:

$$\bar{G} = \frac{(2.51 - 1.13 \times 10^{-5} Z^2 \sqrt{A})\bar{\nu} + 4.0}{-1.33 + 119.6 \frac{Z^{\frac{1}{3}}}{A}} \tag{7.19}$$

where Z and A have their usual meanings as mass and atomic number. The results of applying this distribution to the relevant isotopes are shown in figure 7.12.

The average number of particles emitted for each isotope are given in table 7.15. These

Isotope	Neutrons	γ-rays
^{238}U	2.01	6.44
^{235}U	1.86	6.14
^{234}U	1.81	6.04
^{232}Th	2.14	6.51

Table 7.15: Average multiplicities of neutrons and γ-rays for each fissionable isotope [152].

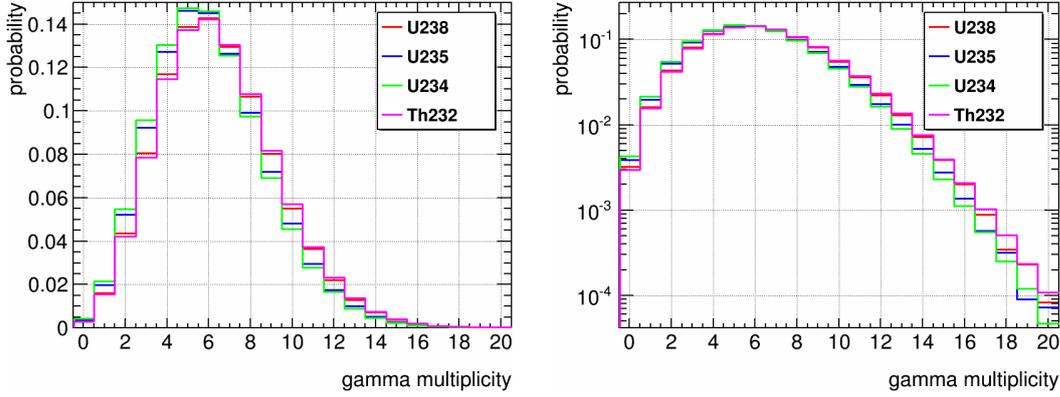demonstrate that 2 emitted neutrons and 6 emitted γ-rays is the most likely scenario for all

Figure 7.12: γ-ray multiplicities for different fissionable isotopes, shown in a log scale on the right.

isotopes considered, but ^{238}U can release as many as 5, ^{235}U and ^{234}U 6 and ^{232}Th 7 neutrons, and there can be as many as 20 γ-rays. All isotopes are also capable of emitting no neutrons and/or no γ-rays. This must be taken into account for the spontaneous fission generator.

As so far spontaneous fission has been treated as single neutron emission, it is reasonable to suggest that the average additional neutron and several additional γ-rays could deposit additional energy that will improve vetoing of the fission event.

7.4.2.4 Particle Energies

The neutron energy distribution follows a Watt spectrum [152]:

$$W(a, b, E') = Ce^{-aE'}\sinh(\sqrt{bE'}) \tag{7.20}$$

where $C = \sqrt{\pi \frac{b}{4a}}\frac{e^{\frac{b}{4a}}}{a}$ and E' is the secondary neutron energy. a and b vary weakly between isotopes. Figure 7.13 shows an example Watt spectrum for induced ^{235}U fission, on a log-log scale.

For γ-rays, where experimental data is sparse (only γ-rays from ^{252}Cf and neutron-induced ^{235}U fission have been measured) a mathematical representation is used [152]:

$$N(E) = \begin{cases} 38.13(E - 0.085)e^{1.648E} & E < 0.3 \text{ MeV} \\ 26.8e^{-2.30E} & 0.3 < E < 1.0 < \text{ MeV} \\ 8.0e^{-1.10E} & 1.0 < E < 8.0 \text{ MeV} \end{cases}$$

which is a fit to the ^{235}U measurements. The representation is plotted in figure 7.14.

FREYA offers options to choose the level of correlation between particle energies and numbers, using different models of energy conservation. However, as there is no data available on the correlations, the default choice in the software is to sample randomly with no correlation or energy conservation. In this work, no correlations are used, based on the assumption that the neutron energy is most important when it reaches the liquid xenon fiducial volume in LZ, by which time it has been significantly reduced and any correlations washed out.
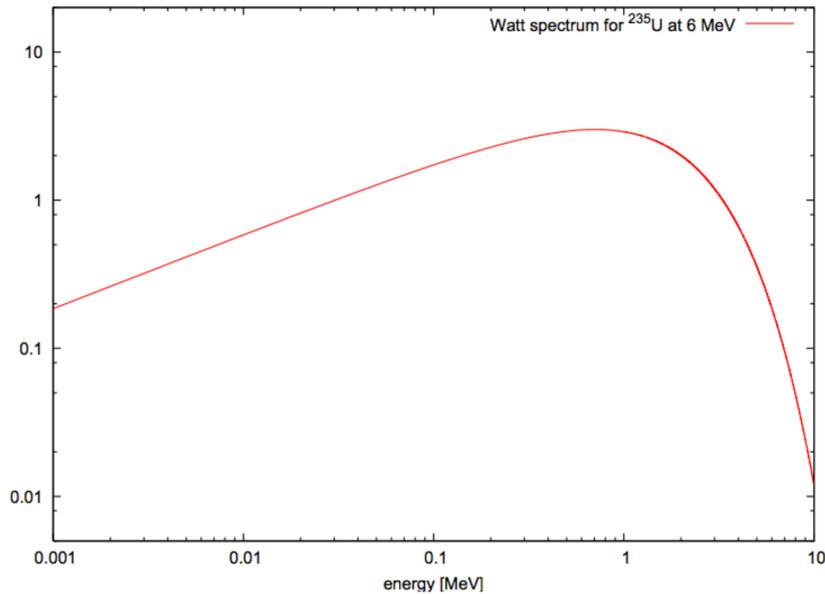
Figure 7.13: Watt spectrum for the energy distribution of neutrons emitted for induced ^{235}U fission.

7.4.3 Neutron Backgrounds in LZ

Neutrons are a dangerous background for LZ due to their capability of mimicking a WIMP signal. If a neutron only scatters once inside the active xenon volume, this will create a single scatter nuclear recoil, just as a WIMP would. Therefore, it is necessary to have a precise understanding of the neutron backgrounds and how many nuclear recoils are expected from them. The main sources of neutrons are spontaneous fission and $(\alpha,$n$)$ reactions in the uranium and thorium decay chains. $(\alpha,$n$)$ reactions are material dependent, and as previously discussed the most prominent sources are expected to be the cryostat and the PMTs. $(\alpha,$n$)$ reactions produce neutrons via interactions of alpha particles from radioactive decay with other nuclei. The process is material dependent and so when modelling backgrounds from these neutrons, the energy spectra calculated require a target atomic number. Currently, LZ background simulations use spectra from SOURCES [150], where the total neutron flux (from both spontaneous fission and $(\alpha,$n$)$ neutrons) is used with one neutron fired at a time. However, this is not realistic for spontaneous fission as it is possible for up to 7 neutrons to be released in one fission event, accompanied by up to 20 γ-rays, depending on the isotope. This raises the question of whether it is possible to veto these events more successfully than single neutrons, suggesting an overestimation of background.

Spontaneous fission neutron energies are completely material independent. This means it is possible to create a spontaneous fission generator that can be placed anywhere in the detector simulation. The generator must produce the correct multiplicities of neutrons and γ-rays as well as using their separate energy spectra.
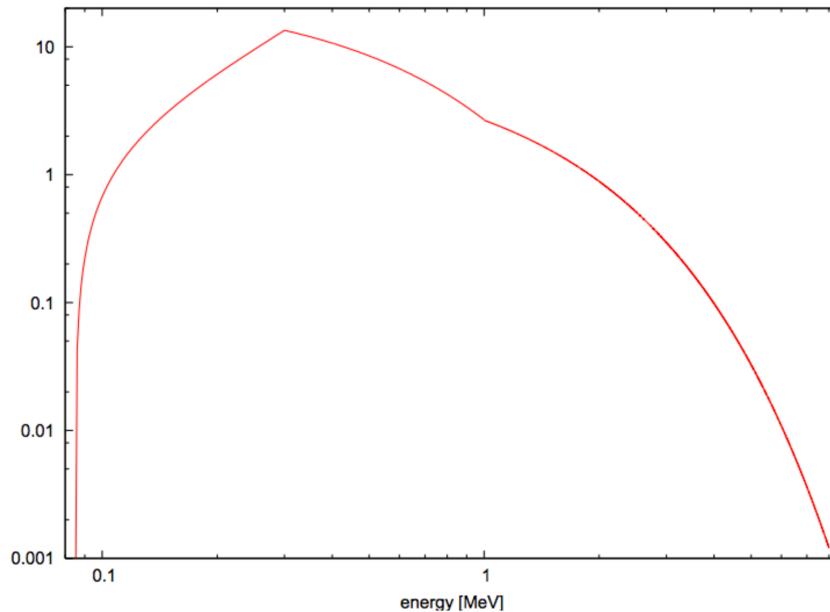
Figure 7.14: γ-ray energy spectrum using a mathematical representation fitted to ^{235}U induced fission data.

7.4.3.1 Neutron yields in different materials

The cryostat and the PMTs are the greatest sources of neutron background in LZ [139]. Table 7.16 shows that for the titanium cryostat, 34.6% of the neutrons are from spontaneous fission. The PMTs are made up of several materials, including quartz (45.9%), aluminium (7.14%), kovar (83.1%), stainless steel and ceramic (14.2%). In the LZ Conceptual Design Review (CDR) backgrounds table the cryostat yields 213.32 neutrons per year, leading to 0.019 NR counts in 1000 days, and the PMTs yield 372.46 neutrons per year, resulting in 0.203 NR counts in 1000 days. It is possible that these numbers could be reduced significantly as they currently assume 0% spontaneous fission rejection. Although table 7.14 and figures 7.11 and 7.12 contain data for thorium, the thorium S.F. half life is so long that (α,n) reactions dominate in all materials (maximum contribution was 1% for brass, stainless steel and Delrin). Therefore, it was decided to only continue work on a uranium generator. Table 7.16 contains the neutron yields for different materials in LZ in neutrons per second per gram per ppb concentration of ^{238}U (n/s/g/ppb). The spontaneous fission yield is of course the same in all materials, but its percentage of the total neutron yield differs.

