

Dark Matter Searches and Study of Electrode Design in LUX and LZ

Adam Bailey

Imperial College London

Department of Physics

A thesis submitted for the degree of Doctor of Philosophy
August 2016

Abstract

There is substantial evidence that over 80% of matter in the universe is dark matter – which is non-baryonic in nature and is thought to be composed of a new, slow-moving, stable particle not found in the Standard Model of Particle Physics. Its presence is inferred from gravitational effects on luminous matter from several independent observations, from the galactic to the cosmological scale. Weakly Interacting Massive Particles (WIMPs) are the leading candidate, which can explain all of the observed effects.

LUX and LZ are dual-phase xenon time projection chambers (TPC), aiming to observe scattering of WIMPs from xenon nuclei. LUX has an active mass of 250 kg of liquid xenon, and took data at the Sanford Underground Research Facility in Lead, South Dakota, between 2013 and 2016. The first WIMP search run of 85 live days in 2013 set world-leading exclusion limits on the spin-independent WIMP-nucleon cross section. This was improved by a reanalysis of those data, and subsequently by a new run yielding 332 live days, which set a minimum exclusion limit of $2.2 \times 10^{-46} \text{ cm}^2$ for a 50 GeV WIMP (90% CL). In addition, the most stringent limit to date on the spin-dependent WIMP-neutron scattering cross section comes from the reanalysis of the 2013 dataset, with a minimum exclusion of $9.4 \times 10^{-41} \text{ cm}^2$ for a 33 GeV WIMP. LZ is a next generation experiment with a 7 tonne active mass to be deployed in the same location as LUX, expected to be 100 times more sensitive.

Work presented in this thesis includes analysis of the 2013 LUX search data to produce the spin-dependent results, evaluating the detector response using a tritium β^- source, and determining the ^{85}Kr background from data. A study was carried out on spurious electron emission phenomena from thin cathodic wires under high electric fields, using LUX engineering data where the grid voltages were increased above nominal operating values; this led to new insights into the microscopic breakdown mechanisms which have affected these (and other) TPC detectors for decades. The detailed understanding of the electroluminescence response gained in LUX was applied to the design of the LZ electroluminescence region; detailed simulation work of electrode geometry was performed to assess the performance of several candidate designs.

Declaration

I declare that material in this thesis is my own work. Any work that is not my own is appropriately stated and referenced. The specific contributions in each chapter are listed below.

Chapters 2 and 3 are describing background information on dark matter and the LUX detector. The latter describes work by the LUX and LZ collaborations.

Run 3 of the LUX experiment is described in Chapter 4. My contributions were the development of a real time data monitor, an analysis to determine the detection efficiency from tritium data, and a ^{85}Kr background analysis. The section on the WIMP search analysis in Chapter 4 is describing collaboration work.

For Chapter 5 I implemented the spin-dependent calculations into the existing LUX analysis code, and assisted in developing the code further. I also assessed the impact of nuisance parameters on the result. The section on the signal estimation describes the statistical test used in the analysis code. The sections on the background models, detector response simulations and the LUX data analysis are describing work done by the collaboration. I was responsible for producing the final results on the spin-dependent scattering cross sections. The exclusion regions on the WIMP-proton and WIMP-neutron coupling constants were done with assistance from Nicole Larsen (Yale).

The study of electron emission phenomena presented in Chapter 6 was an independent analysis. The data used were taken by collaborators during the LUX conditioning campaign.

The simulations in Chapter 7 were done in collaboration with Pawel Majewski (RAL), who provided assistance with the simulation software and creation of the electric field maps. This included creating the original field map for the woven wire mesh and the XENON100 grids. The rest of the simulation work in this chapter was my own.

– Adam Bailey

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Acknowledgements

Firstly, I would like to thank my supervisor Henrique Araújo for invaluable guidance and assistance throughout my PhD and thesis writing. He was always there with useful advice and discussion for any aspect of my studies, and I feel that I have benefited greatly from his experience.

I extend my thanks to coworkers who have assisted me in various ways. To Pawel Majewski for the help with the simulation software for the LZ design work, along with useful feedback on that chapter. Also to Thomas Shutt and Jerry Va'vra for feedback and discussion on that same work. To Daniel McKinsey, Wing To, Nicole Larsen, Alastair Currie, and Aaron Manalaysay for their assistance and input on the LUX Run 3 reanalysis. To Attila Dobi and Matthew Szydagis for useful discussions relating to NEST and tritium, and to Alex Lindote for making the LUX data processing more transparent. To Alfredo Tomas, I learnt many things during the many hours in the lab, although I never want to deal with 50 μm wires again.

My gratitude extends to all members of LUX and LZ, both past and present. It is combined effort of many people that has made the success of the experiments possible. It has been a pleasure to work as part of both collaborations.

I also wish to acknowledge STFC for funding my studies at Imperial College.

There are many people I would like to thank for making the time on-site in South Dakota enjoyable, whether it be eating, drinking, hiking, board games, being stuck in the snow, or even underground in the lab: Doug Tiedt, Mark Hanhardt, Matthew Szydagis (additionally for providing N64), Dev Ashish Khaitan, Claudio Silva, Alex Lindote, Luiz de Viveiros, Carmen Carmona, Jim Dobson, Nicole Larsen, Attila Dobi, Richard Knoche, Mongkol Moongweluwan, Chris Chiller, Angela Chiller, Curt Nehr Korn (where does the krypton go?), Sally Shaw, James Verbus, Kevin O'Sullivan, Elizabeth Boulton, Markus Horn (Forgive me if I have left anyone out). I actually miss it there sometimes!

Lastly, I would like to thank my parents, who have always supported and encouraged me in whatever I wanted to do. And to María, who has been there for support throughout my PhD, and has probably learnt a fair amount about dark matter herself.

Contents

1	Introduction	21
2	Dark Matter – Evidence, Candidates, and Detection	24
2.1	Evidence	24
2.1.1	Galactic Scale	24
2.1.2	Cosmological Scales	27
2.2	Candidates	30
2.2.1	Weakly Interacting Massive Particles (WIMPs)	31
2.2.2	Cosmological properties of WIMPs	32
2.2.3	Non-WIMP particle dark matter	33
2.3	Detection	33
2.3.1	Collider searches	34
2.3.2	Indirect searches	36
2.3.3	Direct detection	38
2.3.4	Complementarity of detection methods	42
3	Direct Detection with Dual-Phase Xenon	44
3.1	Particle Detection in Liquid Xenon TPCs	44
3.1.1	Prompt scintillation signal	46
3.1.2	Electron transport and cross-phase extraction	47
3.1.3	Secondary scintillation light	48
3.1.4	Light yield, charge yield, discrimination and energy	49
3.2	The Large Underground Xenon (LUX) Experiment	51
3.2.1	TPC and cryostat	52
3.2.2	Grids and high voltage delivery	53
3.2.3	Photomultipliers	53

3.2.4	Gas system, circulation and purification	54
3.2.5	Cryogenics and thermosyphon cooling	55
3.2.6	Instrumentation and slow control	56
3.2.7	DAQ and electronics	57
3.2.8	Calibration systems	58
3.2.9	Water tank and veto	60
3.3	The LUX-ZEPLIN (LZ) Experiment	60
3.3.1	The LZ anti-coincidence systems	62
4	LUX Run 3	64
4.1	Operations and Data Processing	64
4.1.1	Detector operations	64
4.1.2	Data processing	65
4.2	Real Time Data Monitoring	67
4.2.1	Event and pulse selection	69
4.2.2	Output and uses	70
4.3	Calibrations	74
4.3.1	Krypton-83m	74
4.3.2	Tritium	75
4.4	^{85}Kr Background Analysis	83
4.4.1	Introduction	83
4.4.2	S1 energy calibration	85
4.4.3	S1 and S2 energy calibration	87
4.4.4	^{85}Kr search analysis	91
4.4.5	S1-only analysis	92
4.4.6	Combined energy analysis	97
4.4.7	Conclusions from ^{85}Kr background analysis	102
4.5	Run 3 WIMP Search Analysis	102
4.5.1	Backgrounds	104

4.5.2	Results	106
4.6	Conclusions	107
5	Spin-Dependent Limits from the 2015 LUX Run 3 Reanalysis	108
5.1	Spin-Dependent Interaction	108
5.1.1	Nuclear spin structure	109
5.1.2	SD WIMP recoil spectrum	114
5.2	Signal Estimation	120
5.3	Models for the LUX PLR	123
5.3.1	Backgrounds and constraints	124
5.3.2	Detector response simulations	125
5.4	Implementation of Nuisance Parameters in the PLR	129
5.5	LUX Data Analysis	136
5.6	Spin-Dependent Results	138
5.7	Conclusions	144
6	Experimental Study of Electron Emission from Thin Wires in Xenon	145
6.1	Electron Emission Mechanisms	145
6.2	Analysis of LUX Electron Emission Data	147
6.2.1	Gate in gas	149
6.2.2	Gate in liquid	159
6.3	Summary of Electron Emission Data	165
6.4	Search for Light Emission from Wires in Liquid Data	168
6.5	Discussion: Which Emission Mechanism?	173
6.6	Conclusion	176
7	Design of the LZ Electroluminescence Region	177
7.1	Design Considerations	179
7.1.1	Photon yields	179
7.1.2	Extraction efficiency	180

7.1.3	Resolution of S2 signals	181
7.1.4	Maximum electric fields on wire surfaces	187
7.1.5	Optical transparency	188
7.1.6	Mechanical and electrostatic considerations	188
7.1.7	Other spurious effects	189
7.1.8	Summary of design considerations and limitations	189
7.2	High Energy Signals: Neutrinoless Double Beta Decay	190
7.3	Towards an LZ Design: Simulation Method	191
7.4	Validation	193
7.5	Simulations of Existing Experiments	195
7.5.1	Creating field maps for previous experiments	195
7.5.2	Mesh density effects on photon and ionisation yields	203
7.5.3	Results from simulations of experimental geometries	205
7.5.4	Comparison to LUX data	210
7.5.5	Comparison to XENON100 data	212
7.5.6	Conclusions from validation studies	213
7.6	Simulations of LZ Grids	214
7.6.1	Creating LZ field maps	215
7.6.2	Parallel and crossed grid results	221
7.6.3	Woven mesh results	225
7.6.4	Effects of grid deflection	233
7.6.5	Conclusions	234
7.7	Energy Resolution for Larger Signals	235
7.8	Summary	239
8	Conclusions and Outlook	240
	Bibliography	241
	Appendix A Figure Copyright	270

List of Figures

2.1	Rotation curve from spiral galaxy NGC 6503	25
2.2	X-ray and matter distribution of the Bullet Cluster merger	27
2.3	CMB power spectrum from Planck data	29
2.4	Feynman diagrams of WIMP interactions	34
2.5	Feynman diagram of WIMP production with ISR	35
2.6	Feynman diagram of Higgs production from vector boson fusion	36
3.1	Energy deposits from nuclear and electronic recoils in xenon	45
3.2	Schematic of an interaction in a liquid xenon TPC	46
3.3	Event display from LUX	46
3.4	The LUX detector in the Davis Cavern	51
3.5	Internal structure of LUX	52
3.6	Diagram of a thermosyphon	56
3.7	LUX DAQ system	57
3.8	LZ detector design	60
4.1	Example waveform containing single electrons and single phes	66
4.2	Flow chart of the DPM	68
4.3	DPM pulse classification	69
4.4	DPM pulse area against pulse length	71
4.5	DPM POD rate	71
4.6	DPM S2 against S1 pulse area	72
4.7	S2/S1 against drift time from DPM	72
4.8	Monitoring of electron lifetime with the DPM	73
4.9	DPM plot of prompt fraction against top bottom asymmetry	73
4.10	^{83m}Kr decay scheme	74

4.11	Example event display from a ^{83m}Kr decay	75
4.12	^{83m}Kr pulse area corrections	76
4.13	^{83m}Kr position corrections	76
4.14	Tritium β^- spectrum	77
4.15	Measured tritium spectrum in S1 phe	78
4.16	S1 detection efficiency from tritium data	79
4.17	Measured tritium spectrum and golden efficiency in S2 phe	80
4.18	Tritium spectrum in reconstructed energy	81
4.19	Final tritium β^- spectrum	82
4.20	Efficiency curve from the LUX Run 3 analysis	82
4.21	^{85}Kr decay scheme	84
4.22	Illustration of ^{85}Kr delayed coincidence event	85
4.23	γ -ray background S1 spectrum	86
4.24	γ selection cut for ^{85}Kr analysis	87
4.25	Background energy spectra for single and multiple scatter γ -ray events .	88
4.26	^{131m}Xe , ^{127}Xe and ^{129m}Xe background peaks in single scatter events . . .	89
4.27	^{214}Bi peak from single scatters	89
4.28	^{129m}Xe background peak in multiple scatter events	90
4.29	^{214}Bi peak from multiple scatter events	90
4.30	Energy resolution for ^{85}Kr analysis	91
4.31	^{85}Kr S1-only analysis, t_{S1} after fiducial and S1 area cuts	93
4.32	Events in peak of t_{S1} distribution	93
4.33	$S1_S$ cut for ^{85}Kr analysis	94
4.34	t_{S1} after applying the fiducial, S1 area, and $S1_S$ cuts	94
4.35	^{85}Kr S1-only. Unidentified area and RMS width after fiducial, S1 area, $S1_S$ cuts	95
4.36	$\log_{10}(\Sigma S2_b/\Sigma S1)$ in ^{85}Kr S1 only analysis	96
4.37	Final t_{S1} distribution for the S1 only ^{85}Kr analysis after all cuts	97
4.38	Distribution of total energy for ^{85}Kr combined energy analysis	99

4.39	Effect of S1 cuts in the combined energy ^{85}Kr analysis	100
4.40	Final cuts in combined energy analysis on S1 sizes	100
4.41	Final t_{S1} distribution for the ^{85}Kr combined energy analysis	101
4.42	Spatial distribution of events after all cuts for the LUX Run 3 analysis . .	103
4.43	ER and NR band calibration in LUX Run 3	104
4.44	$\log_{10}(S2/S1)$ distribution for events passing all cuts in the LUX Run 3 analysis	106
4.45	LUX Run 3 limits on SI WIMP scattering	107
5.1	Diagram of nuclear shell model truncations	110
5.2	Spin-dependent structure factors	113
5.3	SI and SD WIMP signal spectra	118
5.4	Normalised SI and SD WIMP signal spectra	119
5.5	Spin-dependent WIMP signal spectra	119
5.6	Example test statistic distributions from PLR	122
5.7	Example of p-value scan from PLR	123
5.8	Spin-dependent WIMP signal PDFs in S1 vs. S2	129
5.9	Effect of g_2 variation on $\log_{10}(S2)$ vs S1 for a 33 GeV WIMP	131
5.10	Effect of TIB variation on $\log_{10}(S2)$ vs S1 for a 33 GeV WIMP	131
5.11	$\log_{10}(S2/S1)$ vs. S1 after all cuts in the reanalysis of LUX Run 3 data . .	139
5.12	Limits on the spin-dependent WIMP-neutron and WIMP-proton cross sections from the LUX reanalysis	140
5.13	Exclusions on spin-dependent WIMP couplings to protons and neutrons	141
5.14	SD limits with and without nuisance parameters	143
5.15	Comparison of SD limits with different structure functions	143
6.1	Potential barrier in Fowler-Nordheim theory	146
6.2	Potential barrier for emission from insulating impurity	147
6.3	Example waveforms from electron emission data	149
6.4	Emitter locations on Jan. 10th 2014 gas data	150

6.5	Emission rates from Jan. 10th 2014 gate-anode conditioning	151
6.6	Detected phe for emitter B on Jan. 10th 2014 during first stable period . .	152
6.7	Detected phe from emitter A on Jan 10th 2014 during first stable period .	153
6.8	Emitted electrons per pulse for emitters B and E on Jan. 10th 2014 during the first stable period	154
6.9	Positions and pulse rate for emitter H in gas on Jan. 12th 2014	155
6.10	Number of electrons from emitter H in gas on Jan. 12th 2014	155
6.11	Emitter H location and pulse rate on the Jan. 13th 2014 morning test . . .	156
6.12	Emitter H location and pulse rate on the Jan. 13th 2014 evening test . . .	156
6.13	Pulse rate and positions of Jan. 14th 2014 emitters in the gas	157
6.14	Emitter H pulse rate in the gas on the Jan. 22th 2014	158
6.15	Number of emitted electrons for emitter H on Jan. 22nd 2014	158
6.16	Emitters in Jan. 7th liquid data	160
6.17	Position and rates of emitters on Jan. 9th 2014 liquid data	161
6.18	Pulse areas and number of electrons for emitter D in liquid on Jan. 9th 2014	161
6.19	Positions and rates of emitters on Jan. 15th test in the liquid.	162
6.20	Phe and number of electrons from emitter M on Jan. 15th 2014, at 600 s .	163
6.21	Phe and number of electrons from emitter M on Jan. 15th 2014, at 5500 s	163
6.22	Waveform of electron pulses with glow from Jan. 28th 2014 data	164
6.23	Waveform of glow from Jan. 28th 2014 data	165
6.24	Positions and pulse rates of emitters in the liquid on Jan. 28th 2014 . . .	165
6.25	Averaged waveforms from liquid electron emission data	169
6.26	Detected phe from the gate for emitters in the liquid on Jan. 9th 2014 . .	170
6.27	Detected phe from the gate for emitter D in the liquid on Jan. 9th 2014 .	171
6.28	Detected phe from the gate for emitter M in the liquid on Jan. 15th 2014, at 3000 s	171
6.29	Detected phe from the gate for emitter M in the liquid on Jan. 15th 2014 at 5500 s	172
6.30	Estimated photons emitted from the gate per electron	172

6.31	Spheroid and cylindrical projections from wire surfaces	174
7.1	Types of wire grid geometry	178
7.2	Effects of EL variance on overall S2 resolution	185
7.3	Leakage fraction from FastNEST simulations with LZ parameters	186
7.4	Mean of relative leakage fraction for various EL widths	186
7.5	Diagram of NDBD position dependence	191
7.6	Example electron drift tracks	192
7.7	Steps required to create electric field maps	192
7.8	Garfield++ photon yield as a function of reduced field	194
7.9	Garfield++ drift velocity as a function of reduced field	194
7.10	Geometry and fields for LUX, ZEPLIN-III and XENON100	195
7.11	Making woven mesh geometry in GMSH	196
7.12	E field on wire surfaces for LUX simulation with different mesh densities	198
7.13	E field strength in LUX simulation	199
7.14	Layout of XENON100 grids and triangular unit cell	200
7.15	E field on wire surfaces for XENON100 simulation with different mesh densities	200
7.16	E field strength in LUX simulation	201
7.17	Wire roundings for the XENON100 simulation	202
7.18	Electron drift lines in XENON100 simulation	203
7.19	Photon yield for LUX simulation with different mesh densities	203
7.20	Excitations and ionisations for XENON100 simulation with different mesh densities on anode	204
7.21	Excitations per electron in ZEPLIN-III simulation	205
7.22	Excitations per electron in the LUX simulation	206
7.23	Excitations per electron in the XENON100 simulation	207
7.24	Ionisations per electron in the XENON100 simulation	207
7.25	Detected phe of emitter E, from the LUX gate in the gas, on Jan. 10th 2014	211
7.26	Average waveforms for single electron emitter E in the gas on Jan. 10th 2014	212

7.27	Comparison of single electron simulations to XENON100 data	213
7.28	Parameters for different possible LZ scenarios	215
7.29	Crossed mesh unit cell	216
7.30	E field in LZ pre-TDR scenario with 5 mm gate-anode pitch	217
7.31	LZ gate-anode unit cell with equal pitches	218
7.32	Drift lines for LZ pre-TDR scenario with 5 mm gate-anode pitch	218
7.33	Unit cell for woven mesh with anode pitch half of gate pitch	219
7.34	Drift lines for woven mesh with anode pitch half of gate pitch, and aligned with each other	220
7.35	Alignment and focussing for woven mesh with anode pitch half of the gate pitch	220
7.36	Photons per electron for initial simulations with parallel wire grid anode	223
7.37	Photons per electron for initial simulations with crossed wire grid anode	224
7.38	Ionisations per electron for initial simulations with crossed wire grid anode	225
7.39	Pulses from 5 mm pre-TDR geometry	226
7.40	Photons per electron for the geometries with 5 mm anode pitch	227
7.41	Time to secondary pulses for the geometries with 5 mm anode pitch . . .	228
7.42	Excitations and ionisations for the pre-TDR scenario simulation with 7.5 mm anode and gate pitch	228
7.43	Photons per electron for LZ CDR and pre-TDR scenario simulations with 2.5 mm anode pitch	229
7.44	Time difference to secondary pulses for LZ CDR and pre-TDR scenario simulations with 2.5 mm anode pitch	230
7.45	Pulses from 2.5 mm pre-TDR scenario	231
7.46	Effects of grid deflection for LZ pre-TDR scenario	234
7.47	S2 for NDBD events in LZ for different x - y positions	236
7.48	S1 and S2 yields from FastNEST NDBD events	237
7.49	Reconstructed energy of NDBD events simulated in FastNEST with LZ detector parameters	238
7.50	FastNEST reconstructed energy from 5 and 10 keV ER events with LZ parameters	238

List of Tables

3.1	Parameters for LUX Run 3	51
3.2	Parameters of LUX grids in Run 3	54
3.3	Parameters of LZ	61
4.1	^{83m}Kr resolution in S1	86
4.2	S1 resolution for the ^{85}Kr signal	86
4.3	Fiducial and bad area cut for ^{85}Kr analysis	88
4.4	Reconstructed energy for single scatter events from peaks in background spectrum	89
4.5	Reconstructed energy for multiple scatter events from peaks in background spectrum	90
4.6	Energy resolution used for the signal in the ^{85}Kr analysis	91
4.7	Livetime in the ^{85}Kr analysis	92
4.8	Cuts used in the ^{85}Kr S1-only analysis	96
4.9	Events remaining after each cut in the ^{85}Kr S1-only analysis	96
4.10	Energy cuts for ^{85}Kr combined energy analysis	100
4.11	Cuts used in the ^{85}Kr combined energy analysis	101
4.12	Events remaining after each cut in the ^{85}Kr combined energy analysis	101
4.13	Run 3 analysis cuts	103
5.1	Nuisance parameters for the nuclear recoil signal model	128
5.2	Nuisance parameters effects on number of events for SD neutron-only	132
5.3	Nuisance parameters effects on number of events for SD proton-only	133
5.4	Nuisance parameters effects on PDF shape for SD neutron-only	134
5.5	Nuisance parameters effects on PDF shape for SD proton-only	135
6.1	Summary of electron emitters during stable grid conditions	167

7.1	Sources of variance in dual-phase TPCs	181
7.2	LUX EL simulation results with different mesh element sizes	204
7.3	XENON100 EL simulation results with different mesh element sizes . . .	204
7.4	EL simulation results from ZEPLIN-III, LUX, and XENON100	209
7.5	Woven wire mesh parameters for possible LZ anodes	215
7.6	EL simulation results from single anode grids for LZ	222
7.7	EL simulation results from crossed wire anode grids for LZ	222
7.8	EL simulation results from woven wire mesh anodes for LZ	232

Chapter 1

Introduction

Dark matter is a fundamental component of the universe, that is essential in the standard cosmological model. It is very likely to consist of particles, yet is missing from the Standard Model of Particle Physics. Therefore, the exact composition of dark matter currently remains unknown. Several candidate particles are amenable to direct detection, such as the WIMP, which scatters weakly from nuclei. Although WIMPs are naturally difficult to detect, it is possible, and would be a major discovery that confirms the nature of dark matter in the universe. The LUX and LUX-ZEPLIN (LZ) experiments are two major efforts in the field, the leading current experiment and a next-generation system now under construction respectively. In this thesis I describe the search for dark matter interactions in LUX and some aspects of design optimisation for LZ.