7.4.4 A Generator for Spontaneous Fission Events

For the generator, the cumulative probability was calculated for each increasing number of primary particles. A random number generator is used with if and if else statements to select

Material	Density	Total	(α,n)	S.F.
	g/cm^{-3}	(n/s/g/ppb)	(n/s/g/ppb)	(n/s/g/ppb)
Titanium	4.51	3.9×10^{-11}	2.55×10^{-11}	1.35×10^{-11} (34.6%)
PTFE	2.20g/cm^3	8.85×10^{-10}	8.72×10^{-10}	1.35×10^{-11} (1.47%)
PEEK	1.32	3.00×10^{-11}	1.65×10^{-11}	1.35×10^{-11} (45.0%)
Quartz	2.65	2.94×10^{-11}	1.59×10^{-11}	1.35×10^{-11} (45.9%)
Aluminium	2.60	1.82×10^{-10}	1.69×10^{-10}	1.35×10^{-11} (7.14%)
Kovar	8.00	1.63×10^{-11}	2.76×10^{-12}	1.35×10^{-11} (83.1%)
Stainless Steel	7.7	1.85×10^{-11}	4.92×10^{-12}	1.35×10^{-11} (74.6%)
Ceramic	4	9.60×10^{-11}	8.24×10^{-11}	1.35×10^{-11} (14.2%)
BaTiO$_3$	6.02	2.18×10^{-11}	8.29×10^{-12}	1.35×10^{-11} (62.0%)
Cirlex	1.42	3.30×10^{-11}	1.95×10^{-11}	1.35×10^{-11} (40.9%)
Brass	8.5	1.38×10^{-11}	3.10×10^{-13}	1.35×10^{-11} (97.8%)
Copper	8.96	1.38×10^{-11}	3.10×10^{-13}	1.35×10^{-11} (97.8%)
PVDF	1.78	6.41×10^{-10}	6.27×10^{-10}	1.35×10^{-11} (2.18%)
Foam	0.08-0.64	3.24×10^{-11}	1.89×10^{-11}	1.35×10^{-11} (41.7%)
Acrylic	1.18	2.65×10^{-11}	1.30×10^{-11}	1.35×10^{-11} (50.9%)
LAB	0.863	2.86×10^{-11}	1.51×10^{-11}	1.35×10^{-11} (1.47%)
Polyethylene	0.925	2.78×10^{-11}	1.43×10^{-11}	1.35×10^{-11} (48.6%)
Solder	7.5	1.35×10^{-11}	0.00	1.35×10^{-11} (100%)
Epoxy	1.0-1.6	3.81×10^{-11}	2.46×10^{-11}	1.35×10^{-11} (35.4%)
Viton	1.80	8.07×10^{-10}	7.94×10^{-10}	1.35×10^{-11} (1.67%)
Rubber	0.90-2.00	2.88×10^{-11}	1.52×10^{-11}	1.35×10^{-11} (46.9%)
ZrO$_2$	5.68	1.62×10^{-11}	2.71×10^{-12}	1.35×10^{-11} (83.3%)
CuBe	8.10-8.25	6.14×10^{-11}	4.78×10^{-11}	1.35×10^{-11} (22.0%)
Borosilicate glass	2.23	1.70×10^{-10}	1.56×10^{-10}	1.35×10^{-11} (5.88%)
Delrin	1.41	2.44×10^{-11}	1.08×10^{-11}	1.35×10^{-11} (55.3%)

Table 7.16: Neutron yields for different materials within the LZ detector. As spontaneous fission is material independent, all have the same rate but can account for a range of percentages of the total yield [139].

which multiplicity is generated and the command */gps/number N* is applied where N is the number of primary neutrons or γ-rays to create. This is done separately with different random numbers for neutrons and γ-rays.

For neutrons, the total spontaneous fission neutron energy spectrum from both uranium chains was taken from SOURCES, as the energy spectra were so similar for different isotopes. For the γ-rays, a weighted average spectrum from FREYA (as FREYA outputs an energy spectrum for each number of γ-rays emitted, although all were very similar) was used for each uranium isotope. The energy distributions were created with the GEANT4 general particle source commands */gps/ene/type Arb* and */gps/hist/type arb*, with individual points set using */gps/hist/point E W* where E is the energy in MeV and W is the weight of the bin. Figure 7.15 shows the spectra used in the generator.
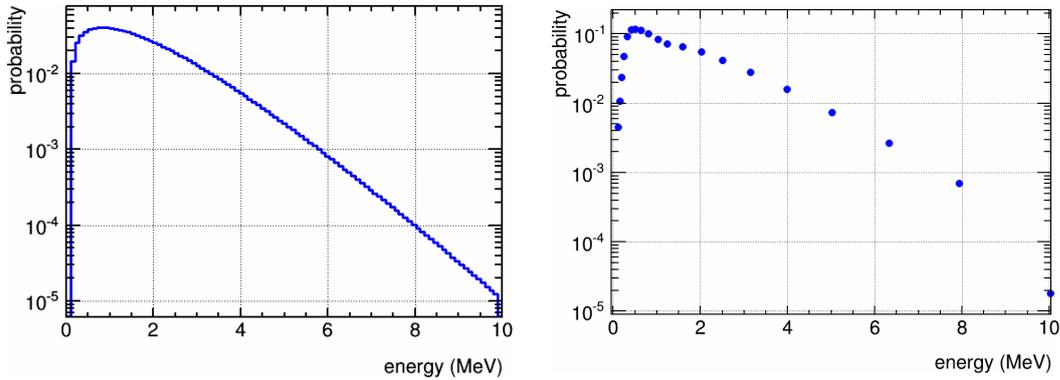
Figure 7.15: Left: Uranium chain neutron energy distribution from SOURCES. Right: Uranium γ-energy distribution from FREYA.

Both the neutrons and γ-rays were given the general particle source command $/gps/ang/type$ iso to obtain isotropic distributions.

7.4.5 Generator Validation

The first step of validation was simply to visualise the events and confirm the generator was behaving as expected. The screenshot shown in figure 7.16 demonstrates that the particles were all emerging from the same vertex simultaneously, with different energies and different directions. In order to confirm the correct energy distribution was being used for neutrons and γ-rays independently, a check was done on simulated data, with *recordLevel* set to 3 in the source volume. For every primary particle, the kinetic energy (*kEnergy_keV*) for its first step was filled into a histogram. The number of primary particles of each type was also recorded to ensure the multiplicity probabilities were working correctly. A check was also done that there was the correct number of events with no neutrons or no γ-rays generated.

As can be seen in figures 7.17 and 7.18, the energy distribution and multiplicities are as expected from the inputs. For the neutron multiplicity, the 6 neutron bin is much lower than it is in figure 7.11, but this is exactly as expected as only about 3 in 10000 fissions are ^{234}U, and less than 1 in 100,000 fissions are ^{235}U, both of which can produce 6 neutrons, whilst ^{238}U does not. Multiplying the original probability of 6 neutrons for ^{234}U from figure 7.11 by the fraction of ^{234}U decays, $0.000333 \times \frac{3}{10000}$ is $\sim 1 \times 10^{-7}$, demonstrating the generator is working correctly.

7.4.6 Simulations

The spontaneous fission generator was written to have the same functionality as other LUXSim generators - it can be placed inside a volume and given an activity using the command $/LUXSim/source/set\ VolumeName\ USF\ Activity\ ActivityUnits$.
For the preliminary studies, it was placed inside the cryostat (volumes *InnerTitaniumVessel* and *OuterTitaniumVessel*) and the PMTs (*Top_PMT_Vacuum* and *Bottom_PMT_Vacuum*). A single

Figure 7.16: Left: tracking output from GEANT4 (edited to only show primary particles). Right: basic visualisation of the event vertex. The red, blue and white lines are axes marking the vertex.

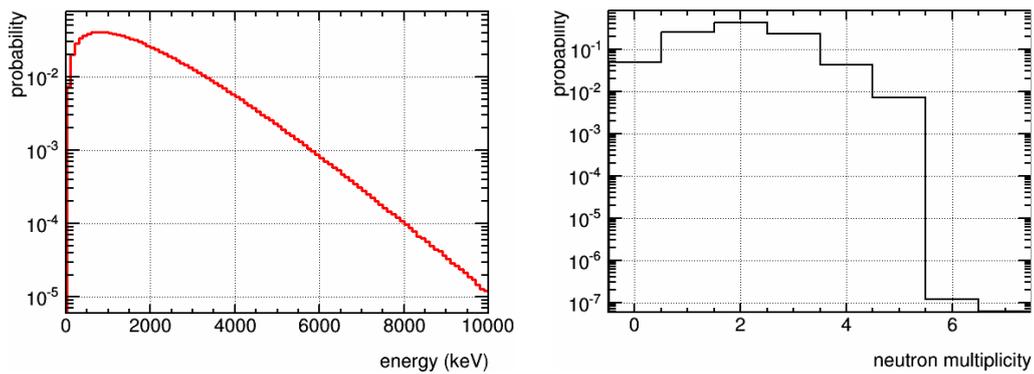Figure 7.17: Left: Primary neutron initial kinetic energy, scaled to number of generated events. Right: Neutron multiplicities obtained by counting the number of primary neutrons and scaling by total events.

neutron source using the LUXSim generator *SingleParticle* and the same energy distribution as the spontaneous fission neutrons was used in a separate run for comparison to the spontaneous fission generator.
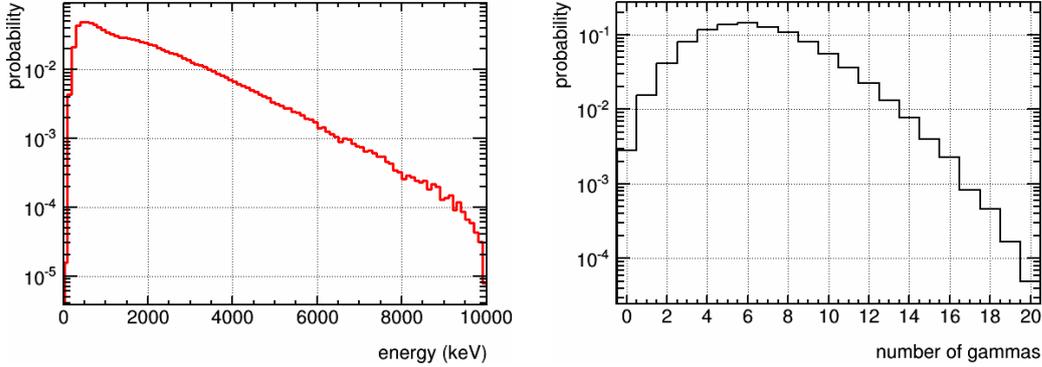
Figure 7.18: Left: Primary γ initial kinetic energy, scaled to number of generated events. Right: γ multiplicities obtained by counting the number of primary neutrons and scaling by total events.

The livetime estimation was calculated using:

$$L(days) = \frac{N_{gen}}{A_c B_{sf} M_c T_{day}} \tag{7.21}$$

where N_{gen} is the number of generated events (i.e. the number given to $/LUXSim/beamOn$), A_c is the measured component activity in Bq/kg, B_{sf} is the branching fraction of spontaneous fission, M_c is the mass of the source component and T_{day} is the number of seconds in a day. For the following, the cryostat mass used is M_c = 1571.152 kg and the PMT vacuum as 0.4kg.

7.4.7 Results

Table 7.17: Surviving counts after each cut, with the percentage of all NR rejected given underneath, for events generated in the cryostat (top two rows) and the PMTs (bottom two rows). *Results scaled up from 1.5×10^7 generated events for easier comparison.

N_{gen}	All NR	+ FV	+ SS	+ Skin	+ OD	all	Total
S.F.	2,790,092	1,387,540	22,756	9,724	45	27	99.9999%
2×10^{7}*	13.95% of N_{gen}	50.27%	99.18%	99.65%	99.998%	99.999%	~1 in 10^6
Neutrons	1,745,909	823,852	17,143	10,632	1,680	923	99.995%
2×10^{7}	8.73% of N_{gen}	52.81%	99.02%	99.39%	99.90%	99.95%	~2 in 10^5
S.F.	5,606,942	1,734,881	37,845	16,996	381	29	99.9999%
2×10^{7}	28.03% of N_{gen}	69.06%	99.33%	99.70%	99.993%	99.999%	~1 in 10^6
Neutrons	4,853,552	1,324,342	25,822	16,958	2,848	1,595	99.992%
2×10^{7}	24.27% of N_{gen}	72.72%	99.47%	99.65%	99.94%	99.97%	~8 in 10^5

Table 7.17 clearly demonstrates the difference in using the spontaneous fission and single neutrons, namely that vetoing by the outer detector becomes significantly more effective. Simulations were done for events generating neutrons in both the cryostat and the PMTs, as they are the two most dominant sources of neutrons. There are overall 1.6× more nuclear recoil events for the spontaneous fission generator, in line with what you might expect when the average number of

neutrons per event is 2. The effect of the fiducial volume on these recoils is similar for both cases, as should be the case. The single scatter cut is also similarly effective, and so is the LXe skin veto. The main improvement comes from the outer detector, which becomes 0.1% more effective, leaving ~40× less events passing.