I begin with a summary of the wide range of astronomical evidence for dark matter in Chapter 2. This is followed by a discussion of the possible candidates, in particular why WIMPs are so well-motivated and remain the leading candidate. This chapter concludes with the possible methods of detecting WIMPs: at colliders, indirect searches, and direct searches.

Chapter 3 describes how dual-phase xenon TPCs are used to search for WIMPs. It discusses particle interactions in liquid xenon, and the mechanisms for detecting the energy deposits from the nuclear or electronic recoils. This includes the primary scintillation signal, electron transport, cross-phase emission, and production of secondary scintillation light (electroluminescence). The yields and energy reconstruction of the recoils are also discussed, and how this results in good rejection of electronic recoil backgrounds. Following this I describe the LUX experiment, with details on the various subsystems required to operate the detector. The chapter concludes with an overview of the LZ experiment, where the main differences from LUX are discussed, in particular the scaling up challenges and the veto systems.

Chapter 4 presents Run 3 of the LUX experiment – the first WIMP search run in the Davis Lab of the Sanford Underground Research Facility (depth 4850'). The spin-independent WIMP search results from the collaboration are presented, with my own contributions highlighted. This includes the development of a real-time data monitor, used throughout data-taking to ensure data quality and check for potential problems with the detector. An analysis of the detection efficiency for electronic recoils was carried out using a tritium β^- calibration. This was one of several methods used for measuring the detector response in Run 3. An internal background dispersed throughout the xenon comes from ^{85}Kr β^- decays, which are typically estimated by sampling

xenon gas from the detector to determine the total krypton content. I employed a delayed coincidence method to search for a rare decay of ^{85}Kr that produces a distinct signal in the detector, and provided an independent confirmation of the result from the sampling.

Improvements to the background model, event selection and position reconstruction along with further calibration motivated a reanalysis of the first LUX data from 2013. I led the spin-dependent analysis, which is presented in Chapter 5 and was the first from the LUX collaboration. This chapter begins with discussion of the spin-dependent interaction and calculation of the differential event rate. This is followed by details of the analysis, such as the background models and the development of the profile likelihood ratio statistical analysis. I implemented the spin-dependent calculations into the existing LUX analysis framework, assisted in developing the code further and assessed the impact of nuisance parameters. I produced the final results on the spin-dependent scattering cross sections. Exclusion regions on the WIMP-proton and WIMP-neutron coupling constants were also presented. The spin-dependent WIMP-neutron scattering cross section exclusion is world-leading at time of writing, with a minimum of $9.4 \times 10^{-41} \text{ cm}^2$ for a 33 GeV WIMP (90% CL).

Dual-phase xenon TPCs require electric fields for electron transport and generation of an optical response from the ionisation released by the interaction, and these fields are created by various wire grids held at a high voltage. In these detectors the maximum voltage is limited by the onset of spurious localised electron emission from cathodic wires, which occurs at a much lower field than expected. In Chapter 6, I present a study of this electron emission phenomenon from the gate electrode in LUX, using data where the grid voltages were raised above their normal operating values. This is the first systematic study of its kind in a xenon TPC, which provides sensitivity to extremely small electron emission currents. The findings contribute to illuminating the nature of these phenomena so that the electrical resilience of future detectors can be improved.

In Chapter 7 a study of the design of the electroluminescence region in LZ is presented. In particular, simulations of electron drifts in the gas phase were carried out for various candidate electrode configurations to assess the impact on the electroluminescence response and resolution of the detector. The various design considerations for the grids are discussed before moving on to describe the simulation methods. Extensive validation studies were carried out, including comparisons to experimental data from LUX and XENON100. I investigated the impact on energy resolution at low energy, which can affect discrimination of electronic recoils, and at high energy, where a good resolution allows spectroscopy of background γ -rays and the possibility to conduct a search for neutrinoless double beta decay.

I conclude with a brief summary of the main findings from this work in Chapter 8.

Currently, as of July 2016, there is no credible detection of WIMPs in any type of experiment. LUX recently presented new results on the spin-independent WIMP-nucleon cross section from Run 4, a 33 500 kg·day exposure [1]. This set the most stringent constraints on the spin-independent WIMP-nucleon cross section, with a minimum of $2.2 \times 10^{-46} \text{ cm}^2$ for a 50 GeV WIMP at 90% confidence. This is a factor of

4 improvement at high WIMP mass over the results from the reanalysis of the Run 3 data. Limits on the spin-dependent WIMP-nucleon cross sections are expected soon, with a similar improvement in sensitivity. Comparable spin-independent results have been presented by the PandaX-II collaboration [2], setting a minimum excluded cross section of $2.5 \times 10^{-46} \text{ cm}^2$ at 40 GeV from an exposure of 33 000 kg·days.

Chapter 2

Dark Matter – Evidence, Candidates, and Detection

In this chapter we begin by summarising the evidence that indicates that over 80% of the mass in the universe is dark matter rather than baryonic matter. The evidence comes from a variety of astrophysical sources, from the galactic to the cosmological scale. Following this the possible dark matter candidates are discussed, of which WIMPs are the leading candidate. WIMP dark matter is consistent with all of the astrophysical observations. Despite this, there has been no conclusive detection of these particles. Lastly, experiments which aim to infer the existence of WIMPs are described, including direct detection, indirect detection, and collider searches.

2.1 Evidence

2.1.1 Galactic Scale

There is strong evidence for the existence of dark matter in individual galaxies from measurements of rotation curves of spiral galaxies. Using Newtonian dynamics, the rotation speed as a function of galactocentric radius is given by:

$$v(r) = \sqrt{\frac{GM(r)}{r}}, \quad (2.1)$$

where $M(r)$ is the mass contained within radius r . As most of the luminous matter is contained within the bulge of the galaxy, $M(r)$ increases very slowly with r away from the central region. Therefore, we roughly expect that $v(r) \propto 1/\sqrt{r}$ and v should be falling with increasing r . However, rotation curves are observed to remain approximately constant out to large distances, even past the outermost stars when only gas is measured. As the mass distribution $M(r) = \int \rho(r)r^2 dr$, for a density profile $\rho(r)$, it implies that the overall mass distribution has $\rho(r) \propto 1/r^2$ and $M(r) \propto r$ above a radius of a few kpc. The disagreement between the expected velocity curve from luminous matter and the observation can be resolved if there is a large mass component that is not visible – dark matter. The contributions to the rotation velocity from each mass component are shown in Fig. 2.1. The dark matter is predicted to be in a roughly ellipsoid shaped halo, which extends above and below the plane of the spiral. There

are many measurements of different spiral galaxies that confirm that there is always at least some dark matter component [3] [4].

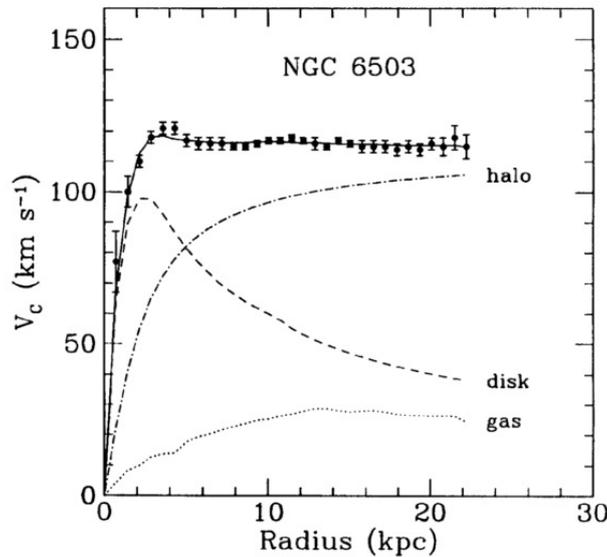


Figure 2.1: The rotation speed of the spiral galaxy NGC 6503. The contributions from the luminous disk, gas, and dark matter halo are indicated. Reproduced from Ref. [5] with permission from EAS Publications Series.

On a larger scale, there is similar evidence from the dynamics of galaxy clusters. Some of the earliest signs of dark matter came from these observations, in particular from the application of the virial theorem to the Coma cluster by Zwicky in the 1930s [6]. The classical virial theorem applies to the situation of a gravitationally bound set of masses provided that the system is in equilibrium, and simply states that the total kinetic energy is equal to (minus) half of the gravitational potential energy. Therefore, by measuring the velocity dispersion of galaxies in the cluster to estimate the kinetic energy contained in the system, the total mass can be inferred. The mass to light ratio was a very surprising 400 solar masses per solar luminosity, meaning that there was a large component of 'dark' matter in the cluster. Many modern measurements also use the virial theorem in determining the mass of galaxy clusters [7] [8] and find similar results.

The mass of a galaxy cluster can also be determined from X-ray emission from the hot inter-galactic plasma. If the cluster is assumed to be an ideal gas in hydrostatic equilibrium a relation between the temperature and mass (within a radius r) can be determined [9]:

$$M(r) = \frac{-kT_{gas}}{Gm_h} \left(\frac{d \log \rho_{gas}}{d \log r} + \frac{d \log T_{gas}}{d \log r} \right), \quad (2.2)$$

where m_h is the mass of a hydrogen atom, T_{gas} the gas temperature and ρ_{gas} is its density. Usually, a parametric description of the cluster gas density profile is used [10] [11]. The parameters of the model are determined by fitting to the surface brightness profile of the cluster. The mass measured by equation 2.2 is much greater than that from the luminous matter (or, equivalently, the gas is much hotter than expected from the luminous matter alone), providing more evidence that there is substantial dark

matter in clusters. X-ray measurements in clusters indicate that dark matter makes up between 80% and 90% of the total mass in clusters [12] [13].

A different method of determining the mass of a galaxy or galaxy cluster without relying on the luminous matter comes from gravitational lensing. In General Relativity light travels along null geodesics in space-time. Massive objects distort space-time around them and light from a background galaxy is bent as it passes by the foreground galaxy, which acts in a similar way to an optical lens. The deflection depends on the exact geometry of the situation and the mass of the foreground galaxy. One regime is strong gravitational lensing, which occurs when there is an almost direct alignment between a background object, a foreground object, and the observer. The light can take multiple paths to reach the observer and this produces multiple distorted arc-shaped images of the background galaxy around the lens. Examples of this effect can be found in Refs. [14], [15], and [16]. A direct alignment can create an Einstein ring around the lens. The required alignment means that strong lensing is a relatively rare phenomenon.

The more common case is weak gravitational lensing, from either a smaller lens or when the background galaxy is far from the lens, which only slightly distorts the source image [17] [18]. The shape of the image is slightly stretched tangentially around the foreground mass. When looking at an individual background galaxy, it is impossible to tell whether the observed shape is due to lensing, or is simply the elliptical shape of the galaxy. Therefore, weak lensing is used with a number of background galaxies. The intrinsic shape of the galaxy is the dominant effect, and this must be averaged out (without any lensing and assuming randomly-oriented galaxies the average shape would be circular). The effect of a lensing distortion can only be seen statistically as a systematic distortion of many galaxies along a particular direction. Weak lensing can be used to map out the mass distribution in individual structures, or to survey the overall mass distribution in the universe. Lensing measurements of individual galaxies find the stellar mass is approximately 5% of the total [19]. Lensing measurements of galaxy cluster masses are in good agreement with those from X-ray emission [20].

A particularly important example is from observations of the Bullet cluster, where two clusters have passed through each other [21], shown in Fig. 2.2. Using X-ray observations the location of the intracluster plasma (by far the dominant contribution to the baryonic mass) was found. Then, with weak gravitational lensing the gravitational potential was mapped out. The measurements found that the two main mass concentrations were separated from the intracluster plasma. This occurs because the ordinary matter is slowed down as it interacts during the cluster collision, whereas dark matter can pass straight through without drag (presumably only interacting gravitationally). The segregation of the mass from the baryonic matter is very difficult to explain with alternative theories, such as modified Newtonian gravity. There have since been other examples of this phenomenon in cluster mergers, see Refs. [22], [23], and [24].

Lastly, and of utmost importance to direct dark matter searches with terrestrial detectors, is evidence of dark matter in our own galaxy and the solar system. Firstly, as the Milky Way appears to be a typical spiral galaxy, there is no reason why it should not be like any other observed and contain dark matter. Although much harder to determine, the measured rotation curve is flat out to the location of the Earth at 8 kpc.



Figure 2.2: Composite image of the Bullet Cluster collision. The pink colour shows the X-ray emission, and the blue shows the mass distribution from the lensing map. Credit: X-ray: NASA/CXC/CfA/M.Markevitch *et al.* [25]; Optical: NASA/STScI; Magellan/U.Arizona/D.Clowe *et al.* [21]; Lensing Map: NASA/STScI; ESO WFI; Magellan/U.Arizona/D.Clowe *et al.*

Measurements in the outer galaxy have larger uncertainties but agree with flatness [26], and this is confirmed at 13 kpc by a more precise measurement from Ref. [27]. The velocity of the fastest stars in the local neighbourhood suggests that the galactic escape velocity is much higher than that expected from the luminous matter alone. The total mass of the galaxy can be measured by studying the kinematics of satellite galaxies or stars in the halo and estimates are of the order $1\text{--}2 \times 10^{12}$ solar masses [28] [29] [30]. An estimate of the stellar mass in the Milky Way indicates 6.4×10^{10} solar masses [31], confirming that the galaxy as a whole is dark matter dominated.

There are several measurements of the local dark matter density. In Ref. [32] ρ_{DM} is between 0.2 and 0.4 GeV, depending on the dark matter halo profile. Other more precise measurements are $0.39 \pm 0.03 \text{ GeVcm}^{-3}$ [33] and $0.32 \pm 0.07 \text{ GeVcm}^{-3}$ [34]. For a review see Ref. [35]. It is important that this is measured precisely for direct detection experiments, as it is degenerate with the WIMP scattering cross section when calculating the scattering rate (see Chapter 5).

In conclusion there is a wealth of evidence for dark matter on the scale of galaxies and galaxy clusters. Techniques for measuring the non-luminous mass using Newtonian gravity, X-rays from intra-cluster gas, and gravitational lensing all give consistent results, using very different experimental techniques.

2.1.2 Cosmological Scales

As well as looking in individual galaxies or clusters, there is evidence of dark matter in the whole universe from the Cosmic Microwave Background (CMB). The CMB radiation comes to us from the epoch when the first atoms formed when the universe was just 380,000 years old. Just before this, the universe was a radiation dominated plasma

mostly consisting of protons, electrons and photons in thermal equilibrium. Once the temperature was low enough to allow recombination to form hydrogen atoms, the universe became transparent to radiation and the photons travelled unimpeded. They appear to us now as a blackbody spectrum with a temperature of 2.726 K (after undergoing a large redshift), coming from a spherical surface around the observer (the surface of last scattering). There are small temperature anisotropies (at the 10^{-5} level) due to nonuniformities in the density at that time, which encode important information about earlier times, and these perturbations would go on to seed the formation of structure as the universe expanded and cooled further. The last scatter of the photons was from the electrons just before recombination, and so the CMB probes the structure of the universe at that epoch.

Before recombination, fluctuations in density cause gravitational potential wells to form, which triggers Baryon Acoustic Oscillations (BAO). Baryons and dark matter gather in these denser and hotter regions, which increases the outward radiation pressure, and this can, at some point, counteract the force of gravity. So, the fluid goes from a period of compression, to one of expansion. Later on, as the region expands and cools, the radiation pressure decreases and compression under gravity restarts. This creates a harmonic oscillator that causes acoustic standing waves in the plasma. These standing waves remain in place until recombination, and are observed in the CMB.

The fluctuations in the temperature spectrum are usually represented as a multipole expansion:

$$\frac{\delta T}{T}(\theta, \phi) = \sum_{l=2}^{\infty} \sum_{m=-l}^l a_{lm} Y_{lm}(\theta, \phi). \quad (2.3)$$

The functions Y_{lm} are the spherical harmonics. The value of interest is the expectation value of the multipole coefficient a_{lm} :

$$C_l \equiv \langle a_{lm} \rangle = \frac{1}{2l+1} \sum_{m=-l}^l |a_{lm}|^2. \quad (2.4)$$

This function is known as the angular power spectrum, which contains all information about the statistical fluctuations in the CMB temperature. The theoretical expectation value is $\langle a_{lm} \rangle = 0$, the same in all directions. The observed spectrum is one sample of this, and will deviate from the expectation. The parameter l is related to the angular size of the anisotropy pattern – small l means a larger angular scale (for example, $l = 2$ corresponds to a scale of 90°) and vice versa.

The most precise measurements of the CMB come from the Planck data [36]. The cosmological parameters are found by fitting a particular model with several parameters to the data. The power spectrum is shown in Fig. 2.3. With these data the second and third peaks are clearly visible.

The various peaks come from the BAO. Structure on the largest angular scale (smallest l) gives the slowest oscillators, and so the first peak corresponds to the scale which only underwent compression once before recombination occurred. The next peak represents the scale which also underwent first expansion. The odd peaks correspond to

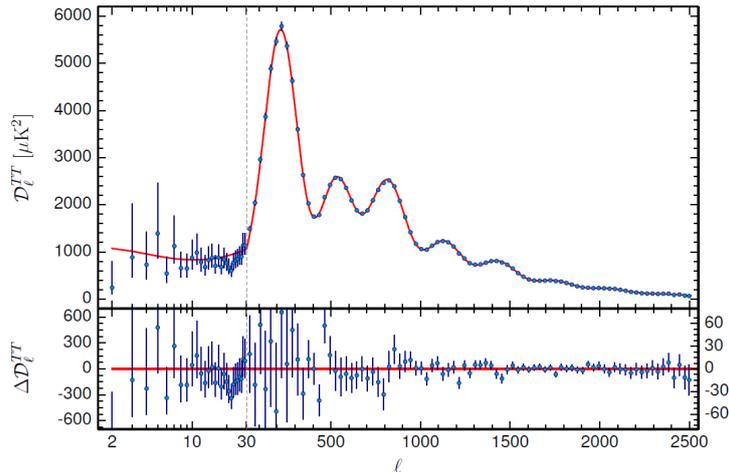


Figure 2.3: The observed CMB power spectrum from Planck and corresponding fit. Figure reproduced from Ref. [36] with permission from ESO.

compression modes and even peaks to expansion modes. The main influence on the first peak is the curvature of the universe. The amplitude of the peaks depends on the overall mass content, as gravity is the driving force of the oscillations. There is sensitivity to how much of the mass is baryonic due to baryon loading: they add mass which enhances the compression peaks (as does dark matter), but for the expansion peaks baryons provide damping as they absorb photons. Therefore, the ratio of the 2nd and 3rd peaks is sensitive to the total baryon content. With the total density and total baryon density, the amount of dark matter can be determined.

For the fit to the power spectrum the energy content is presented as the density parameter, Ω_i : the density of species i divided by the critical density (the density required for flatness). The best fit parameters from Ref. [37] for the energy content of the universe give: baryon density $\Omega_b = 0.049$, cold dark matter density $\Omega_{CDM} = 0.266$ and dark energy density $\Omega_\Lambda = 0.685$. This is consistent with a flat, dark energy dominated universe, described by what has become known as the Lambda-CDM model, the standard model of cosmology – with CDM denoting cold dark matter, which is non-relativistic in the early universe. Dark matter makes up 83% of the matter content in the early universe, which is consistent with the relative amount found in galaxies and clusters.

Another way of studying the large-scale structure of the universe is through gravitational N-body simulations [38]. These simulate the structure formation of the universe from recombination onwards, where the initial conditions are given by the measurements of the CMB anisotropies. The starting point is a number of particles (galaxies, dark matter halos etc.) evenly distributed at some time (redshift). A number of small perturbations are made in the density to represent the fluctuations originating from inflation in the early universe. For each simulation step the force acting on each particle is calculated and the particles are moved. There are methods for reducing the required computation time in this process (rather than simply taking the approach of calculating the force between every pair of particles).

As computing capabilities have increased the scale of the N-body simulations has grown. The simulations allow studies of galaxy formation and the large scale structure of clusters [39] [40] [41]. Therefore, they predict what the structure of the universe should be at a later time, given the conditions in the early universe. The results are compared to galaxy surveys [42] [43] (which can probe galaxies at various redshifts); the simulations reproduce both galaxy formation and overall structure of clusters remarkably well. This confirms that the fluctuations created by the BAO can evolve into the structure observed today, and that cold dark matter is an essential component in structure formation.

2.2 Candidates

Given the evidence that there is a large amount of missing mass in the universe, the next question is, what does it consist of? One of the first possibilities considered was baryonic matter in the form of massive stellar objects. These candidates are collectively known as Massive Astrophysical Compact Halo Objects (MACHOs), which includes black holes, brown dwarfs, neutron stars, white dwarfs and Jupiter sized planets. These non-luminous objects avoid detection by conventional telescopes, but can be searched for using gravitational microlensing. These searches found that MACHOs cannot be the primary constituent of galactic dark matter [44] [45] [46]. Baryonic dark matter is also ruled out by primordial nucleosynthesis in the early universe [47]: almost all of the abundance of light elements can be accounted for in luminous matter.

Another possibility is that instead of missing mass, the theory of gravity needs to be modified. For example, Modified Newtonian Dynamics (MOND) theories were proposed to explain galactic rotation curves, by having different behaviour in the weak field limit. However, it is difficult to explain all observed phenomena with modified gravity theories. They need to explain not only rotation curves, but BAO, structure formation in the universe, and gravitational lensing. In particular, the separation of mass and luminous matter seen in Bullet cluster collision, as described in section 2.1.1, is very difficult to explain without dark matter.

The remaining option is that dark matter is made of elementary particles. We know that these must have mass, interact gravitationally, they must be electromagnetically neutral, and interact very rarely if at all with regular matter. Well over a decade ago the neutrino was still considered a possibility. Due to their small mass they are relativistic, and tend to free stream from overdense regions smaller than ~ 40 Mpc, making it difficult to form structure on that scale [48]. Larger structures could trap the hot dark matter and start forming structure beforehand. However, this would mean supercluster sized structures form before galaxies. This does not agree with observations, which show early formation of galaxies [49] [50]. Therefore, an additional property is that the dark matter particle must be non-relativistic in the early universe, known as cold dark matter. The leading candidate for cold dark matter is the WIMP. There are other candidates of a non-WIMP nature which we discuss later on, after overviewing why WIMPs are so well motivated as a dark matter candidate.

2.2.1 Weakly Interacting Massive Particles (WIMPs)

The WIMP is a stable particle which only interacts with ordinary matter through gravity and very rarely. They are massive, and heavy enough to be non-relativistic in the early universe. These properties are exactly what is required for a fermionic cold dark matter candidate.

None of the Standard Model particles has suitable properties to be a WIMP. Therefore, physics beyond the Standard Model (BSM) is required to motivate suitable WIMP candidates. There are several well-motivated BSM theories, designed to solve various problems that naturally give a WIMP candidate. A selection of these is briefly summarised here.

One of the most well-known BSM theories is supersymmetry (SUSY). It introduces an extra symmetry between fermions and bosons [51] [52]. Each Standard Model boson has a supersymmetric fermion partner, and vice-versa. SUSY was proposed to solve the hierarchy problem in the Standard Model. The mass of the Higgs boson was expected to be at the scale of electroweak symmetry breaking. However, large radiative corrections are expected to the Higgs mass from other Standard Model particles. These corrections are much larger than the Higgs itself, and so there must be very precise cancellations. In the Standard Model there is no reason as to why this fine tuning should occur. With SUSY the contributions from the SUSY partners cancel out those from the Standard Model particles, thus removing the need for fine tuning. Although, with the observed Higgs mass of 125 GeV [53] [54] some fine tuning may still be required [55]. SUSY also facilitates unification of the electromagnetic, weak, and strong forces at high energy. Significantly, the lightest SUSY particle (LSP) is a dark matter candidate if R-parity is conserved. Ordinary particles have $R = +1$, whereas SUSY particles have $R = -1$, so the LSP cannot decay into Standard Model particles. A well motivated particle which could be the LSP is the neutralino, a mixed state between the SUSY partners of the photon, Z^0 boson, and neutral Higgs boson. The neutralino has both scalar interactions (spin-independent) via the Higgs and axial-vector (spin-dependent) interactions mediated by Z^0 or squark exchange.

Another extension to the Standard Model, Little Higgs models were also designed to address the hierarchy problem. In this theory the Higgs is a pseudo-Goldstone boson that arises from a spontaneously broken approximate symmetry. As with SUSY, there are new particles (at the TeV scale) that lead to cancellations of the quadratic corrections to the Higgs mass [56] [57]. However, contributions from fermions are cancelled by new fermions, and bosons cancel with bosons. This is achieved by a collective explicit symmetry breaking. There are two couplings, where each by itself does not break the symmetry (keeping the Goldstone boson massless), but when they are taken together they do break the symmetry. There are constraints on Little Higgs models in the electroweak sector, but these can be alleviated by introducing a new symmetry, T-parity, analogous to R-parity in SUSY [58]. With this, the Standard Model particles are T-even while the new heavy particles are T-odd; the lightest T-odd particle is stable and is likely to be electromagnetically neutral. This particle is most likely a heavy photon (the T-odd partner of the hypercharge gauge boson). This can interact with

ordinary matter at leading order via exchange of a Higgs or \tilde{Q} (a T-odd partner of a quark) exchange, with the latter producing an axial-vector interaction.

A different example of a dark matter candidate comes from theories of universal extra dimensions (UED). These theories postulate the existence of extra compact dimensions, where all fields are allowed to propagate. To a low energy observer these fields appear as an infinite tower of particles, which are quantised modes of the field [59]. Within each tower, all of the particles (Kaluza-Klein particles) have the same quantum numbers, but carry different momenta in the extra dimensions. Ordinary Standard Model particles are the zeroth order modes, which have zero momentum in the extra dimensions. UED models have a discrete symmetry called KK parity, which is even or odd depending on the mode number of the particle. With this, the lightest odd mode must be stable, and if it is both neutral and weakly interacting would be a dark matter candidate [60]. Some possibilities are the Kaluza-Klein mode of a weak gauge boson, neutrino, or neutral Higgs. The Higgs interacts via scalar particles, but the others have axial-vector interactions. The Kaluza-Klein partner to the neutrino can interact via the Z^0 and the gauge boson Kaluza-Klein partner via a quark.

2.2.2 Cosmological properties of WIMPs

One of the most attractive features of WIMPs is that they naturally produce the observed relic density of dark matter in the universe. In the early universe, any species that is interacting sufficiently will remain in thermal equilibrium. WIMPs couple to ordinary matter via annihilation and pair production:

$$\chi\chi \rightleftharpoons qq, \quad (2.5)$$

where q can represent a quark, or any other particle that couples via the weak force. The number density is governed by the Boltzmann equation, which can be written as [61]:

$$\frac{dn}{dt} + 3Hn_\chi = -\langle\sigma v\rangle(n_\chi^2 - n_{eq}^2), \quad (2.6)$$

where σ is the total annihilation cross-section, v the relative WIMP velocity, H is the Hubble constant and n_{eq} is the co-moving number density at thermal equilibrium. The $3Hn$ term represents the dilution due to expansion of the universe. Then $\langle\sigma v\rangle n_\chi^2$ and $\langle\sigma v\rangle n_{eq}^2$ represent the rates of the annihilation and pair production processes, respectively. Production of WIMPs ceases when the universe cools enough that the Standard Model particles do not have enough energy to produce WIMPs of that particular mass. After this point, WIMPs continue to annihilate until the expansion term becomes larger than the annihilation term in equation 2.6. Effectively, the WIMPs become too sparse and no longer interact, resulting in “freeze-out” of the number density. This leads to the relic density we observe today.

Equation 2.6 cannot be solved analytically, but numerical solutions exist [51]. The calculated relic density is accurate to within 10% and is given by:

$$\Omega_\chi h^2 \approx \frac{3 \times 10^{-27} \text{cm}^3 \text{s}^{-1}}{\langle \sigma v \rangle} \sim \frac{10^{-10} \text{GeV}^{-2}}{\langle \sigma v \rangle}. \quad (2.7)$$

For a roughly 100 GeV weakly interacting particle, $\langle \sigma v \rangle = \alpha^2/m_\chi^2 \sim 10^{-9} \text{GeV}^{-2}$, where $\alpha \approx 10^{-6}$ is the weak coupling constant. This gives $\Omega_\chi h^2 \sim 0.1$, which agrees very well with observations. Therefore, a WIMP can naturally produce the required cosmic abundance of dark matter, which is often referred to as the ‘‘WIMP miracle’’ given the extraordinary range of the energy scales which are involved in equation 2.7.

2.2.3 Non-WIMP particle dark matter

There are some particle dark matter candidates that do not fit into the categories mentioned previously. A possible extension to the Standard Model is to add right-handed chirality neutrinos, known as sterile neutrinos [62]. These only interact via mixing (possibly with the standard neutrinos too), and not via the weak force. They could also decay into a Standard Model neutrino and a photon. However, sterile neutrinos are usually light, around the keV scale, and would not be cold dark matter. There are stringent constraints on their properties from cosmological abundance and decay products [63]. There are particular scenarios where they can be cold and account for the structure formation in the universe [64] [65]. There is no solid evidence for the existence of sterile neutrinos, despite some anomalous results that some claim constitute evidence for their existence [66].

Another dark matter candidate is the axion, a particle proposed to solve the strong CP problem [67]: the strong force does not violate CP, but there is nothing in the Standard Model that enforces this. Therefore, it is another fine tuning problem. A solution is to introduce a new symmetry, which is spontaneously broken giving a pseudo-Goldstone boson, the axion. Although the original axion is ruled out as a dark matter candidate by experiment [68], there are viable models (containing ‘‘invisible’’ axions, or axion-like particles) with a large parameter space of masses and couplings which still solve the strong CP problem [69] [70]. In certain scenarios axions can be a cold dark matter candidate [71]. Invisible axions or axion-like particles can interact via the axio-electric effect [72], which is similar to the photoelectric effect. They can also couple to two photons, which results in mixing phenomena between axions and photons in the presence of magnetic field [73]. Both of these effects offer routes to their detection. Current experimental constraints place the axion mass between approximately 1 μeV and 1 meV [74]. Therefore, a large number density is required to make up the galactic density of dark matter, which is possible with a bosonic candidate.

2.3 Detection

There are three ways of detecting WIMP interactions via the weak force, which can be understood from their Feynman diagrams in Fig. 2.4. These are all essentially the

same diagram, but rotated. The first one is the pair production of two WIMPs. This is how collider experiments aim to observe dark matter, by creating it from Standard Model particles. The second diagram is the opposite process, the annihilation of two WIMPs into Standard Model particles. Cosmic-ray telescopes search for extra particles from this process in particularly dense regions of the universe. The last interaction is the scattering of a WIMP, which is what direct detection experiments search for. They detect the energy given to the Standard Model particle by the WIMP. Each method of searching for WIMPs is summarised in the subsequent sections.

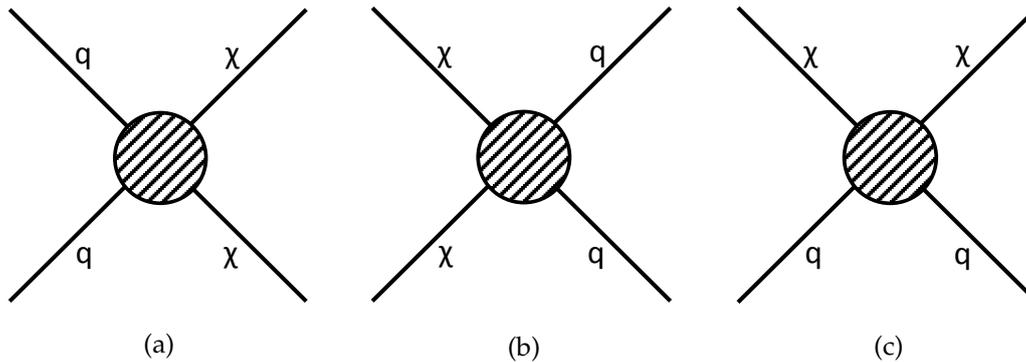


Figure 2.4: Feynman diagrams showing the three interactions used to search for WIMPs; χ represents a WIMP and q a Standard Model particle. The circle represents the interaction, which occurs via some exchange particle depending on the particular theory. (a) Pair production, used in collider experiments. (b) Annihilation, used in indirect searches. (c) Scattering, used in direct detection experiments.

2.3.1 Collider searches

In this section we focus on dark matter searches from the ATLAS and CMS experiments at the LHC. As WIMPs do not interact via the electromagnetic or strong force, they do not leave a signal in detectors at collider experiments. The LHC searches are sensitive to WIMPs and other new “invisible” particles. Their presence from a collision can be inferred in a hermetic detector which completely surrounds the collision point and enables the overall balance of energy and momentum transverse to the beam to be measured. The key quantity used is the transverse energy, in the plane perpendicular to the beam, because the longitudinal momentum of the partons is unknown. If all of the particles from a collision are detected, the transverse energy should be zero. However, if WIMPs or other “invisible” particles are produced and escape detection they show up as missing transverse energy. Collider searches cannot tell if the particle is the astrophysical dark matter, because they are not able to measure the lifetime of the particle or the astrophysical abundance.

To interpret the LHC results an effective field theory (EFT) treatment is often used and the WIMP-Standard Model interaction is described in terms of a contact interaction. There are several possible operators, depending on the type of interaction. Each depends on two parameters; the WIMP mass (m_χ) and the energy scale $M_* = M_{med} / \sqrt{g_\chi g_q}$, where M_{med} is the mediator mass, g_χ is the coupling to WIMPs and g_q is the coupling to quarks. Results are typically presented as lower limits on M_* as a

function of m_χ for each operator [75] [76]. However, this treatment is only valid at low momentum transfer (smaller than the mediator mass). For higher energies more specific models are required. One such treatment is with Minimal Simplified Dark Matter (MSDM) models, which only use four free parameters: m_χ , M_{med} , g_χ , and g_q [77]. This framework also allows comparison between collider and direct searches.

A large class of searches involves the production of a pair of WIMPs in association with initial or final state radiation (ISR/FSR). A general Feynman diagram for the process is shown in Fig. 2.6. The presence of ISR or FSR is required because otherwise the two WIMPs are produced back-to-back, which leaves negligible missing transverse energy. The signature is therefore an object with high transverse momentum and a large missing transverse energy [75] [78].

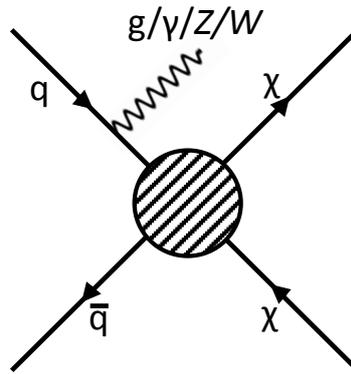


Figure 2.5: Feynman diagram of WIMP production with associated initial state radiation. The circle represents the contact interaction in the EFT approach.

The most common of these analyses are monojet searches, where the WIMP is produced with an ISR/FSR quark or gluon which subsequently hadronises to produce a jet. The requirements are at least one jet (with strict reconstruction quality requirements) with high transverse momentum (thresholds typically 110 to 500 GeV) and large missing transverse energy. An irreducible background is from Z^0 boson production with a single jet, where the Z^0 decays into two neutrinos. The other main background is from W bosons with a single jet, when the W decays into a lepton plus neutrino and the lepton is not reconstructed. These are estimated from side bands in the data. So far the data are in agreement with the Standard Model predictions [79] [80].

Although the monojet searches are the most sensitive, due to having the largest production rate giving high statistics, other searches can provide sensitivity to different models. With an ISR/FSR photon the analysis is quite similar, requiring a single photon with high transverse momentum. The main backgrounds also come from Z^0 or W production with an ISR/FSR photon. Alternatively, there are searches with an associated Z^0 or W boson. These can be either with leptonic or hadronic decays of the boson. If there are observations in different channels, it allows the coupling of WIMPs to be investigated, although no deviation from the Standard Model expectations is currently observed [81] [82].

A different search for dark matter in colliders is from invisible Higgs decays. A Feynman diagram for the production is shown in Fig. 2.6. The Higgs is produced via

vector boson fusion and decays into two WIMPs. The two quarks undergo only a small deflection, and produce jets in the forward regions of the detector while the WIMPs leave a large missing transverse energy. These searches are sensitive to the coupling of WIMPs to the Higgs. Currently the upper limits on the invisible branching fraction to invisibles are 0.28 and 0.24 from ATLAS [83] and CMS [84] respectively.

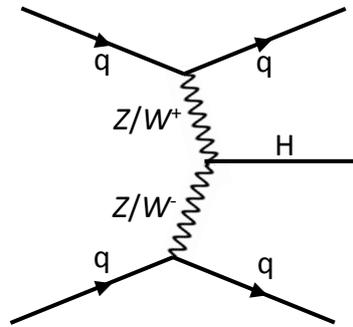


Figure 2.6: Feynman diagram of Higgs production via vector boson fusion. The Higgs can decay into two WIMPs.

2.3.2 Indirect searches

Indirect searches look for a dark matter annihilation signal in the cosmos. This can take place in very dense regions, where gravitational attraction forces dark matter together. Therefore, ground- or space-based cosmic-ray telescopes look for excess emission from regions expected to contain dark matter such as the galactic centre, diffuse emission in the galaxy, the galactic halo, dwarf galaxies, galaxy clusters and even compact objects such as the Sun. Typically, the signals are γ -rays, charged leptons, or neutrinos [85]. The general method is to find either a spectral feature or excess of emitted particles from the astrophysical source. This requires detailed knowledge of background sources from astrophysical objects, which is particularly challenging.

Production of γ -rays from WIMPs can come from annihilation into quarks and the subsequent hadronisation and pion production processes. This would be a continuum spectrum, and so it requires searching for a broad excess of events. Alternatively, a direct annihilation of WIMPs into γ -rays gives a peak in the spectrum, which would be difficult to explain from astrophysical sources [86]. This also has the advantage of being amenable to the use of sidebands in the data to determine the background. As γ -rays are not charged they point back to the astrophysical source.

Neutrinos from dark matter annihilation would also point to the astrophysical source. Due to their low interaction probability they can traverse matter without interacting. To detect them requires large Cherenkov detectors such as Super-Kamiokande or Ice-Cube. Although they can be used to look for WIMP annihilation from objects such as galaxies and clusters, the main target in neutrino searches is the Sun, which can trap WIMPs in its gravitational potential. Only neutrinos are able to directly escape without interacting [87].