Figure 7.19 is a visual representation of the numbers in table 7.17, showing the surviving background from all NR at each step, and figure 7.20 shows the effect of cutting away from the walls. Figures 7.21, 7.22, 7.23 and 7.24 show r^2, z maps of the energy deposits remaining after each cut. The subcathode region contains the majority of the final unvetoed events for both the spontaneous fission generator and single neutrons but there are considerably less events left for the spontaneous fission generator once the outer detector is applied. This makes sense when considering the topology of a spontaneous fission event. Whilst a single neutron may travel into the xenon and deposit energy, a spontaneous fission event is releasing on average 2 neutrons and 6 γ-rays simultaneously and isotropically. If one neutron is making it to the liquid xenon, another one is likely to be heading outwards towards the outer detector.

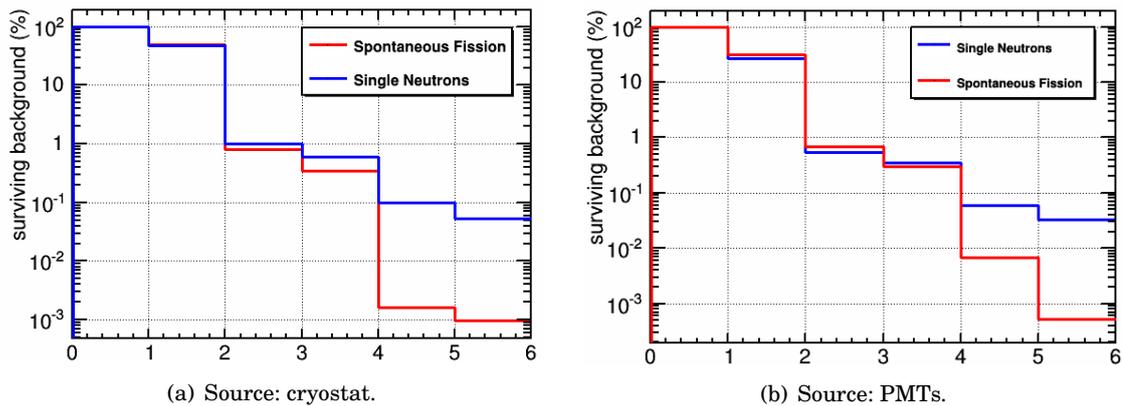(a) Source: cryostat.   (b) Source: PMTs.

Figure 7.19: Content of each bin, left to right: 1. All NR 6–30 keV, 2. as (1) in fiducial volume, 3. as (2) with single hit cut, 4. as (3) with application of LXe skin veto, 5. as (3) with application of outer detector veto, 6. as (3) with application of both vetoes

7.4.8 The Backgrounds Table

In order to investigate the effects of more efficient vetoing on the LZ backgrounds table, the 99.9999% efficiency was applied to the correct proportion of emitted neutrons, calculated for each material, for each background component. The results are shown in table 7.18 and demonstrate that if the same rejection efficiency of 99.9999% is used on every component, the number of counts can be reduced by a factor of 46%. If the NR counts are scaled in the same way, 0.325 expected counts are reduced to 0.157. Whether or not the rejection efficiency will be the same for all components is not certain, but considering the positions of both an internal component (PMTs) and external (cryostat) it is not expected to worsen anywhere else in the detector. Furthermore,

(a) Source: spontaneous fission in the cryostat.

(b) Source: single neutrons in the cryostat.

(c) Source: spontaneous fission in the PMTs.
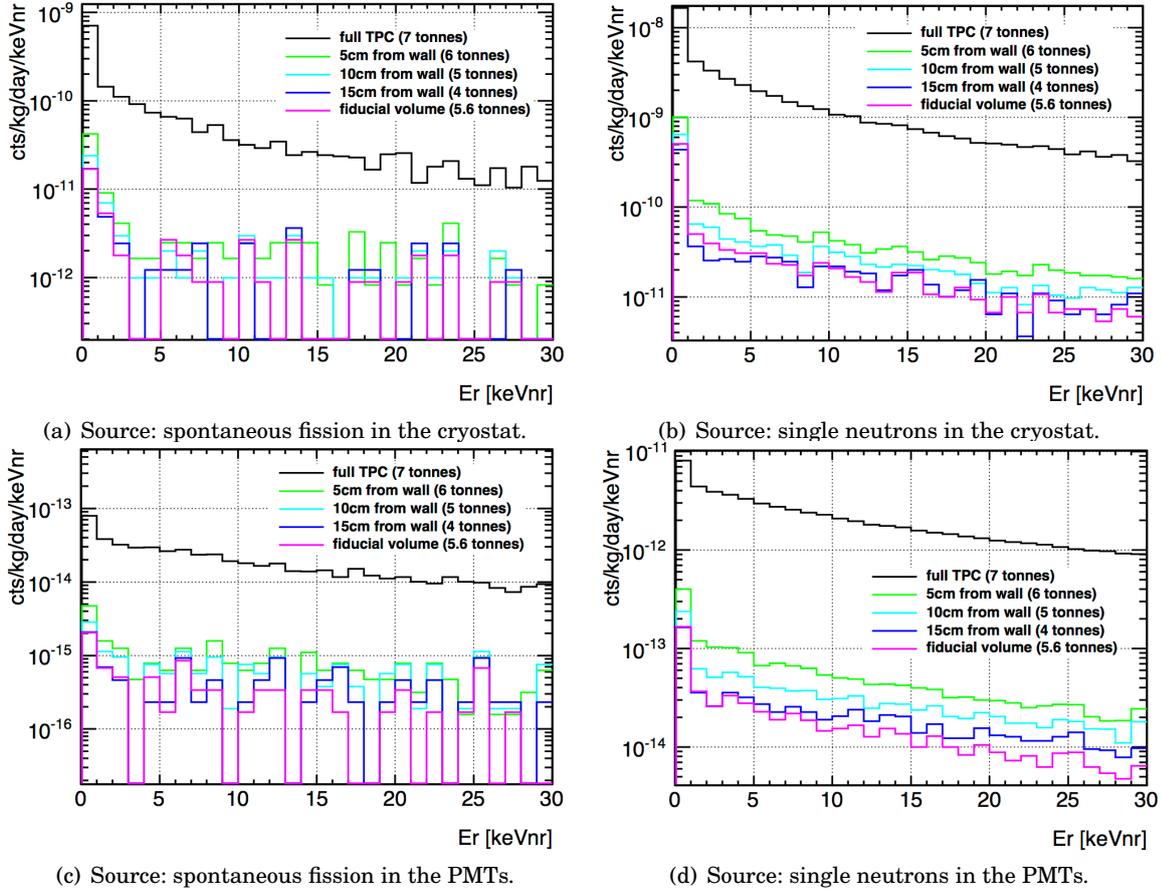
(d) Source: single neutrons in the PMTs.

Figure 7.20: Energy deposited by nuclear recoils in the full TPC (black), 5 cm from the wall (green), 10cm from the wall (cyan), 15cm from the wall (blue) and in the fiducial volume (magenta). The single scatter cut and both vetoes are applied. Events were generated in the cryostat.

they are the two biggest sources of neutrons and so it is reassuring to see such a high level of vetoing efficiency.

7.4.9 Conclusion

It has been demonstrated by simulating both neutrons and γ-rays from spontaneous fission events with the correct multiplicity and energy distribution that the LZ detector should be able to veto these events more efficiently than for single neutrons, which has been the treatment of spontaneous fission backgrounds up to this point. The gain is attributable to the outer detector, which vetoes over an order of magnitude more events. For both the dominant neutron sources in LZ, the cryostat and the PMTs, a spontaneous fission vetoing efficiency of 99.9999% has been obtained. This vetoing is considered good enough that neutrons from spontaneous fission have been excluded from the LZ background model.

Table 7.18: Neutrons per year and NR counts before and after spontaneous fission rejection. The number of neutron emissions per year were from the LZ CDR [139], and have since changed due to new screening results and design changes, but are still shown here to demonstrate the impact of spontaneous fission rejection.

Component	n/yr	NR (cts)	n/yr (after s.f.r.)	NR (cts) (after s.f.r.)
Upper PMT Structure	3.96	0.002	2.67	0.001
Lower PMT Structure	5.493	0.003	4.31	0.002
R11410 3" PMTs	372.46	0.203	141.22	0.077
R11410 PMT Bases	76.73	0.033	56.44	0.024
R8520 Skin 1" PMTs	11.44	0.002	6.97	0.001
R8520 Skin PMT Bases	23.27	0.003	17.27	0.002
PMT Cabling	89.5	0.008	1.97	0
TPC PTFE	24.13	0.007	23.78	0.007
Grid Wires	0.02	0	0.01	0
Grid Holders	6.92	0.003	4.53	0.002
Field Shaping Rings	32.2	0.004	22.34	0.003
TPC Sensors	0.72	0	0.42	0
TPC Thermometers	85.24	0.01	72.88	0.009
Xe Recirculation Tubing	0.37	0	0.36	0
HV Conduits and Cables	15.6	0.001	7.102	0
HX and PMT Conduits	11.91	0	3.86	0
Cryostat Vessel	213.32	0.019	139.51	0.012
Cryostat Seals	40.28	0.001	39.50	0.001
Cryostat Insulation	85.2	0.003	67.46	0.002
Cryostat Teflon Liner	4.97	0	4.90	0
Outer Detector Tanks	101.12	0.0002	59.43	0
Liquid Scintillator	22.94	0	22.60	0
Outer Detector PMTs	20,852	0.022	11,281	0.012
Outer Detector PMT Supports	37.01	0	9.40	0
Total	22,116.80	0.3252	11,989.83	0.157
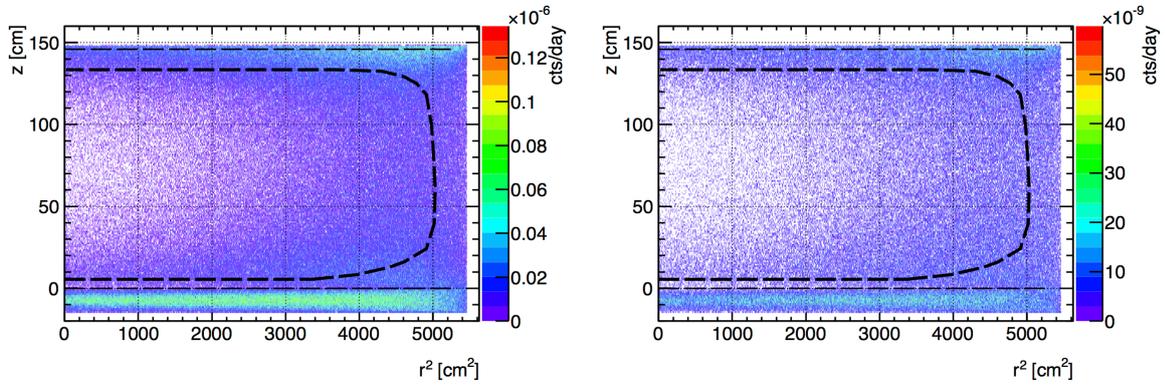
(a) Spontaneous fission, all events 6–30 keV
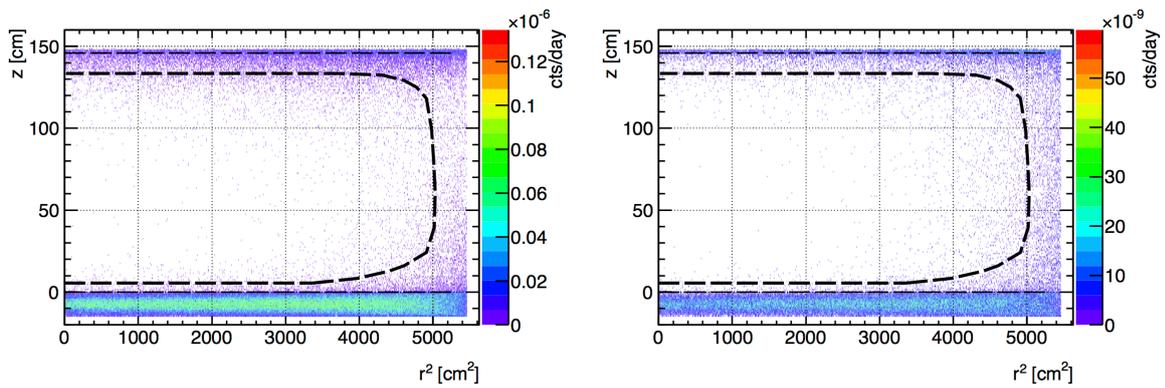
(b) Single neutrons, all events 6–30 keV