With charged particles the searches usually focus on fluxes of antiparticles, as they are much rarer in the universe. Like γ -rays, they can be produced from dark matter annihilation's into Standard Model particles. The signal would be an increased antiparticle to particle ratio. However, charged particles are affected by diffusion in the galaxy, which alters their energy spectrum and means they do not point towards their source. These searches are for excesses in the spectrum of particles that arrive at the Earth. This requires modeling of the sources in the galaxy, and their propagation.

This is a very active field and we now discuss the current experimental status focusing on claimed signals and their possible origins [88].

There have been several observations of excesses in charged leptons with respect to expectations. Firstly, the PAMELA space telescope observed a steep increase in the spectrum of $e^+/(e^++e^-)$ between 10 and 100 GeV [89]. This was confirmed by FERMI [90] and AMS-2, [91] which extended the range to 400 GeV. There are also hints in the e^++e^- spectrum from the combined results of FERMI [92] [93] and HESS [94] [95]. Conversely, no excess was seen in the antiproton channel by PAMELA [96]. If the excess is interpreted as a dark matter annihilation signal the results indicate the properties of the WIMP [88]: mass between 1 TeV and a few TeV, almost exclusively decaying into leptons, and a very large annihilation cross section. However, there is significant tension with other observations. With a standard dark matter profile in the galaxy the best fit for the WIMP mass and annihilation cross section is completely excluded, and there is still tension with different dark matter profiles [97] [98] [99]. The annihilation cross section is also larger than that implied by CMB observations, as the extra particles from annihilation would affect recombination through reionisation [100]. Possible astrophysical explanations for the excess include pulsar sources [101] or nearby supernovae remnants [102].

In the FERMI space telescope data a possible γ -ray line at 130 GeV was observed from an extended region around the galactic centre [103] [104]. There was also an observation of a similar line in galaxy clusters [105]. However, there have been many non-dark matter interpretations of these results, such as instrumental effects or astrophysical sources, or statistical fluctuations [106] [107] [108]. The significance decreased as more data were acquired, and it is likely that the line is an unfortunate combination of instrumental effects and a statistical fluctuation.

Another possible signal is a γ -ray excess at 2–3 GeV from around the galactic centre, originally claimed in Ref. [109] and by others since. The signal is compatible with dark matter annihilation from 31–40 GeV WIMPs with a standard galactic dark matter profile and annihilating into $b\bar{b}$ pairs with $\langle\sigma v\rangle = 7 \times 10^{-26} \text{ cm}^3/\text{s}$ [110]. Although this claim is statistically significant it relies on the accuracy of the particular background model, particularly for the diffuse γ -rays. Astrophysical explanations for the excess include a millisecond pulsar population [111] [112], a spectral break in emission from the central black hole [113], or addition of charged particles in the past producing extra secondary radiation [114] [115]. The strongest constraints on the dark matter nature of the signal come from searches in the antiproton channel, as the excess implies a hadronic annihilation process. With realistic models for galactic propagation and effects of modulation by the Sun's magnetic field, the antiproton data are in tension with the galactic excess signal [114]. Recently, new analyses find a population of unresolved

γ -ray point sources that can account for the excess [116] [117] and the dark matter explanation is ruled out.

An unexplained X-ray line was observed in data from the XMM-NEWTON satellite at 3.5 keV [118] [119]. Around this energy there are many spectral lines from de-excitations of several radioisotopes, and it is possible that the 3.5 keV signal could be from a potassium line [120]. However, some of the observations were in the Andromeda cluster where the plasma temperature is low and no de-excitation lines are expected in this region. The signal could be from decays of a 7 keV mass sterile neutrino, which is not currently ruled out [88], but requires further investigation.

No excess over background from indirect searches from solar neutrinos have been found, but limits are set on the WIMP annihilation cross section. The searches are in various channels, of WIMP annihilations into different Standard Model particles, which subsequently decay into neutrinos. The annihilation cross section is related to the capture rate of WIMPs in the Sun, which in turn depends on the WIMP scattering cross section [85]. Therefore, the results can also be interpreted as limits on the WIMP scattering cross section. The dominant capture process is spin-dependent scattering on hydrogen, and so the indirect limits are very competitive in the proton-only channel. Currently, the most stringent limit on the spin dependent proton-only WIMP scattering is from the IceCube experiment, with a minimum excluded cross section of $3 \times 10^{-41} \text{ cm}^2$ in the $\tau\bar{\tau}$ channel [121]. Limits are also set on the spin-independent scattering cross section, but direct detection experiments are more sensitive in this channel.

There is no indirect signal that is conclusively from dark matter annihilation. Uncertainties in the underlying models of the astrophysical sources, and of propagation through the galaxy in the case of charged particles, coupled to the need to understand the response of these telescopes extremely accurately, means that finding robust evidence for a signal is challenging.

2.3.3 Direct detection

If dark matter is composed of WIMPs, they should be constantly passing through the Earth in great numbers. Direct searches aim to detect the interaction of these WIMPs with a nucleus. The WIMP scatters off a nucleus via the weak interaction, possibly a simple elastic scattering. This causes the nucleus to recoil and interact in the surrounding material. This can produce prompt scintillation, ionisation and heat (phonons). One or more of these signatures can be detected depending on the technology. The energy from a nuclear recoil is very small: for a slow moving WIMP with a mass of 10–1000 GeV it is only $\lesssim 100 \text{ keV}$ [122]. Given existing exclusion limits [1] the interaction rate is less than $\sim 1 \times 10^{-5} \text{ events}/(\text{kg}\cdot\text{day})$. There are several backgrounds which can mimic a WIMP interaction. As the signal rate is so small the sensitivity to WIMPs is often background-limited. Therefore, to unambiguously identify a WIMP signal is challenging, and requires detailed understanding of the backgrounds.

There are various ways of searching for a dark matter signal in direct detection experiments. The first is to measure the nuclear recoil spectrum. The differential event rate, which is calculated later on in this thesis, depends on the velocity of WIMPs with

respect to the Earth, the WIMP scattering cross section, the WIMP and target nucleus masses, the type of interaction, and the details of the nuclear structure. The two interactions typically considered are spin-independent (scalar mediator) and spin-dependent (axial vector mediator). The latter can be further split into a neutron-only or proton-only interaction. There are other possible interactions, and WIMPs could scatter via more than one channel. This is discussed further in Chapter 5. Measuring the recoil spectrum is effectively looking for an excess above background that is consistent with a WIMP signal.

Another way is to search for an annular modulation in the signal [123]. The interaction rate and spectrum depend on the relative velocity of the galactic WIMPs and the Earth. This reaches a maximum when the Earth is travelling in the same direction as the solar system relative to the galaxy, which occurs around June 2nd. Six months later the opposite occurs, giving a reduction in the rate. This results in a few percent variation in the WIMP flux over the year. Most background sources are not expected to show this type of modulation.

The main aspects that affect the sensitivity of a direct detection experiment are the exposure (mass \times live time), the energy threshold for nuclear recoil detection, and background levels in the keV region. Long exposures require a detection technology capable of high operational duty cycle and stability. Increasing the target mass of the detector as much as practicable is also essential. Lowering the detection threshold increases the sensitivity for all WIMP masses and particularly helps for low mass WIMPs. Thresholds of a few keV are possible even in large detectors.

Some backgrounds produce nuclear recoils (NR) and if the interaction is a single scatter these are indistinguishable from WIMPs. Alternatively, and more commonly, electronic recoils (ER) are produced by γ -ray interactions (Compton scattering or photoelectric effect) or from charged particles such as β -decay electrons. There are background rejection techniques that are common to most direct detection experiments. Firstly, due to the low probability of a WIMP interaction, the signal has to be compatible with a single energy deposit and any multiple scatters are rejected. If the detector has the ability to determine the location of an interaction, it can reject regions of higher background, such as near to the edges. This is known as fiducialisation. The response of a material to an interaction depends on the particle that caused the recoil. Different particles deposit different fractions of their energy into scintillation, ionisation, and heat. Therefore, a detector that is sensitive to two of the channels can discriminate between nuclear recoils and other interactions. The timing of the scintillation response alone can also provide excellent discrimination in liquid argon. Being able to discriminate between ERs and NRs with high efficiency is a powerful tool for reducing backgrounds.

Now we discuss some of the backgrounds common to all direct detection experiments, which are typically the external sources, and how these can be mitigated. Internal backgrounds depend on the detector materials and target used, and so are particular to each experiment.

Neutron backgrounds can be particularly dangerous, as single, elastic scatters of neutrons cannot be discriminated against. Cosmic-ray muons create neutrons by spal-

lation reactions on nuclei in the detector or surrounding rock. Direct detection experiments are placed deep underground to reduce this background [124]. An active muon veto is often employed to reject interactions in the inner detector that are coincident with a muon. The veto is typically an instrumented water tank or plastic scintillator panels. Neutrons from radioactivity are also a significant concern. MeV neutron sources include ^{238}U spontaneous fission and (α, n) reactions from the thorium and uranium decay chains [125]. External neutron backgrounds are mitigated by detector shielding, which can be provided by a water muon veto or dedicated passive shielding using light (hydrogen-rich) materials. In active neutron vetoes, these materials are often doped with isotopes having large neutron capture cross sections such as ^{157}Gd . To reduce neutrons from inside the detector materials with extremely low ($< \text{ppb}$) thorium and uranium content must be chosen.

Another background comes from γ -rays, which interact via the photoelectric effect or Compton scattering. These ER backgrounds can only be rejected with a finite efficiency. Several γ -rays originate from the ^{238}U and ^{232}Th chains, from keV up to a few MeV. Other γ -rays come from common radioactive isotopes such as ^{40}K , ^{60}Co , and ^{137}Cs [126]. Similarly to the neutrons, many γ -rays are external to the detector from the surrounding rock or from radon in the air. Shielding the detector with high Z materials or a large water tank can effectively reduce the external background. Radio-pure materials must be selected to reduce γ -rays from inside the detector. High purity germanium detectors are used to screen candidate materials.

Although not currently significant, as detectors become larger and more sensitive they can start to detect neutrino interactions [127]. Neutrinos can scatter elastically and produce ER signals, which provide a background to WIMP searches. The dominant contribution to this process is from solar pp neutrinos. Neutrinos can also interact via coherent elastic scattering off nuclei, which produces NR of a few keV, although this process is yet to be observed. At very low energy, the rate is predominantly from ^8B solar neutrinos, which can mimic a signal from a WIMP with a mass of ~ 5 GeV. At higher energies there are atmospheric and diffuse supernovae neutrinos, which leave similar signals to a ~ 100 GeV WIMP. Therefore, this is an irreducible background and limits the maximum sensitivity of direct detection. This so-called neutrino floor is from solar neutrinos up to WIMP masses below 10 GeV and from atmospheric neutrinos above that [128] [129]. For a xenon target, this limits the sensitivity to approximately 10^{-46} cm^2 for WIMP masses below 8 GeV. It reaches a minimum of $2 \times 10^{-49} \text{ cm}^2$ at 30 GeV before steadily increasing again, reaching $3 \times 10^{-48} \text{ cm}^2$ at 1000 GeV [130].

There are many different experimental techniques used in the direct detection of dark matter and many experiments have performed searches for dark matter. Some examples and key results are briefly mentioned here; a more thorough review can be found in Ref. [131].

The DAMA/LIBRA experiment uses a detector consisting of several low radioactivity NaI(Tl) crystals, with a target mass of 233 kg for the second phase of operation [132]. It is only sensitive to the scintillation signal and there is no discrimination between NRs and ERs. Data have been collected over 14 years, and an annual modulation signal is observed in the 2–6 keV range [133]. The maximum is compatible within 2σ of June 2nd. The best fit is a WIMP with a mass of 10–100 GeV and spin-independent

WIMP-nucleon cross section, σ_{SI} , of 10^{-41} – 10^{-39} cm² [134]. However, this result is in tension with many other experiments that do not observe a compatible signal. Also, the KIMS experiment which used CsI crystals did not find an annual modulation signal [135], which disfavors the DAMA signal being from interactions with iodine. There are new experiments with the same crystal material, NaI, that aim to cross-check the DAMA result. The SABRE experiment [136] is proposed to operate in the same laboratory as DAMA, whereas Anais [137] and DM-Ice [138] will be placed elsewhere.

One of the most sensitive detector technologies at low masses are cryogenic bolometers, which record phonons and either the charge or light signal. The most competitive experiment with this technology is SuperCDMS [139]. It consists of stacks of cylindrical germanium crystals cooled to a few mK, with specifically developed iZIP detectors [140] [141] on their flat surfaces, which detect ionisation and athermal phonons. This allows both the ability to reconstruct the position of an interaction and discrimination against ERs. An advantage of this technology is that the energy threshold is very low, allowing high sensitivity to low WIMP masses. A 577 kg·day exposure sets a minimum excluded limit on σ_{SI} of 1.2×10^{-42} cm² for a 8 GeV WIMP. In addition, a single crystal from SuperCDMS, called CDMSlite, was used to perform a search with a dedicated low threshold. This operated in a mode which enhanced the phonon signal and allowed an even lower energy threshold of 56 eV. However, this mode did not allow a simultaneous measurement of the charge and could no longer discriminate ERs. From a 70 kg·day exposure a minimum exclusion on σ_{SI} of almost 10^{-41} cm² was set for WIMP masses between 4–6 GeV [142], and is the most stringent direct detection limit below 4 GeV.

For WIMP masses above 4 GeV the most constraining limits are from LUX Run 4, a 33500 kg·day exposure, which sets a minimum exclusion on σ_{SI} of 2.2×10^{-46} cm² for 50 GeV WIMP [1]. This detector, a dual-phase xenon TPC, is described in Chapter 3. A similar exclusion limit was also set by another dual-phase xenon TPC, PandaX-II [2]. LUX is also the most sensitive to the spin-dependent WIMP-neutron scattering cross section, $\sigma_{SD,n}$, which is discussed in Chapter 5. The minimum excluded $\sigma_{SD,n}$ is 9.4×10^{-41} cm² at a 33 GeV WIMP mass.

For the other spin-dependent channel, the WIMP-proton scattering cross section $\sigma_{SD,p}$, the most sensitive detector is PICO-60, a bubble chamber. A superheated fluid is kept just below its boiling point and particle interactions cause bubble nucleation. This requires the energy density of the deposit to be larger than some critical value. The bubble is photographed with a CCD camera, so the location of an event can be determined. After each event the detector is reset by a compression and decompression. The bubble nucleation threshold renders these detectors insensitive to minimum ionising backgrounds, such as γ -rays, X-rays and electrons. Acoustic signals are used to distinguish α -particle events, which are louder than nuclear recoils. PICO-2L was a two litre C₃F₈ target [143] and PICO-60 was a 60 litre target of CF₃I [144]. The excellent sensitivity in the spin-dependent proton only channel is due to the fluorine atoms, which have an odd number of protons. The minimum excluded $\sigma_{SD,p}$ is 5×10^{-40} cm² for a 100 GeV WIMP, from a 3415 kg·day exposure [144].

Upcoming experiments are set to further probe the WIMP scattering parameter space. DEAP-3600 [145] is a single phase detector, consisting of a sphere of liquid

argon surrounded by PMTs, which gives excellent light collection efficiency. Only the scintillation signal is detected, but excellent discrimination is still possible. In liquid argon the two excited states that produce the scintillation light have significantly different time constants: 6 ns and 1.6 μ s. ER give a larger yield for the early time constant, whereas for NR a larger fraction is in the late time constant, allowing pulse shape discrimination [146]. ER can be rejected at the 1 in 10^8 level [145]. A three year physics run is due to start in 2016, with an expected sensitivity for σ_{SI} of 1×10^{-46} cm² at 100 GeV from 3000 kg year exposure.

A detector currently undergoing commissioning is XENON-1T [147]. This is a dual-phase xenon TPC with a 1 tonne fiducial volume. The expected sensitivity is a minimum σ_{SI} of 1.6×10^{-47} cm² at 50 GeV for 2 years live time. In addition, the LZ detector is currently under construction [130]. Expected to begin commissioning in 2020 it will have a 7 tonne active xenon target. The sensitivity to σ_{SI} is expected to be 2.5×10^{-48} cm² at 40 GeV for a 1000 day run [130]. The technology of dual-phase xenon TPCs is described in Chapter 3 in relation to LUX and LZ. The upcoming experiments are set to improve over current results, which are shown in Fig. 2.7.

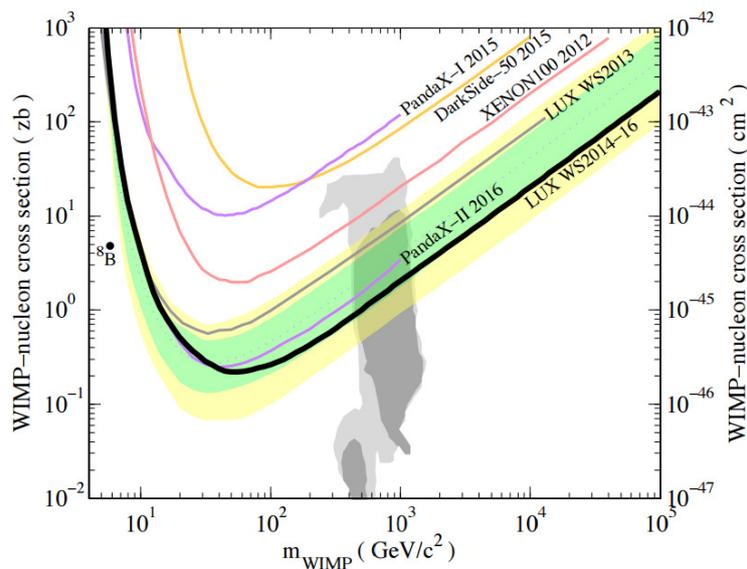


Figure 2.7: Upper limits on the spin-independent WIMP-nucleon cross section. The current world-leading result is from the 2014–2016 LUX data (black) [1]. The reanalysis of the first LUX data from 2013 [148] is the grey line. Other experiments shown are DarkSide-50 [149], PandaX-I [150] and XENON-100 [151]. The dark and light grey regions show the favoured SUSY CMSSM parameter space (1σ and 2σ respectively) [152]. Figure reproduced from Ref. [1].

2.3.4 Complementarity of detection methods

The sensitivity to dark matter from each of the different search techniques depends on the underlying physics model and interactions, and it must be borne in mind that each technique has its own sources of systematic uncertainty. Even in direct detection experiments some target elements are better for particular interactions [153]. For example, in the most common spin independent scattering case the scattering rate is enhanced

(by $\sim A^2$) for heavier elements due to the coherent nature of the interaction across the entire nucleus. However, a spin dependent interaction requires nuclei with non-zero spin, and so it favours elements with an odd number of protons and/or neutrons (such as ^{19}F , ^{129}Xe , ^{131}Xe). This motivates experiments with different target nuclei to cover the possible interactions.

Collider experiments can be interpreted in terms of the WIMP-nucleus scattering cross section [77]. Direct detection experiments are more sensitive than collider searches at high masses, as colliders are limited by the centre of mass energy available. Conversely, colliders are more sensitive to low mass WIMPs, which only leave small energy deposits in direct detection experiments and are more severely affected by the energy threshold. Indirect detection can also set limits on the WIMP-nucleon scattering cross section, with the most sensitive being from the solar neutrino searches (which probe the spin-dependent interaction). Also, signals from collider experiments would not confirm that the particle is long-lived.

Ideally an observation of WIMP dark matter in one type of search would be confirmed by another. This would provide independent confirmation of its existence. In addition, it would allow the properties WIMP to be investigated, such as its mass, spin, and possible interactions.

Chapter 3

Direct Detection with Dual-Phase Xenon

Dual-phase xenon time projection chambers (TPCs) consist of a liquid xenon target with a layer of gaseous xenon above the liquid surface. They detect the scintillation and ionisation signals from NRs and ERs, allowing an accurate reconstruction of energy and position down to O(keV) energy deposits with excellent ER discrimination capability.

Liquid xenon has several advantageous properties for WIMP detection. Primarily, it is itself a scintillator that is transparent to its own light. The scintillation and ionisation yields are very high, allowing low thresholds. For spin-independent WIMP-nucleus scattering there is an A^2 enhancement to the rate due to the coherent nature of the interaction, which gives heavy elements such as xenon an advantage. There are also naturally occurring isotopes with an odd number of neutrons giving spin-dependent sensitivity. The high density means that the interaction lengths of γ -rays, X-rays, and neutrons are short, which gives excellent self-shielding against radioactivity from outside of the detector. A liquid medium is also much easier to scale up than a crystal-based detector. Xenon has no long-lived radioisotopes and most dispersed contaminants can be removed prior to or during operation. These properties make xenon ideal for detecting small signals from WIMP recoils and being able to maintain a low-background experiment.

In this chapter we describe the method of detecting recoils in a liquid xenon TPC and the various processes involved. Following this is a description of the LUX detector and its various sub-systems. The last section covers the LZ detector and the differences between this new experiment and LUX.

3.1 Particle Detection in Liquid Xenon TPCs

Nuclear and electronic recoils deposit energy in the liquid xenon via the same mechanism. The recoiling particle, a xenon nucleus in the case of NR or an electron (either ejected from an atom or from a β^- decay) for ER, causes excitations and ionisations of the surrounding atoms. The deposition from a NR is over a smaller area than for ER [154]. Energy is also lost in collisions below the threshold to cause excitation, which only produce atomic motion and are not detected. The excited atoms relax producing

prompt scintillation light, which is detected at PMT arrays and is known as the S1 signal. Under an applied electric field the electrons released from the ionisations can be separated from the interaction point and drifted upwards through the liquid xenon. At the liquid surface, the field extracts the electrons into the gas phase. The electrons are drifted towards the anode, colliding with atoms in the gas producing electroluminescence light; this secondary scintillation, also detected by the same PMT arrays, is the S2 signal. Alternatively, at the initial interaction the released electrons can recombine with a xenon ion, which then produces a prompt scintillation photon instead of an electron.

The time between the S1 and S2 pulses corresponds to the drift time of the electrons and gives the depth of the interaction. The S2 light will be concentrated on several PMTs around the location of the extracted electrons. This hit pattern is used to reconstruct the position in the horizontal plane of the TPC. The 3D position reconstruction allows fiducialisation: interactions near to the edges of the active xenon volume can be removed, where there are more backgrounds from radioactivity within the detector materials. Interactions involving multiple scatters, such as those from neutrons or γ -rays can also be rejected in a WIMP search.

A schematic diagram of the energy deposit mechanisms is shown in Fig. 3.1, while Fig. 3.2 illustrates how the dual-phase xenon TPC responds to one interaction. An example event from LUX data is shown in Fig. 3.3. The following sections describe the processes in more detail.

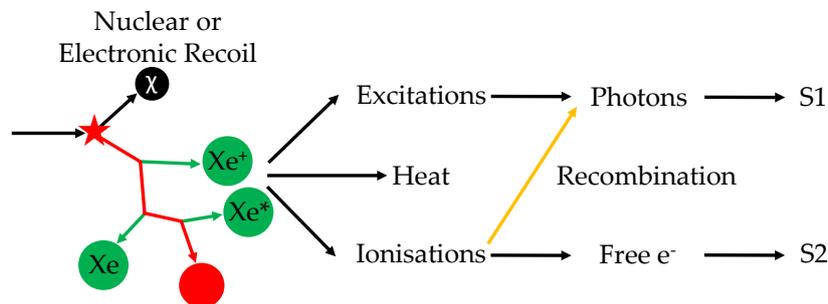


Figure 3.1: Energy deposition from a nuclear or electronic recoil in liquid xenon. The recoiling particle (red), either a xenon nucleus (or atom) for NR or electron for ER interacts with the surrounding xenon atoms.

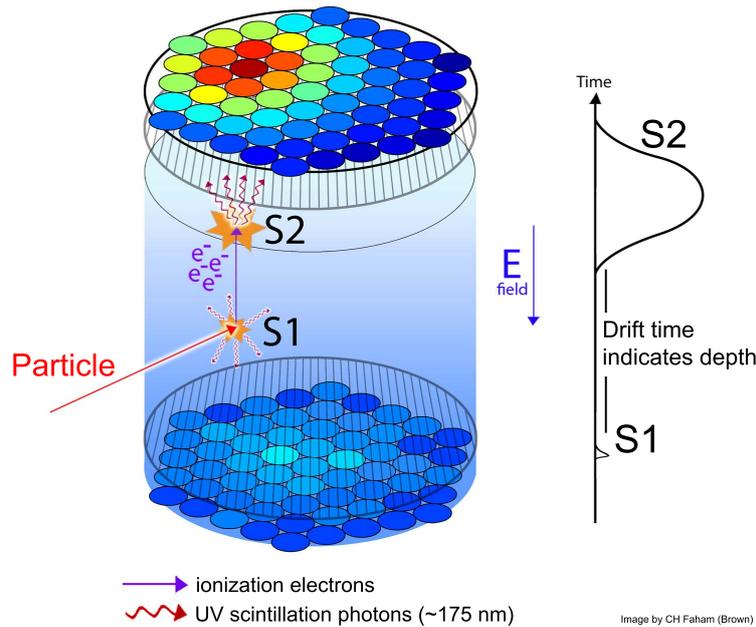


Figure 3.2: Recoil event in a two-phase xenon detector. Primary scintillation (excitations and recombined ions) gives the prompt S1 signal. Ionised electrons escaping recombination are drifted upwards through the liquid and extracted across the surface. In the gas, these produce electroluminescence photons that form the S2. Light is detected at the PMT arrays. The hit pattern of the S2 on the top array gives the x-y position of the interaction, and the drift time between S1 and S2 is proportional to the depth. Figure reproduced from Ref. [155] with permission from Elsevier.

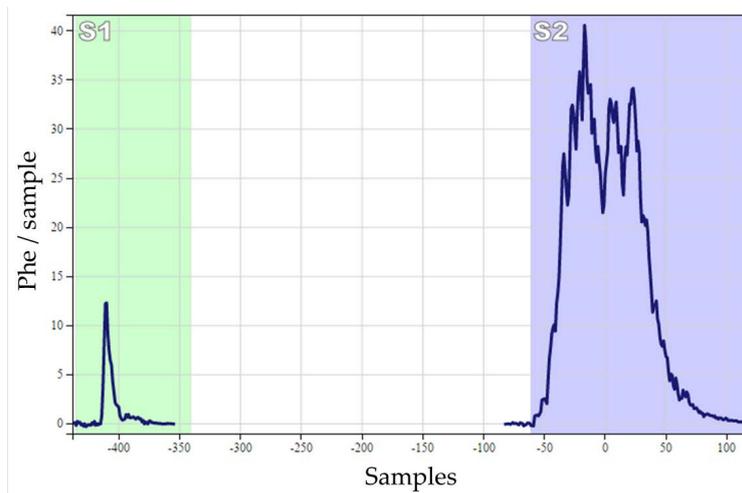
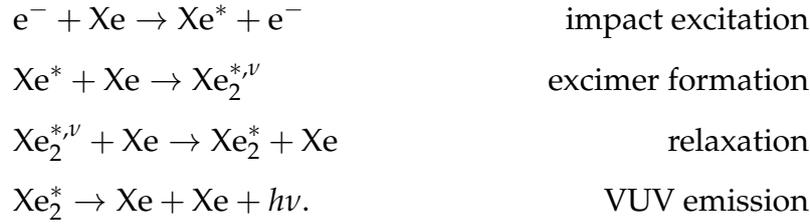


Figure 3.3: Waveform from a single scatter event in LUX. The summed response of all PMTs is shown. One sample is 10 ns.

3.1.1 Prompt scintillation signal

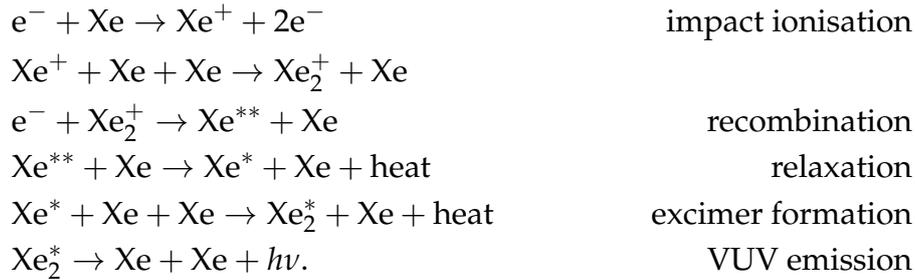
The scintillation light, with a wavelength between 174-178 nm [156] and a spectral width of 14 nm [157] is produced via the formation of the xenon excimer Xe_2^* . The

mechanism is as follows:



The initial excitation forms an excimer dimer with another xenon atom in the second step, which also has vibrational energy (superscript ν). The dimer is in one of the two lowest electronic energy states, the singlet with a decay time constant of 2.2 ns or the triplet state with 27 ns [158]. It de-excites to the dissociative ground state producing a VUV photon. The time constant gives, to first approximation, the characteristic shape of the S1, such as that shown in Fig. 3.3.

The other mechanism that produces scintillation light is recombination of the released electrons with xenon ions:



The electrons recombine with di-atomic xenon ions formed at the interaction. After recombination, the de-excitation process is similar to the direct excitation that followed in the scintillation process, ending with the production of a VUV photon instead of releasing an electron but with a longer time constant of 45 ns [159]. Recombination is both recoil energy and field dependent, with a higher field extracting more electrons from the interaction [154].

As the VUV photons come from the de-excitation of the excimer their energy is lower than the lowest excitation energy of the single xenon atoms. Crucially, the xenon dimer ground state is essentially dissociative (it has a very large equilibrium atomic distance) and so there are no Xe_2 dimers capable of reabsorbing the photon. Therefore, xenon is transparent to its own scintillation light. The light can still be attenuated by impurities in the xenon, but attenuation lengths much greater than the dimensions of a TPC can be achieved [160].

3.1.2 Electron transport and cross-phase extraction

The applied electric field in the liquid serves to separate electrons from the ionised xenon atoms at the interaction site and transport them towards the liquid surface. The

drift velocity of the free electrons in the liquid is between 1–3 mm/ μ s for fields above 100 V/cm [161].

As electrons drift through the liquid they can be lost to electronegative impurities. Even contaminations at the part-per-billion level can significantly attenuate the drifting electrons [162], which sets strict purity requirements. Contaminants are electronegative molecules such as O₂, H₂O, and N₂O. The fraction of attenuated electrons is given by:

$$\frac{N}{N_0} = e^{-t/\tau}, \quad (3.1)$$

for an initial number of electrons N_0 , with N remaining after a drift time t . The mean free electron lifetime τ should ideally be larger than the maximum drift time in the detector (of the order hundreds of μ s).

During transport the electrons also undergo diffusion. With an applied electric field the transverse component (represented by a diffusion coefficient D_T) is larger than the longitudinal component (D_L). The spread due to diffusion is described by:

$$\sigma_{T,L} = \sqrt{2D_{T,L}t}, \quad (3.2)$$

Measurements of the diffusion coefficients are: $D_T \approx 30$ cm²/s for electric fields between 0.5 and 1.2 kV/cm [163] and $D_L = 12 \pm 1$ cm²/s [164]. Typically, this effect smears an ionisation point source by ~ 1 mm over 1 m of drift.

Once electrons reach the liquid surface, they must overcome a potential barrier of a fraction of 1 eV to be extracted into the gas. An applied field is required, as the potential of a free electron in the liquid is lower than that in the gas. The electric field increases the mean of the Maxwellian velocity distribution of the free electrons and gives a preferred direction, some extra momentum perpendicular to the surface. Also, as the electric field is increased the potential barrier becomes smaller and the maximum of the potential moves closer to the liquid surface [165]. Therefore, a high electric field increases the probability of an electron being extracted and prevents electrons being re-absorbed by impurities into the liquid. It takes ~ 6 kV/cm under the liquid surface for a near-unity emission probability.

3.1.3 Secondary scintillation light

Once the free electrons are emitted into the gas phase, the lower density and higher electric field (due to the relative permittivity of liquid xenon of 1.96) means they are easily accelerated, reaching a high enough energy between collisions to excite the xenon atoms and produce electroluminescence light. The process is effectively a random walk, but with a strongly preferred direction after each collision due to the electric field.

For typical operating pressures of dual-phase xenon TPCs, above 1 bar, the excitations produce excimers. The de-excitations are from one of the two lowest states of the excimer to ground, a similar mechanism to that for the primary scintillation in the liq-

uid. Therefore, the emitted photons are in the VUV region with a similar wavelength as the prompt scintillation light, on average $\lambda = 171$ nm [166].

The light yield per electron in the gas phase depends on the electric field and pressure. There are several measurements of the electroluminescence yield available, such as Refs. [167] and [168]. There are differences between measured data from different authors, with the results usually being below the theoretically predicted values [169]. The data from Ref. [167] were taken in xenon vapour at low temperature in equilibrium with the liquid phase and is parameterised in Ref. [170]. This gives the number of photons emitted per cm as a function of the electric field and pressure:

$$\frac{dN_{ph}}{dx} = \alpha E - \beta P - \gamma, \quad (3.3)$$

where α , β and γ are empirical constants. For xenon, $\alpha = 0.137 \text{ V}^{-1}$, $\beta = 177 \text{ bar}^{-1} \text{ cm}^{-1}$ and $\gamma = 45.7 \text{ cm}^{-1}$. The minimum reduced field required to produce electroluminescence in xenon vapour is $E/N = 4 \text{ Td}$ (where $1 \text{ Td} = 10^{-17} \text{ V} \cdot \text{cm}^2$) which corresponds to $\approx 2.5 \text{ kV/cm}$ at a typical gas pressure of 1.6 bar. In dual-phase xenon detectors hundreds of photons can be emitted per electron over less than 1 cm of gas. At this pressure the electron drift velocity ranges between 4 to 12 mm/ μs , for fields between 5 and 10 kV/cm [169] [171]. Therefore the S2 pulse width is of the order 1 μs , much larger than that for S1 pulses. Due to longitudinal diffusion in the liquid not all electrons arrive at the liquid surface at the same time, which smears the S2 pulse shape.

3.1.4 Light yield, charge yield, discrimination and energy

Although the mechanism is essentially the same for ER and NR, the yields of ionisations and excitations are different for the two. For ER the energy deposited in the medium can be written as:

$$E_{ER} = W(n_{ex} + n_i). \quad (3.4)$$

where $W = 13.7 \pm 0.2 \text{ eV}$ [154] is the average energy required to produce either an excitation (contributing to n_{ex}) or an ionisation (n_i) in liquid xenon. The numbers of photons (n_γ) and electrons (n_e) released from the interaction site after recombination (with recombination fraction r) are:

$$\begin{aligned} n_\gamma &= n_{ex} + n_i r, \\ n_e &= n_i(1 - r). \end{aligned} \quad (3.5)$$

The anti-correlation of the recombination process means that it cannot change the total number of quanta. Therefore, the energy can also be described in terms of n_γ and n_e (corrected for electron attachment to impurities), which follows from the above equations:

$$E_{ER} = W(n_\gamma + n_e). \quad (3.6)$$

The energy deposit of NRs is similar, but there is a quenching factor \mathcal{L} , as a significant fraction of the energy is lost to heat (which is less significant for ERs):

$$E_{NR} = \frac{W}{\mathcal{L}}(n_\gamma + n_e). \quad (3.7)$$

Besides the \mathcal{L} factor NR and ER differ in their light and charge yields, which leads to the ability to discriminate between the two types of recoil. The ratio of excitations to ionisations, N_{ex}/N_i , is < 0.2 for ER [172] [173] [174] and ~ 1 for NR [175]. For ER, the theoretical measurement of N_{ex}/N_i predicts 0.06 [154]. At low fields (below approximately 1 kV/cm) $74 \pm 2\%$ of the S1 light comes from recombination [176]. NR have a larger recombination probability due to denser tracks of the recoiling particle. Therefore, as NRs have a larger N_{ex}/N_i and more recombination than ERs, the ratio of S2/S1 is smaller for NRs than for ERs. Discrimination is quantified by the fraction of ER events that lie below the median of the NR band in a plot of $\log_{10}(S2/S1)$ against S1, and over 99% rejection of ER events is possible [177] [178] [151].

Several measurements of the light and charge yields of ERs and NRs are available, and the most accurate come from the LUX experiment which has advanced our understanding of this medium like no other. Simulations with NEST (Noble Element Simulation Technique) [179] are based on global fits to the experimental data; this semi-empirical model is increasingly the tool used by the community to predict the response of the noble liquid detectors.

The measurable S1 and S2 responses are directly related to the underlying light and charge yields at a particular field:

$$S1 = n_\gamma \alpha_1, \quad (3.8)$$

$$S2 = n_e \epsilon N_{EL} \alpha_2, \quad (3.9)$$

where α_1 and α_2 are the photon detection efficiencies in the liquid and gas signals respectively, N_{EL} is the number of electroluminescence (EL) photons produced per extracted electron and ϵ is the electron extraction efficiency from the liquid surface. The values of α , N_{EL} and ϵ can be measured for a given detector at fixed operating conditions. Therefore, equation 3.6 can be used to reconstruct the energy of ER events from the S1 and S2 signal sizes (and similarly for NR):

$$E_{ER} = W \left(\frac{S1}{g_1} + \frac{S2}{g_2} \right), \quad (3.10)$$

where $g_1 (= \alpha_1)$ and $g_2 (= \alpha_2 N_{EL} \epsilon)$ are the conversion factors between the number of quanta released at the interaction site and the measured signals.

3.2 The Large Underground Xenon (LUX) Experiment

LUX is a dual-phase xenon TPC with an active volume of 250 kg liquid xenon (370 kg total), operated at the Sanford Underground Research Facility (SURF) [180] [181] in Lead, South Dakota (USA) between 2012 and 2016. LUX is located in the Davis Cavern, 1.48 km below the surface (4.30 km water equivalent). The detector internals are contained within a cryostat, 101 cm tall with a diameter of 61.6 cm. The active region is shown in Fig. 3.4 is 48 cm tall with a diameter of 47 cm. Light is detected at two PMT arrays, each consisting of 61 2-inch Hamamatsu R8778 PMTs at the top and bottom of the detector. The detector is constructed from low background materials and the cryostat is within a water tank as shielding from external radiation. Diagrams of the detector in the water tank and the internals are shown in Fig. 3.4. Table 3.1 shows some of the other key parameters of LUX. A fuller description of the experiment is given in this section and can also be found in Ref. [155].

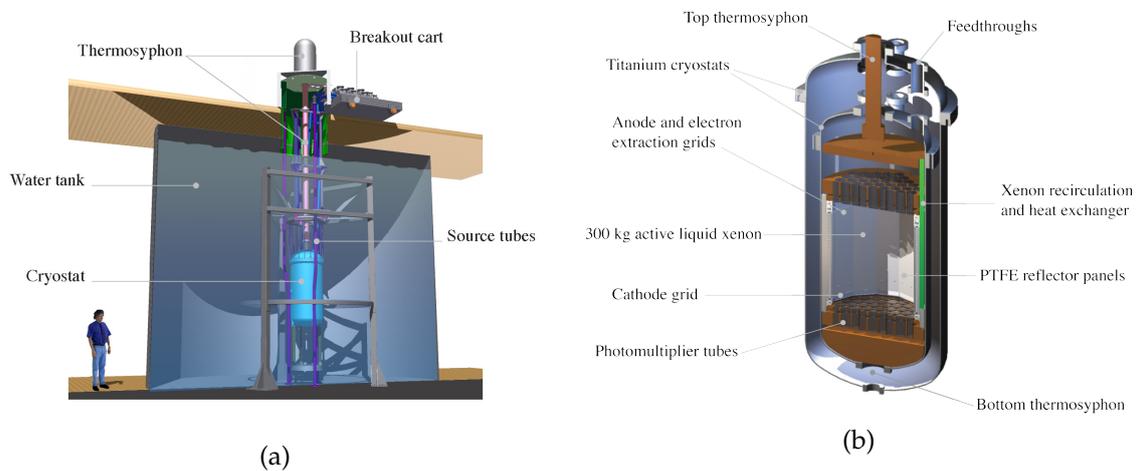


Figure 3.4: Diagram of the LUX experiment in the Davis Cavern. (a) Cryostat in the water tank. (b) Detector internals in the cryostat. Figures reproduced from Ref. [155] with permission from Elsevier.

Table 3.1: Parameters of LUX Run 3.

Parameter	Value
Gate-anode distance	10 mm
Gas gap	5.6 mm
E field surface-anode	5.6 kV/cm
E field gate-surface	2.9 kV/cm
Electron extraction efficiency	49.1%
S2 photon yield	276 photons/electron
Total drift time (cathode-gate)	328 μ s
Drift field	180 V/cm
S1 photon detection efficiency	11.7%

The first phase of the LUX program was LUX 0.1, a prototype setup to test various detector subsystems. This operated at Case Western Reserve University between

2007 and 2009 [182] [183]. LUX was deployed in the surface lab in August 2011, between September 2011 and February 2012 a commissioning run was carried out to test the xenon purification and light collection properties (Run 2) [184]. After this, there were upgrades to the detector and it was reconstructed by July 2012 and moved underground. Gas phase data were taken in 2013 before condensation was completed in February 2013. WIMP search data taking began in March 2013 and was completed by August 2013 [185], with a total live time of 85.3 days (Run 3). The final run accumulated 300 live days, completed in 2016 (Run 4).

The following sections describe the systems required to operate the detector, such as cooling, instrumentation, data acquisition, high voltage delivery, gas circulation and calibration systems.