(c) Spontaneous fission, single scatter events

(d) Single neutrons, single scatter events

(e) Spontaneous fission, LXe skin veto applied

(f) single neutrons, LXe skin veto applied

Figure 7.21: Source: cryostat. Map of all single hit NR deposits 6–30 keV in r^2 vs z, showing the effect of applying vetoes.

(a) Spontaneous fission, scintillator veto applied.

(b) Single neutrons, scintillator veto applied.

(c) Spontaneous fission, all vetoes

(d) Single neutrons, all vetoes

Figure 7.22: Source: cryostat. Map of all single hit NR deposits 6–30 keV in r^2 vs z, showing the effect of applying vetoes.
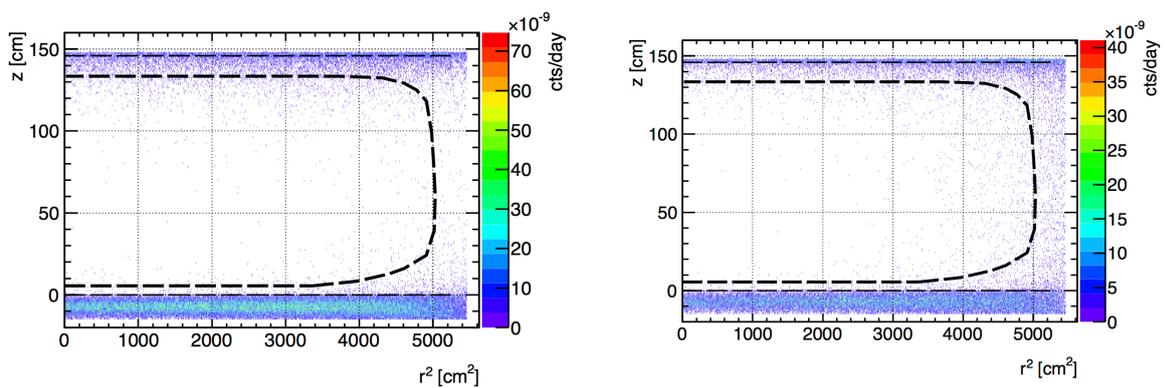
(a) Spontaneous fission, all events 6-30 keV

(b) Single neutrons, all events 6-30 keV

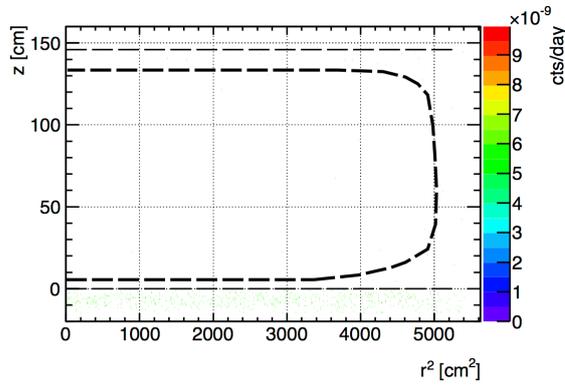(c) Spontaneous fission, single scatter events

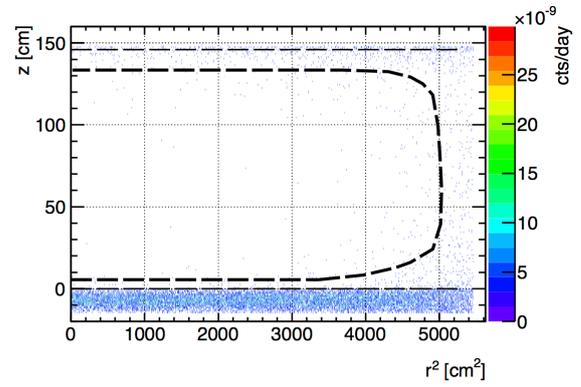(d) Single neutrons, single scatter events

(e) Spontaneous fission, LXe skin veto applied

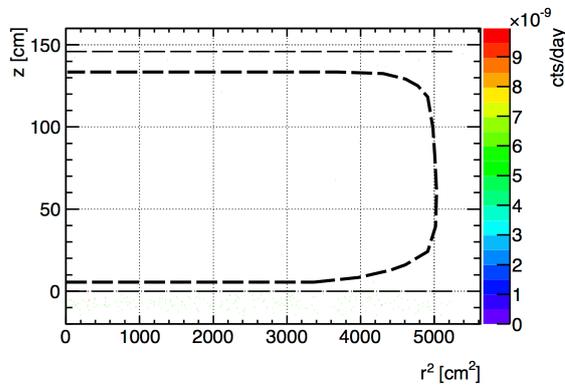(f) Single neutrons, LXe skin veto applied

Figure 7.23: Source: PMTs. Map of all single hit NR deposits 6-30 keV in r^2 vs z, showing the effect of applying vetoes.

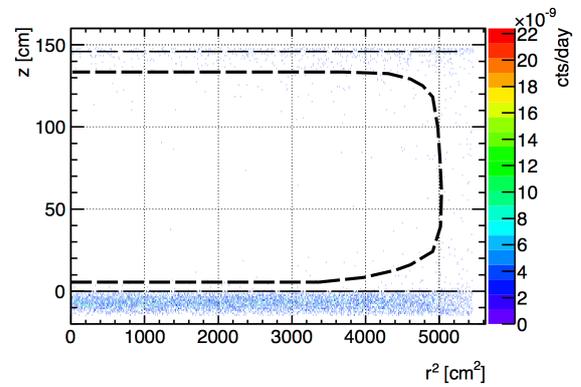(a) Spontaneous fission, OD veto applied.

(b) Single neutrons, OD veto applied.

(c) Spontaneous fission, all vetoes

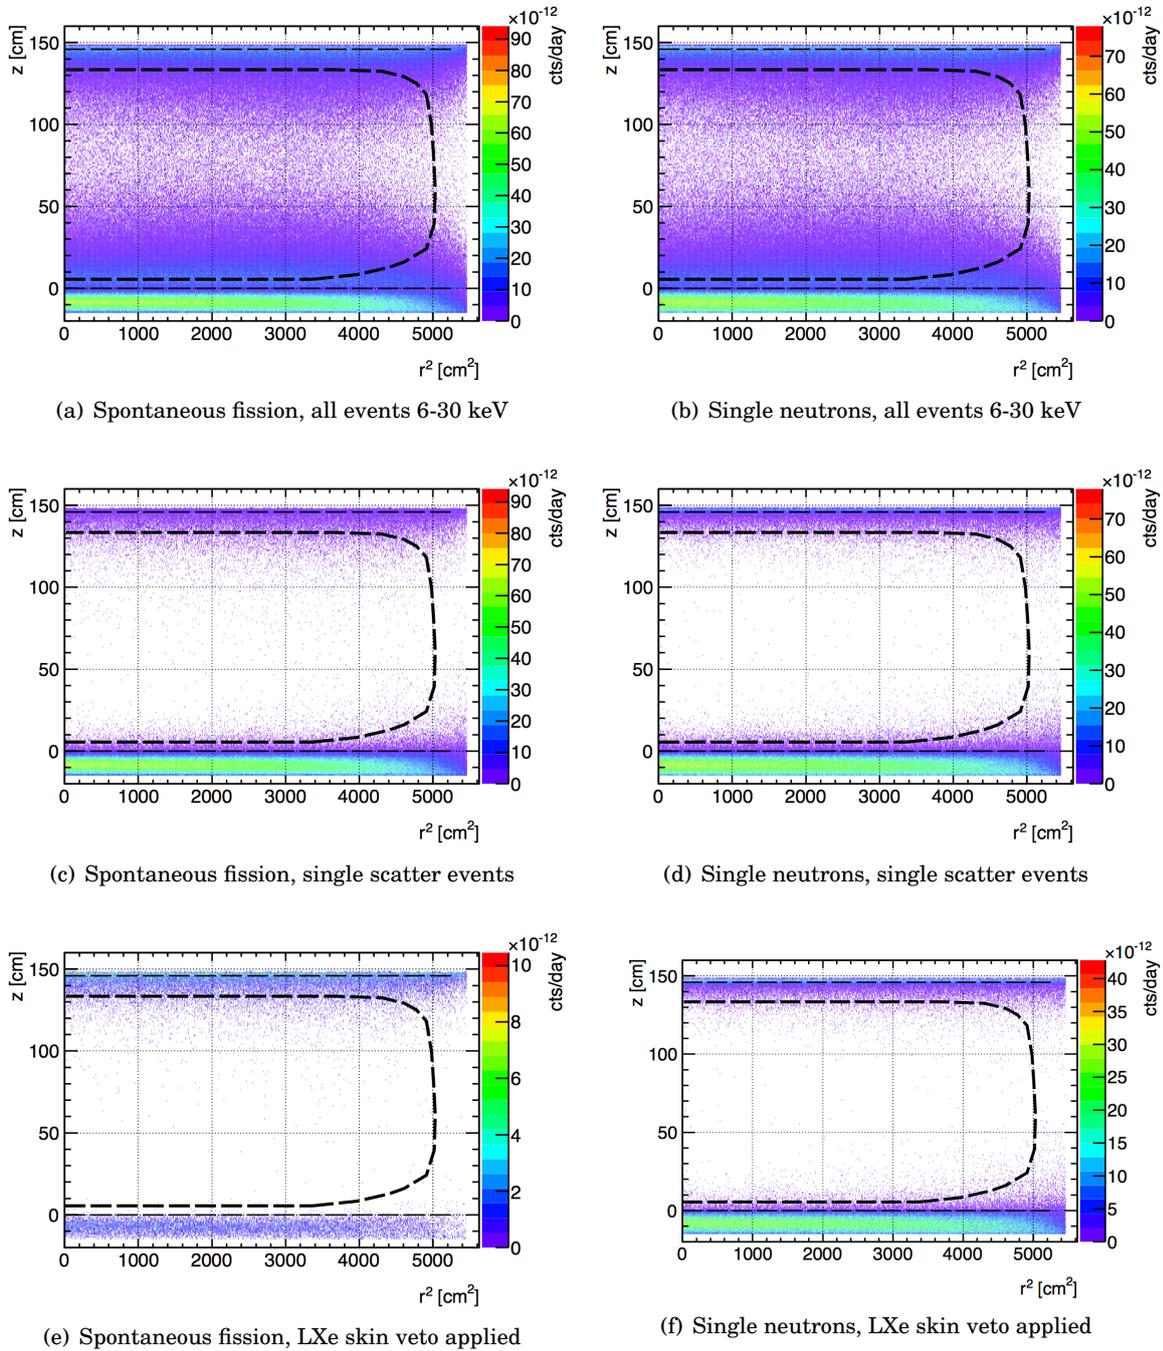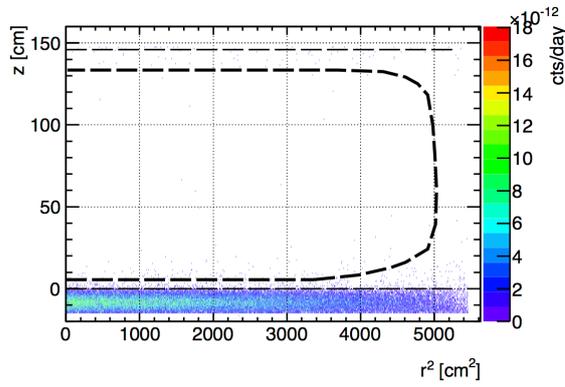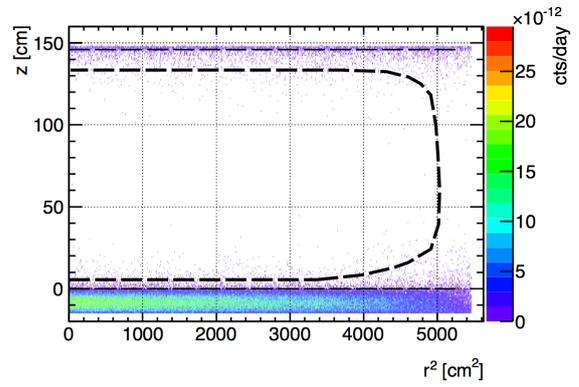(d) Single neutrons, all vetoes

Figure 7.24: Source: PMTs. Map of all single hit NR deposits 6–30 keV, in r^2 vs z, showing the effect of applying vetoes.
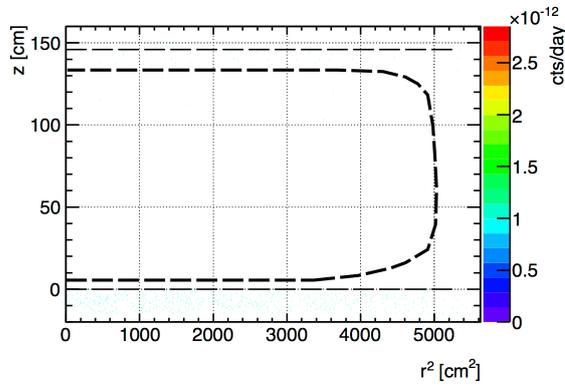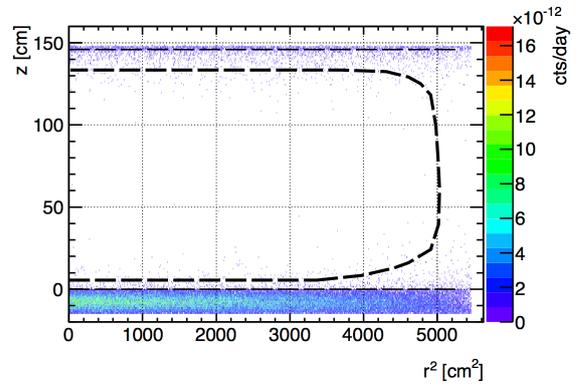
7.5 LZ Sensitivity

The sensitivity of LZ to WIMPs has been analysed using the background simulations detailed above alongside probability density functions (PDFs) for neutrino interactions and WIMP interactions. Differential energy spectra of WIMPs are calculated and turned into a PDF using [154]. The PLR is performed similarly to LUX; signal and background profiles in S1 and S2 are created, but so far, the spatial position of individual events is not used. Instead, a fiducial cut is applied on simulated data. The backgrounds are broken into 11 individual components, which are listed in table 7.19. The detector, environmental background and radon PDFs are obtained from the simulations described in the previous sections. Uncertainties, incorporated as nuisance parameters, are obtained from aspects such as neutrino flux uncertainties and radioassay uncertainties.

Table 7.19: Expected background counts in 5,600 tonne-days for LZ from various sources, along with percentage uncertainties. The final two rows are obtained from the simulations detailed previously in this chapter [155]. **7 ± 3 events expected, but these are at very low energy.

Background	Type	Counts	Uncertainty
^{8}B	NR	0**	10%
HEP	NR	0.21	30%
DSN	NR	0.05	50%
ATM	NR	0.46	33%
PP + ^{7}Be + ^{14}N	ER	255	1%
^{85}Kr	ER	24.5	5%
$2\nu\beta\beta$	ER	67	7%
^{222}Rn	ER	720	10%
^{220}Rn	ER	122	10%
Det. + Env.	ER	43.4	10%
Det. + Env.	NR	0.57	10%

For each WIMP mass, a signal PDF is generated by converting the differential energy spectrum to S1 and S2 via NEST. Figure 7.25 demonstrates the PDFs for 10 GeV, 40 GeV and 1000 GeV WIMPs, overlaid on the main ER and ^{8}B ν profiles. The PLR method is significantly better for background discrimination than using a cut and count method. Using the PLR, LZ can still meet sensitivity requirements in the presence of all the backgrounds listed in table 7.19. Figure 7.26 shows the limit projection for LZ assuming baseline requirements, which reaches a minimum cross section of 2.5×10^{-48} cm^2 at 40 GeV, and the goal limit, which reaches 1.3×10^{-48}. Additionally, the number of ^{8}B ν-A scattering events detected at low energy increases from 7 to 300 when going from baseline to goal, as we begin to cut into the neutrino floor.

The impact of various factors on the limit projection have also been assessed. Radon levels, atmospheric neutrino counts, light collection efficiency, the N-fold coincidence requirement for S1s, the electron extraction efficiency, the electron livetime, drift field and the exposure time have all been altered individually to levels below and above the baseline in order to assess the effect on the limit. It is found that LZ can still reach its goal sensitivity of 3×10^{-48} cm^2 at 40 GeV for

(a) 10 GeV WIMP.   (b) 40 GeV WIMP.   (c) 1000 GeV WIMP.
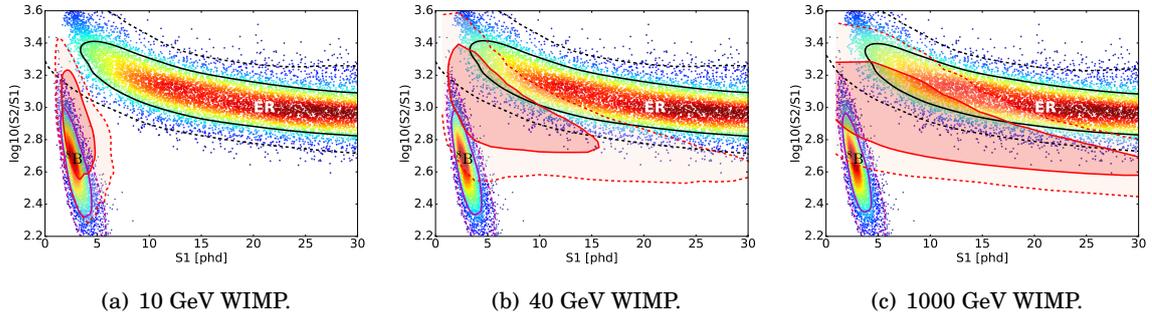
Figure 7.25: PDFs for the ER and ^{8}B backgrounds in the S1 vs log(S2/S1) discrimination space, overlaid with the expected distributions (red) for different WIMP masses. The darker (lighter) red shows the 1σ (2σ) region. Note the similarity of a 10 GeV WIMP with the ^{8}B distribution, and that for higher WIMP masses the signal regions overlap with the ER background region [155].
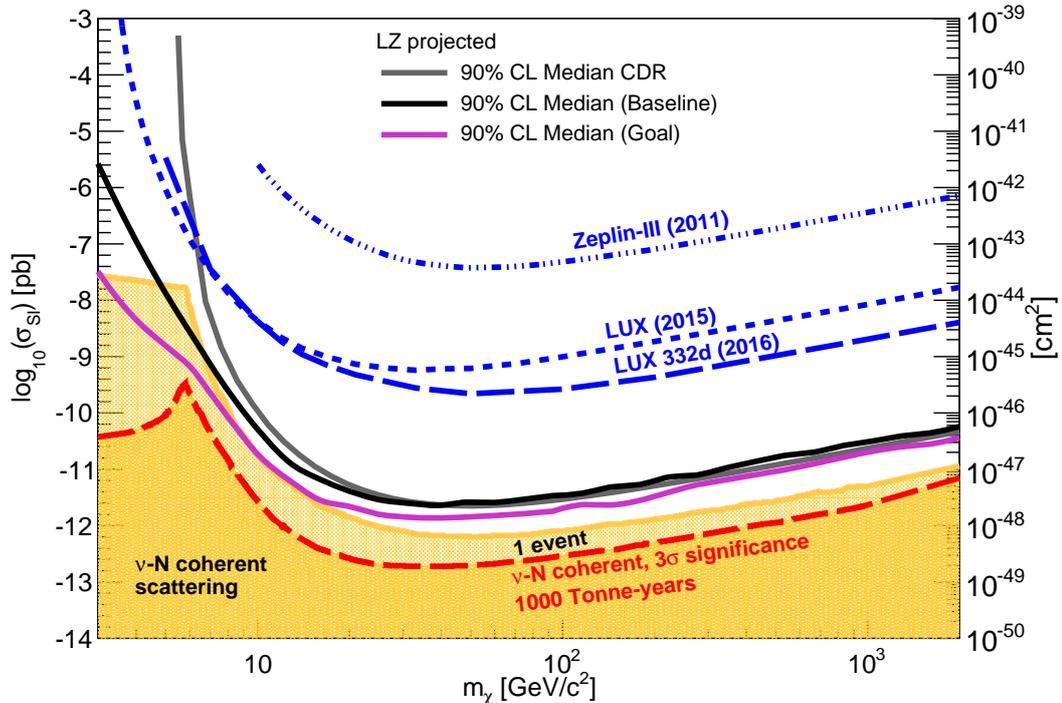
Figure 7.26: LZ sensitivity projection limit curve in the usual WIMP mass vs SI cross section phase space, assuming baseline requirements is met. Here, the limit for the LZ Conceptual Design Report (CDR) is shown in grey, and does not include the lower L_y cut-off obtained from the D-D calibration. The black and purple lines show the baseline and goal limit projections [155].

most eventualities, allowing backgrounds to be higher than baseline. Figure 7.27 demonstrates the reduced, baseline and goal sensitivities, of which the parameters can be found in table 7.20. The reduced case is very pessimistic and does not reach requirement, but it is extremely unlikely that every parameter will turn out to be at the least desirable value of its possible range.