3.2.1 TPC and cryostat

The LUX cryostat is constructed from titanium, and consists of two vessels. The titanium was selected after various materials were screened for radioactivity, and was the one with the lowest (long-lived) activity [186]. The cryostat is supported from the top, and hangs from the outer vessel. The inner vessel contains the detector internals and active xenon, which are shown in Fig. 3.5. Between the inner and outer vessels is vacuum, to insulate the inner cryostat from the room temperature water tank.

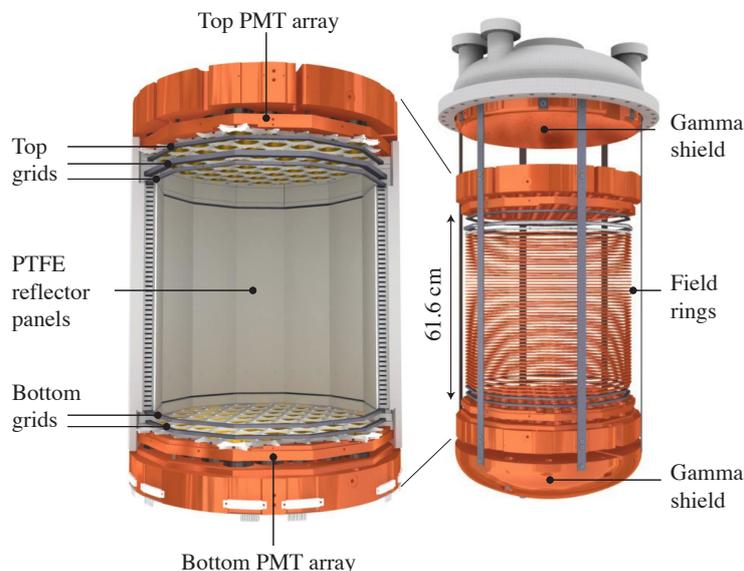


Figure 3.5: Internal structures within the LUX cryostat, reproduced from Ref. [155] with permission from Elsevier .

At the top of the inner vessel is a copper disk (mounted onto the flange), which is a shield against γ -rays. Six titanium bars hang from this to support all other internal structures. At the top and bottom are the two PMT arrays, supported in a copper structure. Below the bottom PMT is another copper γ -ray shield. Attached at the edges of the two PMT supports are panels made of Ultra High Molecular Weight (UHMW) Polyethylene. These support the inner walls of the TPC, 12 polytetrafluorethylene (PTFE) panels, and hold the copper field shaping rings within them.

The dodecagonal structure of the TPC formed by the PTFE panels is designed to minimise leakage of light to outside the detector and to maximise reflectivity. PTFE has outstanding reflectivity for VUV light. The top and bottom of the TPC are formed by the PMT arrays. PTFE trifoil reflectors are placed between the circular PMT windows to maximise the reflectivity. Maximising the reflectivity is essential to improve the photon detection efficiency, which improves the sensitivity particularly to low mass WIMPs.

The other main factor that affects the sensitivity is the background level. To reduce this the materials were selected to minimise intrinsic radioactive contamination levels, and outgassing into the xenon. The PTFE is ultra-high purity grade and the copper is C101 oxygen-free high thermal conductivity.

3.2.2 Grids and high voltage delivery

A dual-phase xenon TPC requires several regions with different electric fields. In LUX, these are created with wire grids held a constant potential, attached to the PTFE panels of the TPC as illustrated in Fig 3.5.

Starting from the bottom the first grid is 2 cm above the bottom PMT array, which shields the PMTs from the high voltage on the cathode. The cathode grid is 4 cm above this, with a potential of -10 kV. Below the cathode is the reverse field region, where the potential decreases so that the PMTs can operate near ground. If an interaction occurs here, the electrons drift downwards and do not produce an S2 signal. The next grid 49 cm above the cathode is the gate, giving a drift field of 180 V/cm in the liquid, defining the active region of the TPC (the WIMP “target”). In order to ensure uniformity of the shape of the field over the TPC height there are copper rings held in grooves in the walls (3.2 mm thick, 12.7 mm width with spacing of 1 cm), these are connected by a resistive voltage ladder between cathode and gate potentials. The next grid is the anode 1 cm above the gate with the liquid level approximately half way between them. The gate allows the field to be adjusted in this region, to give a higher extraction field (gate to liquid surface) and electroluminescence field in the gas. The last grid is the top PMT shield. The field between the anode and top is set to be below the electroluminescence threshold so that interactions here do not produce light.

All of the grids in LUX are made from a plane of single wires, except the anode which is a woven wire mesh. The pitch (wire centre to wire centre) and wire diameters are shown in Table 3.2. The maximum voltage on the cathode was limited by the onset of spurious electron emission. The cathode grid was changed at the end of Run 2 because of this, and the wire diameter was doubled. Even so, the LUX cathode operating voltage is much lower than intended in the original design due the onset of electron emission. Table 3.2 shows the final cathode parameters.

3.2.3 Photomultipliers

LUX has two PMT arrays, one at the top (in the gas) and one at the bottom of the detector (in the liquid) with 61 phototubes each. Most of the S1 light is detected in the

Table 3.2: Parameters of the LUX grids. The voltages are those used during Run 3.

Grid	Pitch (mm)	Wire diameter (μm)	Voltage (kV)
Top PMT screen	5.0	50.8	-1.0
Anode	0.25	30.5	3.5
Gate	5.0	101.6	-1.5
Cathode	5.0	206.0	-10
Bottom PMT screen	10	206.0	-2.0

bottom array due to total internal reflection at the liquid surface, whereas the S2 light is primarily detected in the top. The PMTs are Hamamatsu 2-inch R8778 [187], specifically designed for use in low background liquid xenon experiments by Hamamatsu and the XMASS collaboration. These PMTs have 12 stage multiplication, a 56 mm diameter (with 45 mm photocathode diameter), and average quantum efficiency of 33% for 178 nm xenon scintillation light.

Due to their proximity to the active volume, the PMTs are the leading contribution to the ER background. They are made with low background materials and 20 were screened for radioactivity before use at the Soudan Low Background Counting Facility (SOLO) [188]. The results per unit were 9.8 ± 0.7 mBq in U, 2.3 ± 0.5 mBq in Th, 2.2 ± 0.4 mBq in ^{60}Co and 65 ± 2 mBq ^{40}K [187].

3.2.4 Gas system, circulation and purification

The xenon in LUX is continuously circulated in order to keep the purity required to maintain a large electron attenuation length for drifting electrons in the liquid. The purifier is a commercially available SAES heated zirconium getter, which can purify the xenon of almost any non-noble species. Testing the getter showed it could achieve concentrations of less than 120 parts per trillion (ppt) g/g for O_2 and 950 ppt g/g for N_2 [189].

As the getter operates on gaseous xenon, the liquid xenon must be evaporated, circulated through the gas system, then recondensed back into the detector. Liquid is removed via a spill weir, an additional volume next to the main active volume of xenon. The top of the weir is at the liquid level, so additional liquid spills over into the weir. This also serves to keep a constant liquid level. The liquid from the weir is fed into a dual-phase heat exchanger, and evaporated. One of two pumps moves the gas through the system. Each is on a separate path in the gas system and it is possible to quickly swap from one to the other. This means that if one pump requires maintenance the detector can continue running with the other. The gas is passed through the getter, before going back into the detector via the other part of the heat exchanger where it is cooled to near its condensation point. After condensing the xenon is returned into the bottom of the detector. The continuous purification prevents build-up of impurities that are introduced while running, such as from outgassing from detector components or small leaks. The flow rate during data taking was approximately 35 standard litres per minute (slpm) and the turn-over time for the 370 kg of xenon is approximately 40 hours [178].

Xenon purity is monitored throughout running by taking samples from the detector. Xenon gas can be collected from various locations into a dedicated sampling system. The sample is analysed by mass spectrometry with a Residual Gas Analyser (RGA). A typical RGA can detect concentrations greater than 1 part-per-million, but impurities in xenon start to degrade the electron lifetime at the part-per-billion level [162]. The RGA measures the partial pressures of gases in the sample, which is proportional to the flow rate and absolute concentration. The RGA signals can be enhanced by increasing the flow rate, but this is limited by the maximum total pressure for the sample (usually $\sim 10^{-5}$ Torr). The RGA is saturated by the xenon, because the partial pressures of the impurities are small in comparison. The sensitivity is increased by using a LN cold trap to freeze the xenon (with a pressure > 1.8 mTorr) [190] [191]. This greatly reduces the xenon concentration while still allowing the impurities to pass through. With this method, the sensitivity to oxygen, nitrogen, and methane is 0.66×10^{-9} , 9.4×10^{-9} , and 0.49×10^{-9} mol/mol respectively [190]. Sensitivity to krypton was also observed at the 5×10^{-13} mol/mol (Kr/Xe) level [191]. The operation of the sampling system is highly automated; samples are taken every week so that any problems are detected early. Using these systems LUX was able to achieve and sustain the required electron lifetime for a WIMP search (several hundreds of microseconds) and diagnose any arising problems quickly.

The gas system also has several safety features. The Storage and Recovery Vessel (SRV) is beside the detector, and is able to contain all 370 kg of xenon. It is continuously kept cold when the detector is running in case of an emergency. If there were problems in evacuating the xenon to the SRV there is also a large recovery balloon. This is a last resort as the xenon would be contaminated with air. Besides these, there is an automatic pump shutoff system which can stop circulation in the event of an overpressure or other problems.

3.2.5 Cryogenics and thermosyphon cooling

The detector temperature is cooled and held at ~ 175 K using thermosyphons [192]. The LUX thermosyphons are closed pipes filled with nitrogen gas with the top immersed in a LN₂ bath (located in the “breakout” area above the water tank). This is connected to the bottom of the thermosyphon, which is attached to the detector. The thermosyphons are vertical, so as the gas rises to the top it cools and condenses. Then it trickles down to the bottom, where it evaporates, thus removing heat from the detector. The gas then rises again to repeat the cycle. The efficiency of the thermosyphons at the operating temperature was found to be very high (55 kW/K/m) with a nitrogen pressure slightly above atmosphere [192]. There are four thermosyphons in LUX. There is one at the top and bottom of the inner vessel used to cool the detector from room temperature. Two more are attached to the shielding at the sides of the detector and maintain a vertical thermal gradient during operations.

The recirculation process means that xenon needs to be continuously heated and cooled between 175 K and room temperature. The power required for this is greatly reduced by the use of heat exchangers. The first is a dual-phase heat exchanger, which takes liquid from the weir. This is pumped on by the circulation system which evaporates the liquid and creates a cooling surface. The cold gas then passes through a

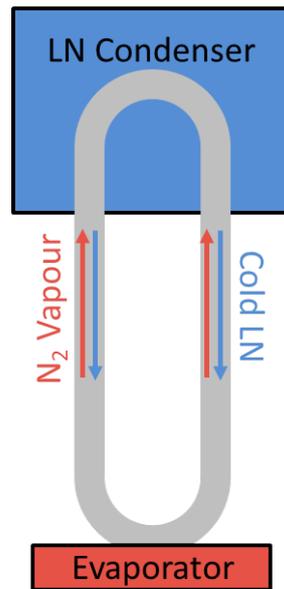


Figure 3.6: Diagram of thermosyphon operation. The condenser is a liquid nitrogen bath and the evaporator is a cold head attached to the detector. They are connected via a stainless steel pipe.

single-phase heat exchanger to take it to room temperature. After passing through the gas system the single-phase heat exchanger cools the gas to near its condensation point. Then, the dual-phase heat exchanger condenses the gas and the liquid is reintroduced to the detector. Although the thermosyphons alone could deliver enough cooling power (up to 450 W [192]), the amount of liquid nitrogen required (240 litres/day) would be prohibitively expensive and difficult to manage in an underground environment.

The main store of LN₂ is four 450-litre dewars just outside the Davis cavern, and delivery lines supply LN₂ to the detector systems such as the LN₂ reservoir for the thermosyphons at the top of the detector. It also supplies the SRV, which is continuously kept cool. Both of these are refilled automatically from the 450-litre dewars. A line is available for syphoning LN₂ into a dewar to cool the cold trap in the sampling system. The 450-litre dewars are refilled from a tanker that is transported to and from the surface via the main shaft.

3.2.6 Instrumentation and slow control

The detector is monitored for any change in conditions during operations. This is important for detecting any potential failures of components and ensuring stability of the detector so that there is no effect on the data.

Temperature is monitored in 61 separate places using thermometers made from thin film platinum resistors on a ceramic substrate. These are mounted on Cirlex[®] boards to provide strain relief when the ceramic is pressed against the surface being measured and ensure a good thermal contact. Including the wiring and electronics of the thermometers, they were found to be accurate to within 170 mK in LUX [155].

Internal detector pressure is measured with several sensors. There are differential, capacitance manometer high pressure and ion/convection gauge high-vacuum pressure sensors. The pressure is also monitored in the vacuum space of the cryostat and throughout the gas system.

Two types of liquid level sensor are used in LUX. The first are parallel-wire sensors, which measure the capacitance between the two wires. The capacitance changes depending on the length of wire submerged in the liquid. This type is used in circulation plumbing, the main chamber, weir, condenser and the liquid line to the detector. The other type is a parallel plate sensor, which measures the capacitance between two plates. Three of these are built into the gate-anode grid structure and are used to ensure the liquid is level with respect to the grids.

All of the detector parameter measurements are monitored on the LUX slow control. As well as the sensors mentioned above, the slow control also monitors: the grid high-voltage, flow controllers on the gas system, PMT voltages and the LN₂ systems. The slow control is accessed via a web browser to view or retrieve data. The data are permanently stored in an SQL database to enable access to past data. The slow control also serves as an alarm system. Alarms can be set on any of the sensors. If a parameter goes above/below the maximum/minimum value, alarms are triggered both on-site and sent via email or text message. Many detector operations are also performed via the slow control.

3.2.7 DAQ and electronics

The LUX DAQ system [193] and readout is designed to provide a high signal-to-noise ratio for single photoelectrons while maintaining a large dynamic range for large signals. The PMTs provide a gain of 3.3×10^6 , which generates a pulse area of 13.3 mVns at the PMT base with 50 Ω termination. This is further multiplied by a factor 5 at the pre-amplifier. The post-amplifier splits the signal into three outputs, each with a different gain. A diagram of the DAQ system is shown in Fig. 3.7.

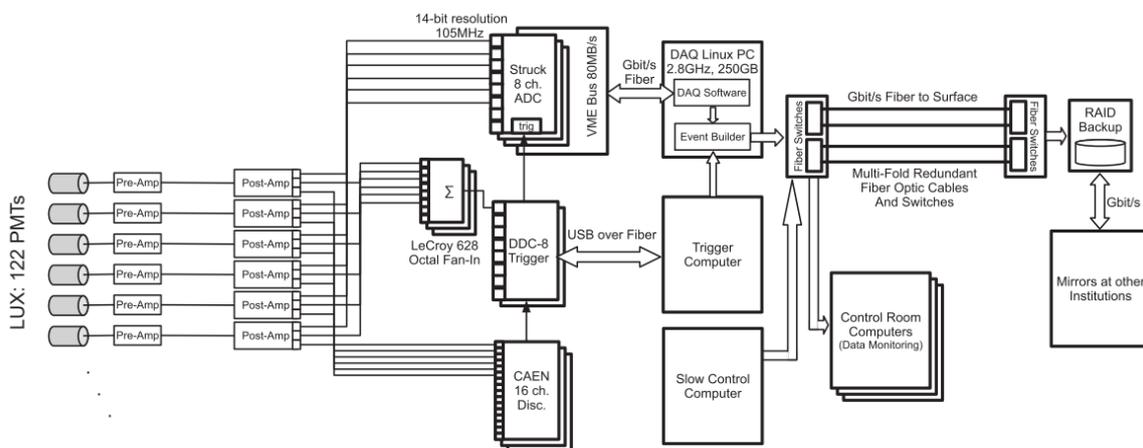


Figure 3.7: Data flow through the DAQ of the LUX experiment. Signals from the PMTs are amplified and split into three outputs. Figure reproduced from Ref. [193] with permission from Elsevier .

The first output from the post-amplifier is to 16 Struck 8-channel analog-to-digital (ADC) modules, which are the readout for the data. The gain for these is a further factor of 1.5 and each PMT is read out by one channel. The size of a single phe at the Struck output is 100 mVns, with a height of 4 mV. The Struck boards operate in Pulse-Only Digitisation (POD) mode, which acquires only when there is an excursion above threshold. This means that the long periods of quiet time are not recorded, and vastly reduces the storage requirements. When a pulse is detected a set amount of time around the pulse is also recorded. The Struck channels operate in pairs: even if only one channel is above threshold, both are recorded. The pulse detect threshold is set to 13 ADC counts to give a 95% single phe detection efficiency. A 5σ fluctuation of background noise is 6.5 ADC counts on the struck.

The second output from the post-amplifier is for the trigger system [194]. This is an 8-channel digital-to-digital conversion (DDC), with a gain of $\times 2.8$. The PMTs are summed into 16 groups (of 7 to 8 PMTs each) for the trigger. The DDC-8 firmware has the ability to identify S1 and S2 pulses, measure the size of the S1, and record the hit pattern of the S2. It can trigger on either S1s, S2s, or S1 + S2. For either the S1 or the S2 mode, whenever the respective type of pulse is found a trigger signal is sent to the DAQ. After each trigger there is a hold-off time (between 4 μ s and 65 ms) before the next trigger can be sent to the DAQ. This prevents overlap of events and triggering on the tails after large interactions. In the combined S1 + S2 mode, the drift time can also be used to select events at certain depths. Once an S1 is found, it looks for an S2 within a certain time window afterwards. Trigger information is recorded and combined with the data stream from the Struck modules.

The third output from the post-amplifier is the discriminator with a gain of $\times 18$. This is required because the Struck modules have two memory banks, and there is a dead time ($< 290 \mu$ s) when one is full and it switches to the other. This time is insignificant unless there is significant noise in one of the Struck channels, which would cause the memory to fill quickly and force switches. This type of noise would not be coincident with other pulses. The discriminator implements a Valid Pulse Trigger Gate (VPTG), which is only turned on for pulses that are detected in more than one channel. Only when the gate is on are pulses recorded.

There are two modes of operation of the DAQ. The first is “multi-event mode”, where all PODs are recorded whether or not there was an associated trigger. The trigger and event building are done in software and so changes can be made offline for data previously taken. In “single-event mode” only pulses within events are recorded to the DAQ computer, which is using the full trigger capability to discard unwanted events. For the Run 3 WIMP search data the trigger operated in S2 mode, with a 1–4 μ s hold-off and the DAQ was set to multi-event mode, effectively recording every detected photon to disk.

3.2.8 Calibration systems

Calibrations are essential to understand the detector response to both ER and NR. They determine the positions of the ER and NR bands in the S2/S1 vs. S1 plane for the discrimination and to apply position corrections. This section describes the deployment

of the internal and external sources, and the use of the calibration data is described in the following chapter.

External calibration sources are deployed in one of six square tubes at the sides of the cryostat, within the water tank. The source is contained within a cuboid holder that is lowered down the tubes with a winch. There is a collimator in front of the source and tungsten backing behind it. The tungsten is to reduce photons backscattering into the detector. The external sources include several gamma-ray sources from hundreds of keV to MeV energies (e.g. ^{137}Cs , ^{208}Tl) and neutron sources (^{252}Cf , AmBe).

A different external source is the Deuterium-Deuterium (DD) neutron generator, which is a source of 2.45 MeV neutrons and produces nuclear recoils down to keV energies in the xenon. The machine is placed outside of the water tank and a plastic tube containing air is placed within the water tank to allow a collimated, mono-energetic neutron beam to reach the detector.

The gas system also allows for the deployment of internal calibration sources, which LUX utilised for the first time — these novel sources are partly responsible for the success of the experiment. The first source is ^{83m}Kr [195] [196]. This is used for normalising the S1 and S2 responses in position and time and measuring the electron lifetime. ^{83m}Kr is produced from ^{83}Rb infused into zeolite located in a section of the gas system and sealed by valves. ^{83}Rb decays via electron capture with a half-life of 86.2 days into an excited state of ^{83}Kr , which rapidly decays into ^{83m}Kr . This noble gas radioisotope is the calibration source, which produces decays in two stages producing a 32.1 keV deposit (half-life of 1.83 hours), followed by one of 9.4 keV (half-life of 154 ns), which are mostly a combination of internal conversion and Auger electrons with a small contribution from gamma-rays and X-rays. The Kr emanates out of the zeolite and fills the volume of the pipe section. To inject the source into the xenon, the gas in the section containing the Rb source is first pumped out. Then, some of the xenon from the main circulation path is allowed to flow through this section for a set time (a few minutes). Pumping out is required to ensure an appropriate rate of ^{83m}Kr events (~ 10 Hz). The source enters via the normal recirculation path through the getter (as a noble gas Kr is not removed) and enters the detector at the bottom, dispersing throughout the TPC in a few minutes. The short half-life of ^{83m}Kr means that the event rate becomes negligible after a few hours, so this calibration can be repeated routinely throughout running (every few days).

The other internal source is tritium [178], which is primarily used for calibrating the low energy ER response with non-resonant interactions that look like the dominant background. However, tritium has a long half-life of 12.3 years and it attaches to surfaces. To remove this from the detector it is used in the form of tritiated methane (CH_3T). This is commercially available and is removed via the standard chemical purifier. The methane is mixed with purified xenon to make the source. To get only a fraction of the source from the bottle, the gas is allowed to expand into one of several available expansion volumes, via an additional purifier. The bottle is then sealed and gas from the main circulation path, but after the getter, is allowed to flow through the volume to inject the tritiated methane into the detector. The methane concentration in the detector decreases exponentially with a time constant of 5.90 ± 0.07 hours by action of the purification system.

3.2.9 Water tank and veto

The cryostat is immersed in a water tank to shield backgrounds outside the detector: neutrons from radioactivity in the cavern walls and neutrons from muon interactions in the rock. Water is more practical than lead or polyethylene and can be purified. With a standard commercial purifier water also has lower γ -ray backgrounds. The cylindrical tank is 6.1 m tall with a diameter of 7.6 m. Nitrogen gas is used to purge the space above the water to prevent the build-up of radon gas.

The water tank is instrumented with 20 10-inch PMTs to detect Cherenkov radiation of incoming muons. A nuclear recoil in the detector that is coincident with a muon in the water tank can be vetoed. There are four sets of PMTs aligned vertically, with the last four at the bottom of each column. The tank is lined with Tyvek reflectors to optimise the light collection. However, the water veto was not operated during Run 3.

3.3 The LUX-ZEPLIN (LZ) Experiment

LUX-ZEPLIN (LZ) [197] [130] is a next-generation dual-phase xenon TPC for detecting WIMP dark matter, but scaled up from a 250 kg active mass to 7 tonnes. The general principle of operation is same as that for LUX, but with several improvements. In particular, the instrumented xenon “skin” between the outer wall of the TPC and the inner vessel of the cryostat, and a scintillator veto outside of the cryostat (discussed in section 3.3.1) allow the rejection of events with energy deposits both in the main volume and the veto. Construction started in 2015; LZ will be installed in the Davis cavern at SURF in 2018/19. The expected sensitivity is 2×10^{-48} cm² at a WIMP mass of ~ 50 GeV for a 1000 day run [130]. This is approximately a factor 100 more sensitive than the LUX Run 3 result. Figure 3.8 shows the LZ experiment set up.

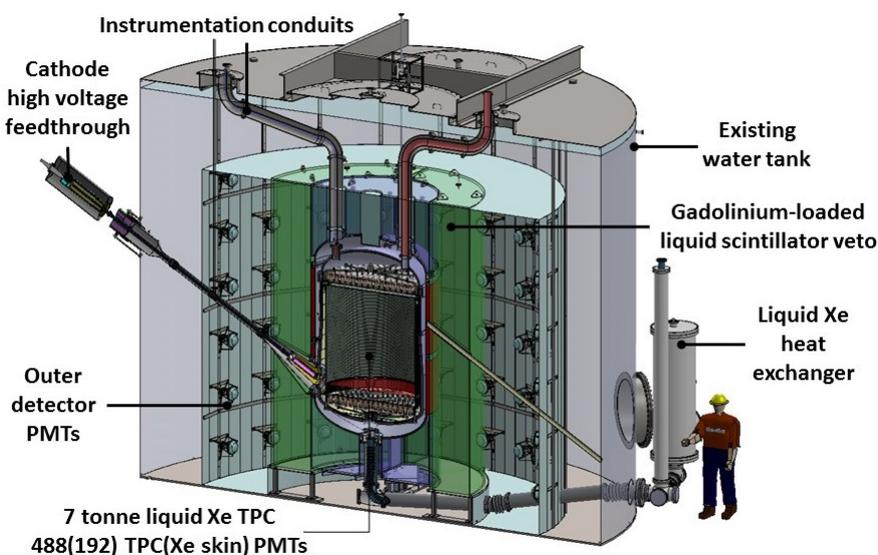


Figure 3.8: The LZ experiment, reproduced from Ref. [130].

The inner cryostat has a diameter of 1.58–1.66 m and a height of 2.52 m, with a cathode-gate distance of 1.456 m. More details on the grids can be found in Chapter 7. The PMT arrays at the top and bottom consist of 488 Hamamatsu R11410-20 3-inch PMT tubes, specially designed with the manufacturer to be low-background. The instrumented layers outside of the cryostat in Fig. 3.8 form the veto system, discussed in the following section. This is all contained within the same water tank that was used for LUX. The cathode high voltage feedthrough from air into xenon is placed outside of the detector at room temperature and delivered via a vacuum insulated xenon filled umbilical. Recirculation of the gas for purification goes through the conduit on the right of Fig. 3.8, through the external heat exchanger. Some of the key parameters of LZ are shown in Table 3.3, and more details can be found in the design reports [197] [130].

Table 3.3: Parameters of the LZ experiment.

Parameter	Value
Gate-anode distance	13 mm
Gas gap	8 mm
E field surface-anode	10.2 kV/cm
E field gate-surface	5.2 kV/cm
Electron extraction efficiency	97.6%
S2 photon yield	820 photons/electron
S2 width	1.2 μ s
Total drift time (cathode-gate)	802 μ s
Drift field with 50/100 kV cathode	310 / 720 V/cm
S1 photon detection efficiency	7.5%

The larger xenon target is motivated by the increased sensitivity along with providing improved self-shielding for external backgrounds. The fiducial volume in LZ will be $\sim 80\%$ of the active xenon, compared to $\sim 50\%$ in LUX. For γ -rays and neutrons the mean interaction length is of the order 10 cm [198] [199], so they are very likely to multiple scatter before reaching into the core of the detector.

A goal of the experiment is for the most significant background to be from ERs caused by elastic scattering of solar pp neutrinos, which can be discriminated by their S2/S1 ratio to 2 events in the nominal exposure. An exception is from solar ^8B neutrinos, which are expected to contribute 7 ± 3 NR events, but are at very low energies. However, the dominant background could well be due to radon, which is present in the bulk of the xenon and in detector materials. Radon gives ER backgrounds in the bulk from decays of the ^{214}Pb daughter, and plate-out on the walls can give events that are mis-reconstructed towards the centre. The goal can be achieved by controlling the exposure of the components to air and storing them in radon-clean environments. The largest other sources of material radioactivity are from the PMTs and the titanium cryostat, but are expected to be sub-dominant. Intrinsic backgrounds such as ^{85}Kr and ^{37}Ar can be mitigated by purification of the xenon by chromatographic separation [200]. As the solar-pp neutrinos and radon backgrounds are ER, the discrimination must be high to reject as many of these as possible. Even if the radon backgrounds only meet the conservative requirement, not the goal, only a total of 6.4 and 0.5 ER and NR events respectively are expected in 1000 days (assuming 95% discrimination, 50% NR acceptance).

The larger maximum drift length in LZ (corresponding to $\sim 800 \mu\text{s}$ electron lifetime) sets more stringent requirements on the electron lifetime. This requires continuous purification using a hot zirconium getter, as used in LUX, to avoid build-up of electronegative impurities from outgassing. A higher capacity getter is used in order to have a faster flow rate through the gas system of 500 slpm. In EXO an electron lifetime of over 5 ms was attained [201], which should also be achievable in LZ.

Various calibration sources will be used in LZ. This includes external neutron and γ -ray sources in one of three source tubes in the vacuum space of the cryostat and a DD-neutron generator. However, the improved self shielding means that external calibration sources are less effective, as the particles can no longer reach the centre of the detector. Therefore, internal sources are particularly important for LZ. $^{83\text{m}}\text{Kr}$ can be used to monitor spatial effects and can be deployed routinely during the run, due to its 1.8 hour half-life. However, the mixing in the LZ xenon will be slow, and the krypton may not become uniformly distributed throughout the volume before it decays away. Therefore, an additional source $^{131\text{m}}\text{Xe}$ will be used. This provides 164 keV gamma-ray with an 11 day half-life. However, the longer half-life means this can only be deployed occasionally. The internal tritium source developed for LUX will also be deployed in LZ to calibrate the ER band.

3.3.1 The LZ anti-coincidence systems

The LZ veto strategy includes the LXe Skin and the liquid scintillator Outer Detector. It serves two purposes: to reject events which leave a deposit in both the main xenon volume and the surrounding material, and to characterise backgrounds. Besides neutrino sources, the major backgrounds are neutrons or γ -rays from the detector materials. The skin region can also tag events that occur outside the detector, but leak light into the TPC and fake S1s.

The liquid xenon skin fills the volume between the wall of the TPC and the cryostat. This consists of 4–8 cm of xenon at the sides of the detector and a thicker region below the bottom PMT array; the primary function of this xenon is dielectric insulation of the HV components of the TPC. Clearly, there is a need to instrument this region with light readout as any scintillation leaking into the TPC can add to S1 light from a fiducial interaction and therefore create a difficult background topology. The side region is instrumented with 90 upward-facing and 90 downward-facing Hamamatsu R8520-406 1-inch PMTs, and a further 12 R11410-22 3 inch PMTs in the dome. Some 2 tonnes of liquid xenon make up the skin.

The Gd-doped liquid scintillator in the Outer Detector is held in 10 clear acrylic vessels surrounding the cryostat. The tanks are viewed by 120 Hamamatsu R5912 8-inch PMTs in the water space, 84 cm away (see Fig. 3.8). The scintillator material, linear-alkylbenzene (LAB) is chosen to give a good detection efficiency for γ -rays or neutrons and was previously used in reactor neutrino experiments [202] [203]. It is doped with Gd due to its huge neutron capture cross section, even a small amount of Gd (0.1%) reduces the mean capture time from 200 to 30 μs . Neutrons captured on Gd produce a cascade of 3-4 γ -rays with total energy of 8 MeV. Alternatively, a neutron capture on hydrogen in the scintillator produces a 2.2 MeV γ -ray.

The veto system is primarily to reduce the neutron and γ -ray backgrounds to be less than 10% of the solar pp-neutrino background over a larger fraction of the active volume. The xenon skin is essential for vetoing γ -rays, as they do not always penetrate to the scintillator. The energy threshold is set to be much less than the typical γ -ray energy (\sim MeV). With a 100 keV threshold a 95% veto efficiency of γ -rays can be achieved in the skin. A 100 keV threshold is also required in the liquid scintillator, which gives a 97% veto efficiency for neutrons (typically 0.5 to 5 MeV) [130]. The ability to veto a large fraction of these backgrounds significantly increases the size of the fiducial volume to 5.6 tonnes, whereas without the veto system it would be only 3.3 tonnes [130].

This hermetic Outer Detector also provides a reliable way of measuring the neutron background, by measuring neutrons that scatter in both the liquid xenon and veto system. This would provide essential evidence to support a claim for a WIMP discovery.

Chapter 4

LUX Run 3

Run 3 was the first WIMP search of the LUX experiment, with a total of 85.3 live days of data taken between April and August 2013. It followed a run in the surface laboratory at SURF, within a smaller water tank, which confirmed the general operation of the detector, while the Davis laboratory was being refitted. Run 3 was both an engineering run to study the detector performance and also a (non-blind) science run.

This chapter summarises the work required to operate LUX and to produce a WIMP search result, highlighting my own contributions: a preliminary analysis of the tritium calibration data, a direct search in data to constrain the ^{85}Kr background in Run 3, and the description of a real-time detector monitoring tool. The result from the ^{85}Kr analysis was published as part of Ref. [204] and the tritium work was one of two independent analyses that confirmed the ER detection efficiency for Run 3. The WIMP search result from Run 3 concludes this.

4.1 Operations and Data Processing

4.1.1 Detector operations

The detector was conveyed underground in late 2012 and recommissioned alongside all ancillary systems, and the water tank was filled. The liquid xenon target was condensed in February 2013; a period of optimisation of the operating parameters followed (cryogenic conditions, recirculation flow, electrode voltages). Thereafter calibrations were undertaken, with WIMP search data acquisition taking place whenever possible.

There are several routine tasks to be performed while the detector is running. All LUX personnel participate in on-site shifts to operate the detector coordinated by a Detector Operations Manager and a Shift Manager on site.

^{83m}Kr calibrations occurred approximately once or twice a week, resulting only in a short down-time from taking WIMP search data. This initially required manual operation of the valves to pass krypton into the xenon circulation path, but was later automated and operated from the slow control. External calibrations using the source tubes were less frequent. Tritium calibrations require days to be removed from the detector, and so were performed before and after taking the WIMP search data.

Another regular task is sampling the xenon to monitor purity throughout data taking. During Run 3 this was done approximately 3 times per week, from one of several locations in the detector. This was mostly operated via the slow control, but required the LN₂ cold trap to be filled manually.

The LN₂ supplies must also be monitored. Although the SRV and thermosyphon refills take place automatically, they cannot be done just after LN₂ has been taken for the sampling system. Therefore, they can also be started manually to make sure they do not require LN₂ at the same time as the sampling. The 450-litre dewars are refilled from a tanker by a LUX shifter and the tanker is transported underground two or three times a week by SURF personnel.

On-site personnel are responsible for monitoring the detector and responding to alarms. The slow control provides information from all of the sensors. Many are alarmed, but it is possible to spot potential problems before they occur by monitoring the data. It is important that the shifters can respond to any potential problems quickly. Slow control alarms are also sent to two dedicated mobile phones carried by LUX shifters, so that someone is alerted even if it occurs overnight, and underground access outside of normal hours is available in case of an emergency.

The quality of the data taken every day was monitored by on-site and off-site shifters, using both online and offline analysis tools. This is important to allow any potential problems to be quickly identified and resolved. The Detector Physics Monitor, described in section 4.2, is key software that enables this.

4.1.2 Data processing

Raw data from the detector are stored as binary “.dat” files. At this stage, the data consist of the signal in each of the 122 PMTs, digitised at 100 MS/s with any excursion above threshold recorded (pulse only digitisation, or POD). The binary data are processed by an “event builder”, which finds the pulse that triggered and saves the data around this time into an event. Selectable pre- and post-trigger times determine how much time to record before and after the trigger respectively. For WIMP search data this is set to be over twice the maximum drift time of the TPC to ensure that both the S1 and S2 from an interaction are in the same event. This produces binary “.evt” files.

The next stage is to convert the .evt files into reduced quantities (RQs). These are a set of parameters describing parameters of the event and pulses within the event. At the evt stage, the data still contain the waveform of the PMT signals. In the RQs, the data are numbers that quantify the event and any pulses within it. For example, the pulse areas are stored, not the whole waveform. It is the RQs that are typically used in data analysis.

The processing into RQs has several steps. The LUX processing framework for evts is modular, and each task is performed by one module. This allows flexibility, as individual modules can be modified without affecting the whole processing chain, and alternatives can be easily compared. Modules can be written in C++, python, or MATLAB. The framework to run the modules is a python script. This reads an xml

file which specifies which modules to use and the order to run them in. The xml file also specifies the inputs required by each module. Next, the main modules used for processing data are described.

In the evt file each event contains the individual PMT signal as a function of time. The first module is the POD calibration in terms of physical units of photoelectrons: the raw POD signals are converted from the raw signal in mV to phe per sample. This requires the gain of each PMT in mVns. The next module is the POD summer which simply adds the contribution from individual PMTs into a sumPOD. This can be seen in Fig. 4.1, the top panel shows the individual PODs and the bottom the sumPODs. After this is the pulse finder module, which identifies the start and end times of pulses. This is followed by a further pulse timing module. Here other timing parameters are calculated, for example the point at which 1% of the total pulse area has been reached or the time of the maximum pulse height.

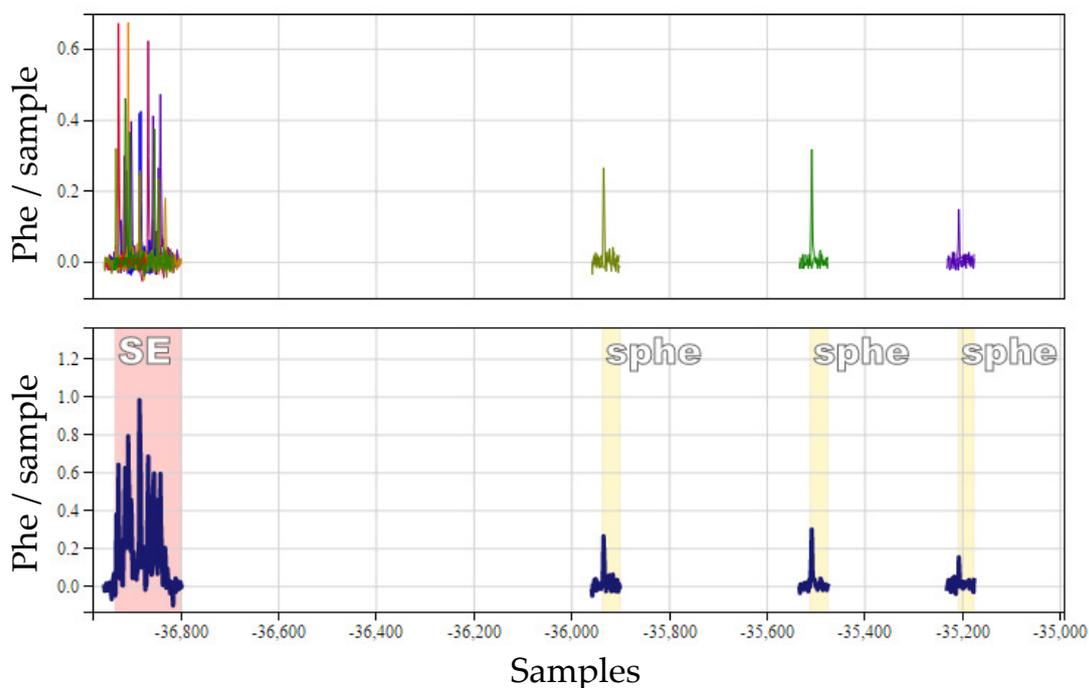


Figure 4.1: Example single electron (SE) and single phe (sphe) pulses. The top panel shows the PODs, where each colour represents a different PMT channel. The bottom panel is the summed POD (sumPOD), which also shows the classification of each pulse. One sample is 10 ns.

The next module calculates pulse quantities. This returns a large number of parameters that describe each pulse. Some examples of the calculated quantities are the pulse area, prompt fraction (area within the first 12 samples of the pulse divided by the total pulse area), and top-bottom asymmetry ($[\text{area in top array} - \text{area in bottom array}]$ divided by total area). These are required for the next module, the pulse classifier. This determines whether a pulse is an S1, S2, single photoelectron or single electron. There is another category for pulses that are not classified as one of the other types. Example pulse classifications for single electrons and single phes can be seen in Fig. 4.1. S1s and single phes are short (in time), have large prompt fractions and tend to have negative top-bottom asymmetry. S1s also have a PMT coincidence criterion, that light is

detected in at least two PMTs within 100 ns. S2s and single electrons are longer with a smaller prompt fraction. The difference between S2 and single electrons is defined by a simple pulse area cut. With the pulses classified, the next module can classify the whole event. A single scatter event requires one S1 followed by one S2 (but ignores additional S1s after the S2 or equivalently additional S2s before the S1). It also classifies multiple scatter events with more than one S2.

The final processing stages are the position reconstruction and corrections. The position reconstruction calculates the (x, y) position from each S2 pulse using the Mercury algorithm [205]. Corrections are made to the pulse areas based on the position. This is to ensure a uniform response across the detector. S1 and S2 sizes vary across the x - y plane of the detector, and as a function of z . The detector response is evaluated using a monoenergetic source (^{83m}Kr) uniformly distributed throughout the detector. Pulse areas are corrected to the values they would have if the interaction were at the centre of the detector. This is discussed further in section 4.3.1. S1 areas are position dependent due to differences in light collection efficiency. This is also the source of the S2 (x, y) position dependence. The S2 z position dependence comes from loss of electrons due to impurities. After these modules, the RQs now contain the classification of each pulse and the position-corrected pulse areas.

4.2 Real Time Data Monitoring

Real time data monitoring is important for ensuring that potential problems in the data can be identified and resolved quickly. It also allows monitoring of routine operations, such as deployment of calibration sources. Monitoring for LUX is done with the Detector Physics Monitor (DPM). This analyses raw data as soon as they are taken and produces plots on a web page within approximately 2 minutes that can be accessed by the collaboration; these are also typically displayed underground. This is complementary to the slow control, which monitors parameters such as pressure and temperature but not the data from the PMTs. The DPM runs on the local on-site cluster, on the same machine as the event builder.