The background simulations have provided important input to limit setting—only the neutrino

Table 7.20: Parameters used in LZ sensitivity projections [155].

Parameter	Reduced	Baseline	Goal
Photon detection efficiency	0.05	0.075	.012
Drift field (V/cm)	160	210	650
Electron lifetime (μs)	850	850	2800
e^- extraction efficiency	50%	95%	99%
N-Fold PMT coincidence	4	3	2
Livedays	1000	1000	1000
^{222}Rn (mBq in active LXe)	13.4	13.4	0.67

counts have had no input from simulation. This motivates the extensive improvements and validation detailed previously in this chapter.

Currently, there are plans to alter the PLR framework to include WIMP mass as a parameter of interest (POI). This would allow construction of a confidence region in the case of a WIMP discovery. Additionally, Background PDFs will continue to be updated and improved as aspects such as screening measurements and Monte Carlo results are updated, position information (e.g. r, z) will be included as observables, and LUX WIMP search and calibration data will be used to inform ER and NR modelling with NEST.
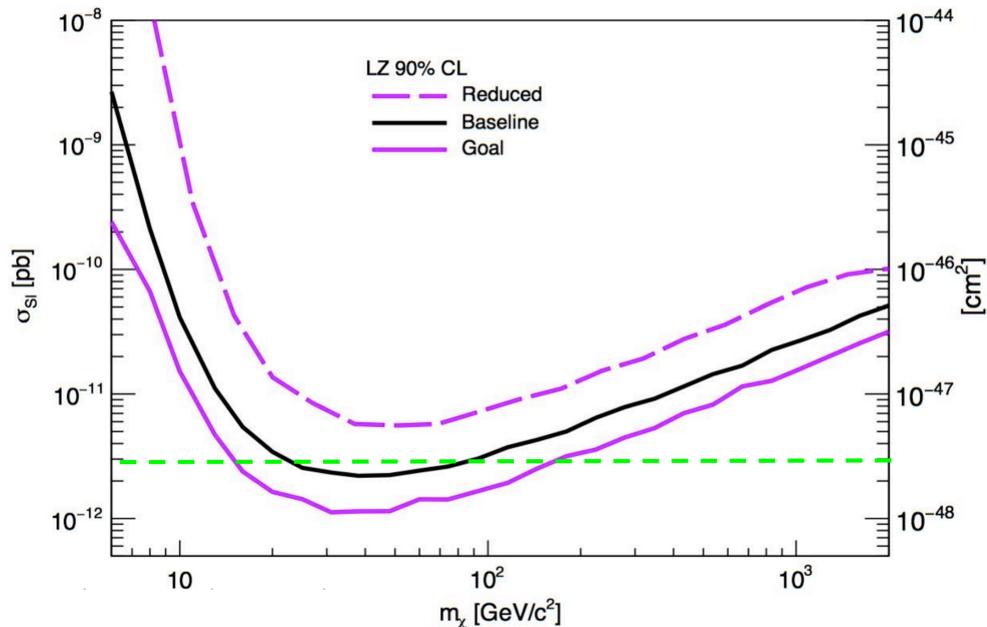

Figure 7.27: LZ sensitivity projections for reduced, baseline and goal cases. The requirement sensitivity of reaching a cross section of 3×10^{-48} is shown as a dashed green line [155].

CONCLUSIONS

This thesis has presented the contribution of the author to the LUX and LZ dark matter searches. LUX's world-leading dark matter constraints and new measurements of liquid xenon properties have been reported, including spin-dependent and spin-independent WIMP-nucleon scattering cross sections and measurements of L_y and Q_y, the scintillation and ionisation yields of liquid xenon. LUX has measured a spin-independent WIMP-nucleon cross section upper limit of 2.2×10^{-46} cm^2 at a WIMP mass of 50 GeV/c^2 and showed the region 3.3 - 5.2 GeV/c^2 to be detectable in liquid xenon for the first time. L_y has been measured at 1.1 keV and Q_y at 0.7 keV with a novel neutron calibration technique; these are lower energies than ever previously achieved, and there is ongoing work to lower them further. Such low energy signal identification was only achievable with high quality, well validated data processing algorithms. A parameterisation of the signal efficiency of these algorithms has also been reported and was essential for limit setting.

Whilst LUX is now being decommissioned, there is still much to be done with its data. For example, there is work on an ultra-low energy calibration of LUX using ^{127}Xe electron capture, and a search for two-neutrino double electron capture on ^{124}Xe. Many of the analyses that were performed on run 3 data will now benefit from the extra livetime of run 4. Furthermore, we have learnt many things from LUX in terms of construction, commissioning, calibration and day-to-day operation, all of which will be transferable to LZ, which is currently under construction. LZ will begin taking data in 2019; to reach its design sensitivity, its backgrounds must be extremely low and well characterised. This thesis has also presented the design and implementation of a simulated background model, along with its results which were included in LZ's Technical Design Review. Work on LZ continues; the background model will be continuously improved and refined to reflect changes in design and to incorporate new screening results as they are attained.

LZ is a generation II dark matter search that builds upon its predecessors, LUX and ZEPLIN; its sensitivity is so high that it will detect the previously unobserved neutrino-nucleon coherent scattering process. We can only hope that this will be the generation that reports a dark matter detection, as we approach this possibly unsurpassable neutrino floor, and that the third generation of detectors may be used to study the properties of WIMPs in greater detail. LZ will probe favoured theoretical phase-space, and in the case of a detection, will require strong confirmation from other experiments of the same direct-detection nature as well as from indirect and collider experiments. In the case of a null-detection, it will constrain electroweak parameter space, ruling out potential models and perhaps motivating new ones. The entire dark matter community has awaited such a detection for over 20 years, and it remains to be seen just how long they must continue to wait.

[1]   F. Zwicky, "Die Rotverschiebung von extragalaktischen Nebeln," Helv. Phys. Acta **6**, 110 (1933).

[2]   J. H. Oort, "Some Problems Concerning the Structure and Dynamics of the Galactic System and the Elliptical Nebulae NGC 3115 and 4494," Astrophys. J. **91**, 273 (1940).

[3]   V. C. Rubin and W. K. Ford, "Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions," Astrophys. J. **159**, 379 (1970).

[4]   E. Karukes, P. Salucci, and G. Gentile, "The Dark Matter Distribution in the Spiral NGC 3198 out to 0.22 R_{vir}," Astronom. and Astrophys. **578**, 13 (2015).

[5]   M. Milgrom, "A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis," Astrophys. J **270**, 365 (1983).

[6]   A. Aguirre, J. Schaye, and E. Quataert, "Problems for MOND in Clusters and the Ly-α Forest," Astrophys. J. **561**, 550 (2001).

[7]   F. J. Sánchez-Salcedo, E. Martinez-Gómex, V. M. Aguirre-Torres, and H. M. Hernández-Toledo, "Low-mass disc galaxies and the issue of stability: MOND vs dark matter," MNRAS (2016).

[8]   F. Iocco, M. Pato, and G. Bertone, "Testing modified Newtonian dynamics in the Milky Way," Phys. Rev. D **92**, 084046 (2015).

[9]   P. Ade *et al.* (Planck), "Planck 2013 results. I. Overview of products and scientific results," Astron. Astrophys. **571**, A1 (2013).

[10] P. Ade *et al.* (Planck), "Planck 2015 results. XIII. Cosmological parameters," (2015), arXiv:1502.01589.

[11] D. J. Eisenstein *et al.*, "Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies," Astrophys. J. **633**, 560 (2004).

[12] W. J. Percival *et al.*, "Measuring the Baryon Acoustic Oscillation scale using the SDSS and 2dFGRS," Mon. Not. R. Astron. Soc **381**, 1053 (2007).

[13] D. Harvey, R. Massey, T. Kitching, A. Taylor, and E. Tittley, "The non-gravitational interactions of dark matter in colliding galaxy clusters," Science **347**, 1462 (2015).

[14] D. Clowe *et al.*, "A direct empirical proof of the existence of dark matter," Astrophys. J. **648**, L109 (2006).

[15] M. Markevitch, "Chandra observation of the most interesting cluster in the universe," ESA Spec. Publ. **604**, 723 (2005).

[16] E. J. Tollerud, J. S. Bullock, L. E. Strigari, and B. Willman, "Hundreds of Milky Way Satellites? Luminosity Bias in the Satellite Luminosity Function," Astrophys. J. **688**, 277 (2008).

[17] A. M. Brooks, M. Kuhlen, A. Zolotov, and D. Hooper, "A Baryonic Solution to the Missing Satellites Problem," Astrophys. J. **765**, 22 (2013).

[18] J. Penarrubia, A. J. Benson, M. G. Walker, G. Gilmore, A. McConnachie, and L. Mayer, "The impact of dark matter cusps and cores on the satellite galaxy population around spiral galaxies," Mon. Not. R. Astron. Soc. **406**, 1290 (2010).

[19] M. Boylan-Kolchin, J. S. Bullock, and M. Kaplinghat, "The Core-Cusp Problem," Mon. Not. Roy. Astron. Soc. **415**, L40 (2011).

[20] W. J. de Blok, "The Core-Cusp Problem," Adv. Astro. **2010**, 789293 (2009).

[21] D. Cerdeno, "Dark Matter 101: from production to detection," (2016), Lecture series, Higgs Centre School of Theoretical Physics.

[22] J. L. Feng, "Dark Matter Candidates from Particle Physics and Methods of Detection," Annu. Rev. Astron. Astrophys. **48**, 495 (2010).

[23] J. F. Navarro, C. S. Frenk, and S. D. M. White, "The Structure of Cold Dark Matter Halos," Astrophys. J. **462**, 563 (1997).

[24] G. Jungman, M. Kamionkowski, and K. Griest, "Supersymmetric Dark Matter," Phys. Rept. **267**, 195 (1996).

[25] M. Beltran, D. Hooper, E. W. Kolb, and Z. A. C. Krusberg, "Deducing the nature of dark matter from direct and indirect detection experiments in the absence of collider signatures of new physics," Phys. Rev. D. **80**, 043509 (2008).

[26] V. Chepel and H. Araújo, "Liquid Noble Gas Detectors for Low Energy Particle Physics," JINST **8**, R04001 (2013).

[27] M. Aaboud *et al.* (ATLAS), "Search for new phenomena in final states with an energetic jet and large missing transverse momentum in $p\bar{p}$ collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector," (2013), arXiV:1604.07773.

[28] M. Boudaud *et al.*, "A new look at the cosmic ray positron fraction," Astron. Astrophys. **575**, A67 (2015).

[29] M. Boudaud *et al.*, "Wino Dark Matter in light of the AMS-02 2015 Data," Phys. Rev. D. **91**, 111701 (2015).

[30] L. Accardo *et al.* (AMS), "High Statistics Measurement of the Positron Fraction in Primary Cosmic Rays of 0.5-500 GeV with the Alpha Magnetic Spectrometer on the International Space Station," Phys. Rev. Lett. **113**, 121101 (2014).

[31] L. Accardo *et al.* (AMS), "Precision Measurement of the Proton Flux in Primary Cosmic Rays from Rigidity 1 GV to 1.8 TV with the Alpha Magnetic Spectrometer on the International Space Station," Phys. Rev. Lett. **114**, 171103 (2015).

[32] M. Ackermann *et al.* (Fermi LAT), "Limits on Dark Matter Annihilation Signals from the Fermi LAT 4-year Measurement of the Isotropic Gamma-Ray Background," JCAP **1509**, 008 (2015).

[33] M. G. Aartsen *et al.* (IceCube), "Improved limits on dark matter annihilation in the Sun with the 79-string IceCube detector and implications for supersymmetry," JCAP **04** (2016).

[34] S. Adrián-Martíneza *et al.* (ANTARES), "Limits on Dark Matter Annihilation in the Sun using the ANTARES Neutrino Telescope," Phys. Lett. B759 (2016) **759**, 69 (2016).

[35] G. J. Alner *et al.* (ZEPLIN), "First limits on nuclear recoil events from the ZEPLIN I galactic dark matter detector," Astropart. Phys. **23**, 444 (2005).

[36] D. S. Akerib *et al.* (LUX), "Results from a search for dark matter in the complete LUX exposure," (2016), arXiv:1608.07648.

[37] P. Benetti *et al.* (WArP), "First results from a dark matter search with liquid Argon at 87K in the Gran Sasso underground laboratory," Astropart. Phys. **28**, 495 (2008).

[38] P. Agnes *et al.* (DarkSide), "Results from the first use of low radioactivity argon in a dark matter search," Phys. Rev. D **93**, 081101 (2016).

[39] G. J. Alner *et al.* (ZEPLIN), "First limits on WIMP nuclear recoil signals in ZEPLIN-II: a two phase xenon detector for dark matter detection," Astropart. Phys. **28**, 287 (2007).

[40] K. Abe *et al.* (XMASS), "Direct dark matter search by annual modulation in XMASS-I," (2015), arXiv:1511.04807.

[41] D. Y. Akimov *et al.* (ZEPLIN), "WIMP-nucleon cross-section results from the second science run of ZEPLIN-III," Phys. Lett. B **709**, 14 (2012).

[42] J. Angle *et al.* (XENON10), "First Results from the XENON10 Dark Matter Experiment at the Gran Sasso National Laboratory," Phys. Rev. Lett **100**, 021303 (2007).

[43] E. Aprile *et al.* (XENON100), "Dark matter results from 225 live days of XENON100 data," Phys. Rev. Lett **109**, 181301 (2012).

[44] X. Xiao *et al.* (PandaX), "Low-mass dark matter search results from full exposure of the PandaX-I experiment," Phys. Rev. D **92**, 052004 (2015).

[45] A. Tan *et al.* (PandaX), "Dark Matter Search Results from the Commissioning Run of PandaX-II," Phys. Rev. D **93**, 122009 (2016).

[46] D. S. Akerib *et al.* (LUX), "Improved WIMP scattering limits from the LUX experiment," Phys. Rev. Lett. **116**, 161301 (2016).

[47] F. Aubin *et al.* (PICASSO), "Discrimination of nuclear recoils from alpha particles with superheated liquids," New J. Phys **10**, 103017 (2008).

[48] C. Amole *et al.* (PICO), "Improved Dark Matter Search Results from PICO-2L Run-2," Phys. Rev. D **93**, 061101 (2016).

[49] C. Amole *et al.* (PICO), "Dark matter search results from the PICO-60 CF3I bubble chamber," Phys. Rev. D **93**, 052014 (2016).

[50] S. Archambault *et al.* (PICASSO), "Constraints on low-mass WIMP interactions on ^{19}F from PICASSO," Phys. Lett. B **711**, 153 (2012).

[51] M. Felizardo *et al.* (SIMPLE), "The SIMPLE Phase II dark matter search," Phys. Rev. D , 072013 (2014).

[52] E. Aprile *et al.* (XENON100), "Limits on spin-dependent WIMP-nucleon cross sections from 225 live days of XENON100 data," Phys. Rev. Lett **111**, 021301 (2013).

[53] K. Choi *et al.*, "Search for Neutrinos from Annihilation of Captured Low-Mass Dark Matter Particles in the Sun by Super-Kamiokande," Phys. Rev. Lett **144**, 141301 (2015).

[54] L. Roszkowski, R. R. de Austri, and R. Trotta, "Implications for the Constrained MSSM from a new prediction for $b \rightarrow s\gamma$," JHEP **07**, 75 (2007).

[55] F. Mayet *et al.*, "A review of the discovery reach of directional Dark Matter detection," J. Phys. Rep. **627**, 1 (2016).

[56] J. B. R. Battat *et al.*, "First background-free limit from a directional dark matter experiment: Results from a fully fiducialised DRIFT detector," Phys. Dark Universe **9**, 1 (2015).

[57] D. Santos *et al.*, "MIMAC: A micro-tpc matrix for dark matter directional detection," JPCS **460**, 012007 (2013).

[58] K. Miuchi *et al.*, "First underground results with NEWAGE-0.3a direction-sensitive dark matter detector," Phys. Lett. B **686**, 11 (2010).

[59] J. P. Lopez *et al.*, "Background Rejection in the DMTPC Dark Matter Search Using Charge Signals," Nucl. Instr. Meth. A **696**, 121 (2012).

[60] R. Bernabei *et al.*, "Dark Matter search," Riv. Nuovo. Cim. **26N1**, 1 (2003).

[61] R. Bernabei *et al.*, "Final model independent result of DAMA/LIBRA-phase1," Eur. Phys. J. C. **73**, 2648 (2013).

[62] B. Roberts, V. Dzuba, V. Flambaum, M. Pospelov, and Y. Stadnik, "Dark matter scattering on electrons: Accurate calculations of atomic excitations and implications for the DAMA signal," Phys. Rev. D. **93**, 115037 (2015).

[63] J. Davis, "Fitting the annual modulation in DAMA with neutrons from muons and neutrinos," Phys. Rev. Lett. **113**, 081302 (2014).

[64] J. Klinger and V. Kudryavtsev, "Muon-induced neutrons do not explain the DAMA data," Phys. Rev. Lett. **114**, 151301 (2015).

[65] C. E. Aalseth *et al.*, "Search for An Annual Modulation in Three Years of CoGeNT Dark Matter Detector Data," (2014), arXiv:1401.3295.

[66] C. E. Aalseth *et al.*, "Maximum Likelihood Signal Extraction Method Applied to 3.4 years of CoGeNT Data," (2015), arXiv:1401.6234.

[67] E. B. de Souza *et al.* (DM-ICE), "First Search for a Dark Matter Annual Modulation Signal with NaI(Tl) in the Southern Hemisphere by DM-Ice17," (2016), arXiv:1602.05939.

[68] C. Tomei (SABRE), "SABRE: Dark matter annual modulation detection in the northern and southern hemispheres," Nucl. Instr. Meth. A (2016), in press, available online.

[69] R. Gaitskell, "Devil in the Detail—Calibration of Dark Matter Detectors," (2016), presented at the 11th International Identification of Dark Matter Conference, Sheffield, U.K.

[70] A. J. Anderson, J. M. Conrad, E. Figueroa-Feliciano, K. Scholberg, and J. Spitz, "Coherent Neutrino Scattering in Dark Matter Detectors," Phys. Rev. D **84**, 013008 (2011).

227

[71] J. Billard, L. Strigari, and E. Figuerola-Feliciano, "Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments," Phys. Rev. D **89**, 023524 (2014).

[72] S. J. Asztalos *et al.* (ADMX), "SQUID-Based Microwave Cavity Search for Dark-Matter Axions," Phys. Rev. Lett **104**, 041301 (2010).

[73] J. R. Brownstein and J. W. Moffat, "Galaxy Rotation Curves without Nonbaryonic Dark Matter," Astrophys. J. **636**, 721 (2006).

[74] J. R. Brownstein and J. W. Moffat, "Galaxy Cluster Masses Without Non-Baryonic Dark Matter," Mon. Not. Roy. Astron. Soc **367**, 527 (2006).

[75] J. R. Brownstein and J. W. Moffat, "The Bullet Cluster 1E0657-558 evidence shows Modified Gravity in the absence of Dark Matter," Mon. Not. Roy. Astron. Soc **382**, 29 (2007).

[76] J. W. Moffat and V. T. Toth, "Modified Gravity: Cosmology without dark matter or Einstein's cosmological constant," (2007), arXiv:0710.0364.

[77] E. Dahl, *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 (2009).

[78] T. Doke, A. Hitachi, J. Kikuchi, K. Masuda, H. Okada, and E. Shibamura, "Absolute scintillation yields in liquid argon and xenon for various particles," Jap. J. Appl. Phys. **41**, 1538 (2002).

[79] A. Baldini *et al.*, "Absorption of scintillation light in a 100 l liquid xenon γ-ray detector and expected detector performance," Nucl. Instrum. Meth. A **545**, 753 (2005).

[80] C. Silva, J. P. da Cunha, A. Pereira, M. Lopes, V. Chepel, V. Solovov, and F. Neves, "A model of the reflection distribution in the vacuum ultra violet region," Nucl. Instrum. Meth. A **619**, 59 (2010).

[81] B. Lenardo, K. Kazkaz, A. Manalaysay, J. Mock, M. Szydagis, and M. Tripathi, "A Global Analysis of Light and Charge Yields in Liquid ," IEEE Trans. Nucl. Sci **62**, 3386 (2015).

[82] J. Lindhard, M. Scharff, and H. Schiott, "Range concepts and heavy ion ranges," Matt. Fys. Medd. Dan. Vid. Selsk **33**, 1 (1963).

[83] F. Bezrukov, F. Kahlhoefer, and M. Lindner, "Interplay Between Scintillation and Ionization in Liquid Xenon Dark Matter Searches," Astropart. Phys. **35**, 119 (2011).

[84] D. S. Akerib *et al.* (LUX), "An ultra-low background PMT for liquid xenon detectors," Nucl. Instrum. Meth. A **703**, 1 (2013).

[85] M. L. Cherry, M. Deakyne, K. Lande, C. K. Lee, R. I. Steinberg, B. Cleveland, and E. J. Fenyves, "Multiple muons in the homestake underground detector," Phys. Rev. D **27**, 1444 (1983).

[86] F. Gray, C. Ruybal, J. Totushek, D.-M. Mei, K. Thomas, and C. Zhang, "Cosmic ray muon flux at the Sanford Underground Laboratory at Homestake," Nucl. Instrum. Meth. A **638**, 63 (2011).

[87] D. S. Akerib *et al.* (LUX), "Data acquisition and readout system for the LUX dark matter experiment," Nucl. Instrum. Meth. A **688**, 1 (2011).

[88] D. S. Akerib *et al.* (LUX), "FPGA-based Trigger System for the LUX Dark Matter Experiment," Nucl. Instrum. Meth. A **818**, 57 (2015).

[89] S. Agostinelli *et al.*, "GEANT4 – a simulation toolkit," Nucl. Instrum. Meth. A **560**, 250 (2003).

[90] J. Allison *et al.*, "GEANT4 developments and applications," IEEE T. Nucl. Sci. **53**, 270 (2006).

[91] D. S. Akerib *et al.* (LUX), "LUXSim: A Component-Centric Approach to Low-Background Simulations," Nucl. Instrum. Meth. A **675**, 63 (2012).

[92] A. C. Sinnock and B. L. Smith, "Refractive Indices of the Condensed Inert Gases," Phys. Rev. **181**, 1297 (1969).

[93] V. Solovov, V. Chepel, M. I. Lopes, A. Hitachi, R. F. Marques, and A. J. P. L. Policarpo, "Measurement of the Refractive Index and Attenuation Length of Liquid xenon for its Scintillation Light," Nucl. Instrum. Meth. A **516**, 462 (2004).

[94] G. M. Seidel, R. E. Lanou, and W. Yao, "Rayleigh Scattering in Rare Gas Liquids," Nucl. Instrum. Meth. A **489**, 189 (2002).

[95] M. Szydagis *et al.*, "NEST: A Comprehensive Model for Scintillation Yield in Liquid Xenon," JINST **6**, P10002 (2011).

[96] I. M. Obodovskii and K. T. Ospanov, "Scintillation output of liquid xenon for low-energyγ-quanta," Instruments and Experimental Techniques **37**, 42 (1994).

[97] L. Baudis *et al.*, "Response of liquid xenon to Compton electrons down to 1.5 keV," Phys. Rev. D **87**, 115015 (2013).

[98] E. Aprile *et al.*, "Measurement of the Scintillation Yield of Low-Energy Electrons in Liquid Xenon," Phys. Rev. D. **86**, 112004 (2012).

[99] M. Szydagis *et al.*, "Enhancement of NEST Capabilities for Simulating Low–Energy Recoils in Liquid Xenon," JINST **8**, C10003 (2013).

[100] A. Hitachi, "Properties of liquid xenon scintillation for dark matter searches," Astropart. Phys **24**, 247 (2005).

[101] P. Sorensen and C. E. Dahl, "Nuclear recoil energy scale in liquid xenon with application to the direct detection of dark matter," Phys. Rev. D **83**, 063501 (2011).

[102] G. Plante *et al.*, "New measurement of the scintillation efficiency of low-energy nuclear recoils in liquid xenon," Phys. Rev. C **84**, 045805 (2011).

[103] J. Angle *et al.* (XENON100), "First Dark Matter Results from the XENON100 Experiment," Phys. Rev. Lett **105**, 131102 (2010).

[104] M. Horn *et al.* (ZEPLIN), "Nuclear recoil scintillation and ionisation yields in liquid xenon from ZEPLIN-III data," Phys. Lett. B **705**, 471 (2011).

[105] J. Mock *et al.*, "Modeling Pulse Characteristics in Xenon with NEST," JINST **9**, T04002 (2014).

[106] D. S. Akerib *et al.* (LUX), "First results from the LUX dark matter experiment at the Sanford Underground Research Facility," Phys. Rev. Lett. **112**, 091303 (2014).

[107] C. H. Faham, V. M. Gehman, A. Currie, A. Dobi, P. Sorensen, and R. J. Gaitskell, "Measurements of wavelength-dependent double photoelectron emission from single photons in VUV-sensitive photomultiplier tubes," J. Instrum. **10**, P09010 (2015).

[108] "COMSOL Multiphysics®," http://www.comsol.com.

[109] D. S. Akerib *et al.* (LUX), "Radiogenic and Muon-Induced Backgrounds in the LUX Dark Matter Detector," Astropart. Phys **62**, 33 (2014).

[110] D.-M. Mei, C. Zhang, and A. Hime, "Evaluation of (α,n) Induced Neutrons as a Background for Dark Matter," Nucl. Instrum. Meth. A **606**, 651 (2009).

[111] J. Back and Y. A. Ramachers, "ACTIVIA: Calculation of Isotope Production Cross-sections and Yields," NIM **A586**, 286 (2008).

[112] D. Mei and A. Hime, "Muon-induced background study for underground laboratories," Phys. Rev. D **73**, 053004 (2006).

[113] D. E. Groom, N. V. Mokhov, and S. Striganov, "Muon Stopping Power and Range Tables 10 MeV-100 TeV," Atomic Data and Nuclear Data Tables **78**, 183 (2001).

[114] D. S. Akerib *et al.* (LUX), "The Large Underground Xenon (LUX) Experiment," Nucl. Instrum. Meth. A **704**, 111 (2013).

[115] A. Derbin *et al.* (Borexino), "The Main Results of the Borexino Experiment," (2016), arXiv:1605.06795.

[116] K. Asakura *et al.* (KamLAND-Zen), "Results from KamLAND-Zen," AIP Conf. Proc. **1666**, 170003 (2016).

[117] A. W. Bradley, *LUX Thermosyphon Cryogenics and Radon-Related Backgrounds For The First WIMP Result*, Ph.D. thesis, Case Western Reserve University (2014).

[118] D. S. Akerib *et al.* (LUX), "Tritium calibration of the LUX dark matter experiment," Phys. Rev. D **93**, 072009 (2015).

[119] A. Dobi, *Measurement of the Electron Recoil Band of the LUX Dark Matter Detector with a Tritium Calibration Source*, Ph.D. thesis, University of Maryland (2014).

[120] H. Park, K.-O. Choi, J.-M. Lee, M. S. Hahn, and M. Krank, "Absolute Measurement of the Neutron Emission Rate with a Manganese Sulphate Bath System," J. Korean Phys. Soc. **47**, 603 (2005).

[121] Scott Dewey, M. and Thompson, A. K., *Radioactive Neutron Sources Emission Rates*, Tech. Rep. 44010C-44020C (Ionizing Radiation Division, NIST, 2010).

[122] J. Verbus *et al.*, "Proposed low-energy absolute calibration of nuclear recoils in a dual-phase noble element TPC using D-D neutron scattering kinematics," (2016), arXiv:1608.05309.

[123] D. S. Akerib *et al.* (LUX), "Low-energy (0.7-74 keV) nuclear recoil calibration of the LUX dark matter experiment using D-D neutron scattering kinematics," (2016), arXiv:1608.05381.

[124] T. Joshi *et al.*, "First Measurement of the Ionization Yield of Nuclear Recoils in Liquid Argon," Phys. Rev. Lett. **112**, 171303 (2014).

[125] E. Aprile *et al.*, "Simultaneous Measurement of Ionization and Scintillation from Nuclear Recoils in Liquid Xenon for a Dark Matter Experiment," Phys. Rev. Lett **97**, 081302 (2006).

[126] A. Manzur, A. Curioni, L. Kastens, D. N. McKinsey, K. Ni, and T. Wongjirad, "Scintillation efficiency and ionization yield of liquid xenon for monoenergetic nuclear recoils down to 4 keV," Phys. Rev. C. **81**, 025808 (2009).

[127] E. Aprile *et al.*, "New measurement of the relative scintillation efficiency of xenon nuclear recoils below 10 keV," Phys. Rev. C. **79**, 045807 (2009).

[128] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, "Asymptotic formulae for likelihood-based tests of new physics," Eur. Phys. J. C **71**, 1554 (2011).

[129] R. H. Helm, "Inelastic and Elastic Scattering of 187 MeV Electrons from Selected Even-Even Nuclei," Phys. Rev. **104**, 1466 (1956).

[130] R. Agnese *et al.*, "Search for low-mass weakly interacting massive particles with Super-CDMS," Phys. Rev. Lett. **112**, 241302 (2014).

[131] R. Agnese *et al.*, "New results from the search for low-mass weakly interacting massive particles with the CDMS low ionization threshold experiment," Phys. Rev. Lett. **116**, 071301 (2016).

[132] D. S. Akerib *et al.* (LUX), "First spin-dependent WIMP-nucleon cross section limits from the LUX experiment," Phys. Rev. Lett. **116**, 161302 (2016).

[133] Z. Ahmed *et al.*, "Results from a Low-Energy Analysis of the CDMS II Germanium Data," Phys. Rev. Lett. **106**, 131302 (2011).

[134] S. C. Kim *et al.*, "New Limits on Interactions between Weakly Interacting Massive Particles and Nucleons Obtained with CsI(Tl) Crystal Detectors," Phys. Rev. Lett **108**, 181301 (2015).

[135] J. AngleL *et al.* (XENON10), "Limits on Spin-Dependent WIMP-Nucleon Cross Sections from the XENON10 Experiment," Phys. Rev. Lett **101**, 091301 (2008).

[136] V. N. Lebedenko *et al.* (ZEPLIN), "Limits on the Spin-Dependent WIMP-Nucleon Cross Sections from the First Science Run of the ZEPLIN-III Experiment," Phys. Rev. Lett **103**, 151302 (2009).

[137] C. Savage, G. Gelmini, P. Gondolo, and K. Freese, "Compatibility of DAMA/LIBRA dark matter detection with other searches," JCAP **04** (2009).

[138] S. Malik *et al.*, "Interplay and characterization of Dark Matter searches at colliders and in direct detection experiments ," Phys. Dark Universe **04**, 51 (2015).

[139] D. S. Akerib *et al.* (LZ), "LUX-ZEPLIN (LZ) Conceptual Design Report," (2015), arXiv:1509.02910.

[140] D. R. Tovey, R. J. Gaitskell, P. Gondolo, Y. Ramachers, and L. Roszkowski, "A new model-independent method for extracting spin-dependent cross section limits from dark matter searches," Phys. Lett. B **488**, 17 (2000).

[141] Z. Ahmed *et al.*, "Search for Weakly Interacting Massive Particles with the First Five-Tower Data from the Cryogenic Dark Matter Search at the Soudan Underground Laboratory," Phys. Rev. Lett. **102**, 011301 (2009).

[142] A. Manalaysay, "Dark-matter results from 332 new live days of LUX data," (2016), presented at the 11th International Identification of Dark Matter Conference, Sheffield, U.K.

[143] E. A. Bagnaschi *et al.*, "Supersymmetric Dark Matter after LHC Run 1," Eur. Phys. J. C **75**, 1 (2015).

[144] S. Shaw *et al.* (LZ), "Identification of Radiopure Titanium for the LZ Experiment," (2016), *in preparation*.

[145] C. Ghag *et al.* (ZEPLIN), "Performance of the veto detector incorporated into the ZEPLIN-III experiment," Astropart. Phys **35**, 76 (2011).

[146] M. Chadwick *et al.*, "ENDF/B-VII.1: Nuclear Data for Science and Technology: Cross Sections, Covariances, Fission Product Yields and Decay Data," Nucl. Data Sheets **112**, 2887 (2011).

[147] X. R. Liu and J. Mott, "Low background techniques for SuperNEMO," JPCS **598**, 012022 (2015).

[148] *ASM Handbook Volume 15: Casting* (ASM International, 2008).

[149] A. V. Voronkov, V. I. Zhuravlev, and E. G. Natrusova, "The Group Neutron Data Library (GNDL)," International Atomic Energy Agency (IAEA) (1987).

[150] W. B. Wilson, R. T. Perry, W. S. Charlton, and T. A. Parish, *SOURCES4A: A code for calculating* (α,n)*, spontaneous fission, and delayed neutron sources and spectra*, Tech. Rep. (1999) lA-13639-MS, Los Alamos.

[151] F. Gönnenwein, "Neutron and Gamma Emission in Fission," LANL FIESTA Fission School and Workshop (2014).

[152] J.M. Verbeke, C. Hagmann and D. Wright, *Simulation of Neutron and Gamma Ray Emission from Fission and Photofission* (2014).

[153] J. Terrell, "Distribution of Fission Neutron Numbers," Physical Review **108**, 783 (1957).

[154] C. McCabe, "The Astrophysical Uncertainties Of Dark Matter Direct Detection Experiments," Phys. Rev. D **82**, 023530 (2010).

[155] D. S. Akerib *et al.* (LZ), "LUX-ZEPLIN (LZ) Technical Design Report," (2016), *in preparation*.