The DPM reads the binary .dat files as they are produced, rather than the .evt files which take at least 10-15 minutes to be processed. The DPM is written in C++ with ROOT and based on the LUX event builder, which reads the binary waveform, builds the sumPODs and finds the trigger pulse and event window. However, instead of creating the .evt files the DPM performs an analysis on the sumPODs and events. Another advantage of using .dat files is that all PODs can be analysed, not only those within an event window. The DPM also identifies events, and so it can produce plots for all PODs and on the event level. As the DPM must analyse the sumPODs quickly, it cannot be as advanced as the actual data processing. The DPM reconstructs a few basic quantities of each sumPOD, and has a quick pulse classification. After classifying the PODs golden events (one S1 and one S2) can be found, and their properties reconstructed.

The DPM code consists of an infinite loop that continuously looks for incoming data and updates the plots. The structure of the code is summarised in Fig. 4.2. It starts by initialising all of the required variables and histograms, which is done only

once. The first step within the infinite loop is to find the latest dataset available on the DAQ computer. Then it searches the folder to find the list of files from that dataset. If there are new files available, it reads the last N files in the list (one-by-one) that have not already been analysed. N can be adjusted, and is set so that the DPM cannot fall behind the incoming data (it effectively sets the refresh rate of the plots).

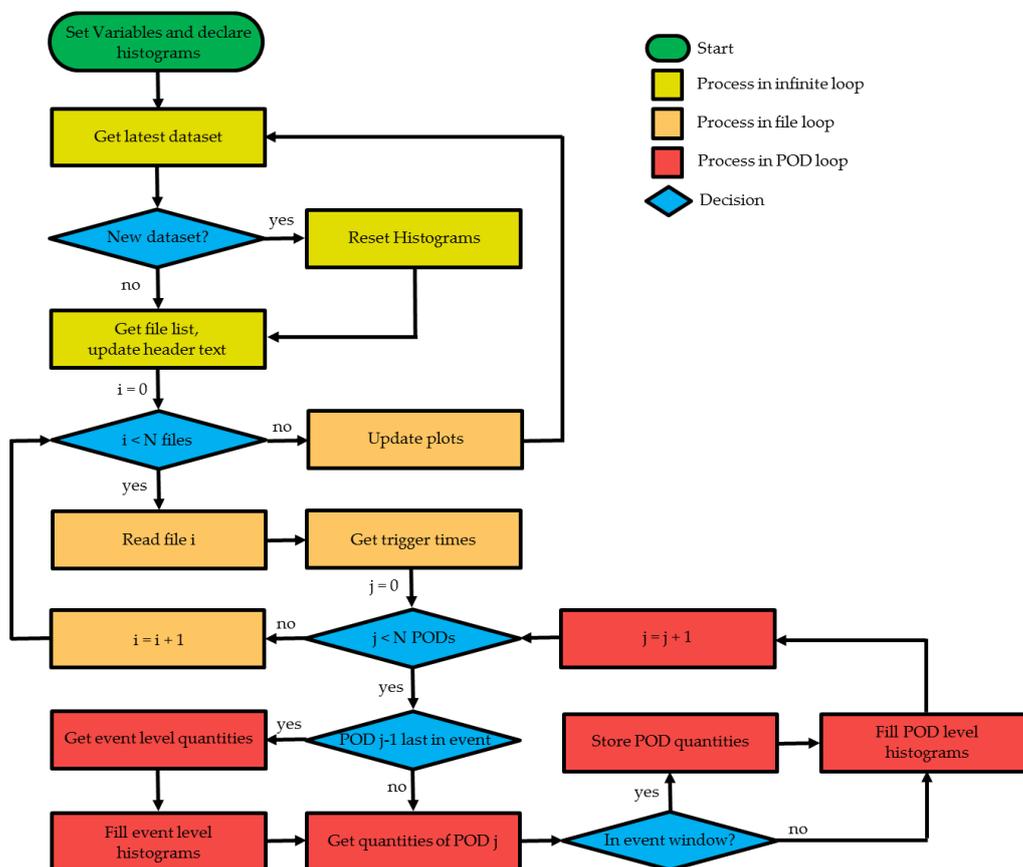


Figure 4.2: Flow chart of the DPM code.

For each file, it reads the trigger times of the events, before looping through all of the sumPODs. The first step is to check whether the previous POD was the last pulse within a trigger window. This is done here because it requires the pulse quantities of the previous POD to already be constructed, and the time of the current POD to know that the event ended. If it was not the last POD in an event, it carries on to reconstruct the POD quantities (area, prompt fraction, pulse classification). For PODs that lie within an event, the POD quantities are stored for use later on. Then any POD level histograms (either all PODs, all S2 PODs etc.) are filled and it returns to the start of the POD loop.

When it reaches the end of an event, it goes back and accesses the previously stored pulse quantities. It determines if the event had one S1 and one S2 and, if so, it calculates event level quantities and fills the histograms for the event.

Once all N files have been analysed, the last step in the infinite loop is to redraw the plots. The new plots are then saved, overwriting the previous ones. At this point if the latest .dat files are still from the same dataset, it carries on as before. If there is a new

dataset it saves the final plots from the previous dataset to an archive, before resetting the histograms. Then it starts the analysis again on the new dataset. It automatically goes to the new dataset even if there are still files that were not analysed from the old one, which is to ensure that it cannot fall behind the incoming data.

The saved images are served to a web page via a simple bash script. This periodically looks in the folder where the DPM code saves the plots and copies them to the web page. It also occasionally looks in the folder where the plots from the completed datasets are stored and transfers those to the archive. Both the current plots and archive can be accessed via the same web interface.

The following sections describe the reconstruction of the pulse and event quantities and the plots produced by the DPM.

4.2.1 Event and pulse selection

Every POD has several quantities reconstructed: pulse area, prompt fraction and length. A quick pulse classification is applied, based on the full version of the pulse classifier used in the data processing, but with the minimum required cuts. Out of the different parameters used in the full pulse classifier, the prompt fraction and pulse area were found to provide an acceptable separation between the different types. To set the cuts in the DPM classifier, the prompt fraction vs pulse area is plotted from some WIMP search data with the full processing chain.

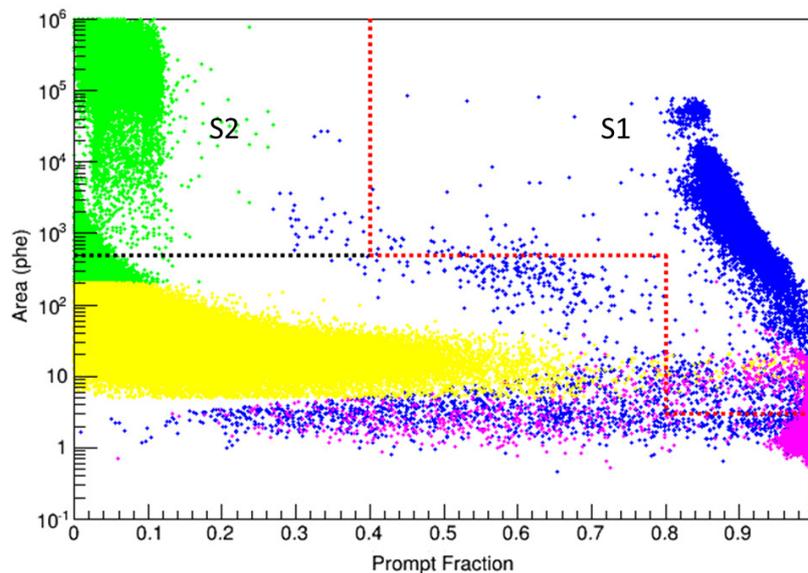


Figure 4.3: Pulse classification from the full data processing chain used to create cuts for the DPM classifier. Blue points are classified as S1 from the data processing, green are S2, pink are single phe and yellow single electrons. The region labelled S1 and determined by the red line defines S1s in the DPM. The top left region is for S2s in the DPM. Pulses in the other region are either single electrons or single phe, depending on the pulse area and PMT coincidence.

The points on Fig. 4.3 are from the full data processing, and the colours are for the different pulse classes. The selection cuts to reproduce this classification are:

- S1
 - $3 < \text{area} < 500$ phe, coincidence in ≥ 3 PMTs, prompt fraction > 0.8 .
 - Or, $\text{area} > 500$ phe, prompt fraction > 0.4 .
- S2
 - $\text{Area} > 500$ phe, prompt fraction < 0.4 .
- Single phe
 - $\text{Area} < 3$ phe.
- Single Electron
 - $3 < \text{Area} < 500$ phe, coincidence in ≥ 3 PMTs, prompt fraction < 0.8 .

These cuts are drawn on Fig. 4.3. The S2 cut has been set higher than in the actual data processing due to ‘electron trains’, a tail of delayed electron emission following large S2 signals. The electron trains were causing many small S2s to be found in the DPM after large events. Below 500 phe, the separation between S1s and SEs is based only on the prompt fraction. If a pulse either has an area below 3 phe, or does not pass the coincidence requirement for SEs or S1s, it is classed as a single phe. For pulses above 500 phe, only the prompt fraction is used to determine if it is an S1 or S2. The DPM classification reproduces the full processing classification reasonably well, but will misclassify some of the small S2s or S1s.

4.2.2 Output and uses

The DPM produces a total of 20 plots; the main ones are summarised here. The first two are histograms of the sumPOD areas, split between one for pulses shorter than $2 \mu\text{s}$ and the other for those longer than $2 \mu\text{s}$. There is a similar plot of the area against length. These plots include all sumPODs in the data, and are therefore useful for spotting any extra populations of pulses. Fig. 4.4 shows examples from WIMP search data and $^{83\text{m}}\text{Kr}$ calibration data. The population at the bottom left is single photoelectrons, and the one just above (1,1) is single electrons. The two extra populations from krypton S1 and S2 pulses can be seen by comparing the two plots.

The total POD rate is also measured over time. There can be increases in the POD rate for various reasons, such as grid-related or other spurious light sources within the chamber. If there is a spike in the POD rate then the data should be excluded. The DPM is able to see POD rate spikes, but there is also another tool that specifically tracks this and provides output of the rate from particular PMTs. The DPM plots the total number of PODs divided by the total livetime of the analysed data. Therefore, any livetime losses in the DPM are accounted for. Fig. 4.5 shows examples of the POD

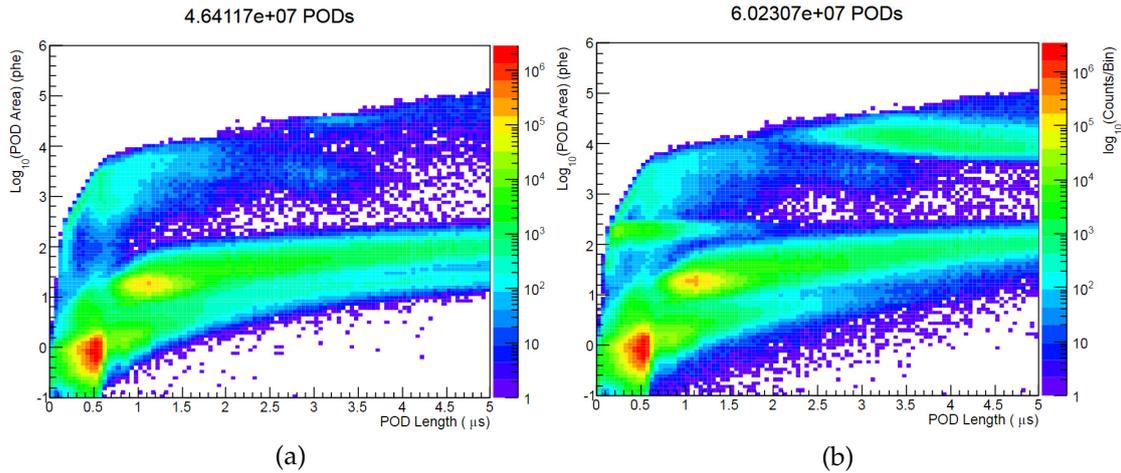


Figure 4.4: DPM plot of the pulse area against pulse length. (a) During WIMP search data. (b) during ^{83m}Kr calibration.

rate during a krypton injection and spurious POD rate spikes. The average area of PODs is also tracked over time, which will change if any extra population of pulses appears. The livetime is also plotted, by adding up the total time analysed. For 100% livetime the total time analysed should be equal to the bin width in time.

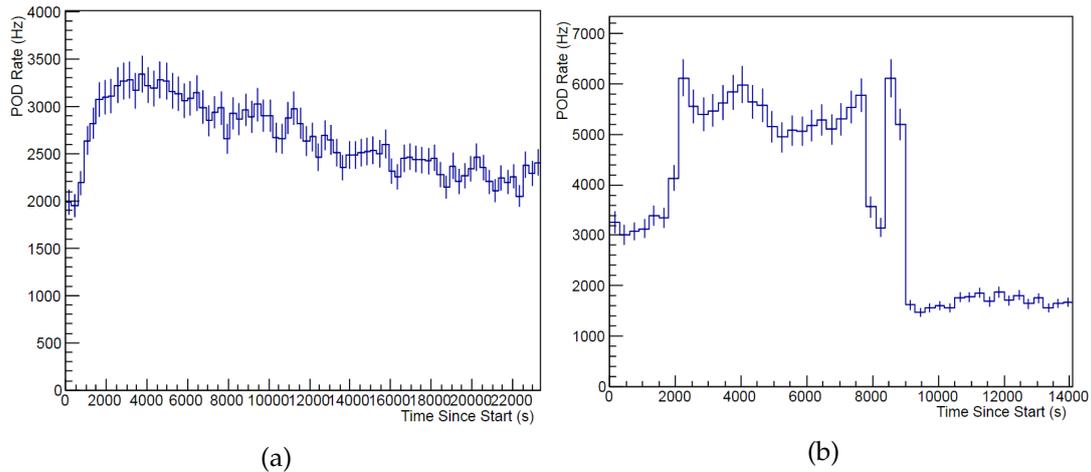


Figure 4.5: POD rate from the DPM (a) during ^{83m}Kr calibration and (b) during spurious light emission in the detector.

As well as looking at individual PODs, golden events can be investigated. The plot of the S2 against S1 size is useful to identify different populations of events and is one of the first to show events from calibration data. It is useful for ensuring that a calibration was properly deployed. Fig. 4.6 is shortly after a ^{83m}Kr injection took place, which shows as the population around $\log_{10}(S1) = 2.3$.

The next event level quantity is $\log_{10}(S2/S1)$ against drift time. This plot shows the effects of the purity on the data. An approximate measurement of the electron lifetime is made by fitting a straight line to $\log_{10}(S2/S1)$, but it is not completely accurate as only a basic S1 z-position correction is made. The DPM plot is still useful for seeing changes in the purity. Fig. 4.7 shows a dataset from before and after a loss of purity due

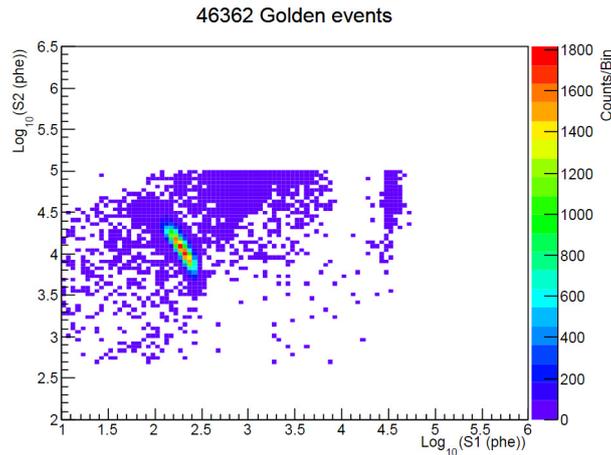


Figure 4.6: DPM plot of S2 against S1 pulse area for golden events. The hot-spot reveals the ^{83m}Kr calibration source.

to an operational incident (circulation stoppage) and there is a drop in the measured electron lifetime. Therefore, this allows any loss in purity to be seen very quickly. It was also used to measure the purity over time, as shown in Fig. 4.8. Various drops in purity can be seen here, most of which were due to stops in circulation. The large drop at the start was due to bypassing the getter during a tritium injection. A better purity measurement is from KrypCal, the analysis used to measure the electron lifetime and apply corrections. The DPM agrees reasonably well with the analysis apart from when the purity is very high.

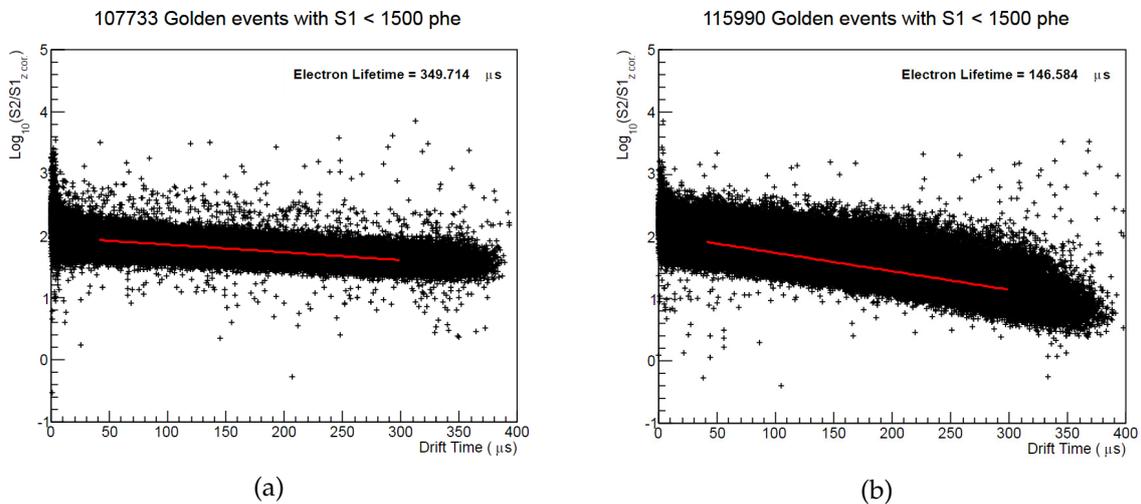


Figure 4.7: Scatter plots of the S2/S1 ratio against drift time from the DPM. The S1 size has a basic correction for z position in this plot. (a) From a dataset with high purity. (b) Dataset after a loss of purity.

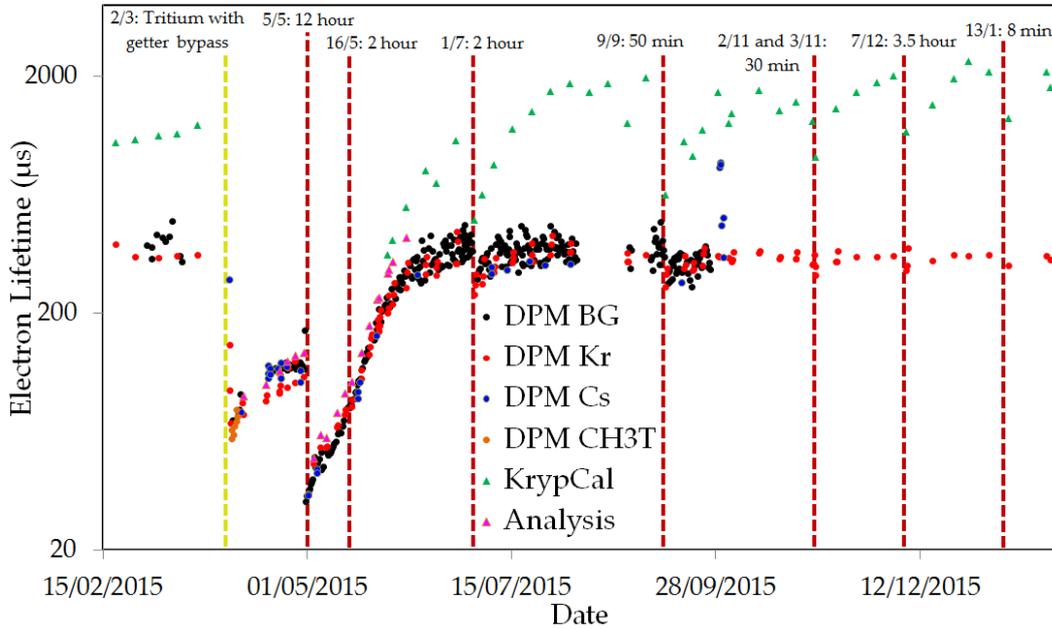


Figure 4.8: Measured electron lifetimes. Points are included from the DPM, from various types of data. Also included are results from KrypCal and from analysis of the data. The yellow vertical line is the time of the purity loss from the getter bypass. The red vertical lines are xenon circulation outages. The date and duration of the outage are labelled.

A spurious light source was identified by a LUX collaborator, the pulses were SE-like, between approximately 10 and 50 phe, but with a larger prompt fraction than normal. Therefore, a plot of prompt fraction against top bottom asymmetry was added to the DPM in order to monitor this. The light is sporadic, and shows as the top population in the “snowman” in Fig. 4.9. The rate of events in this region is also monitored over time. This is used to identify datasets that may be removed from the WIMP search dataset collection.

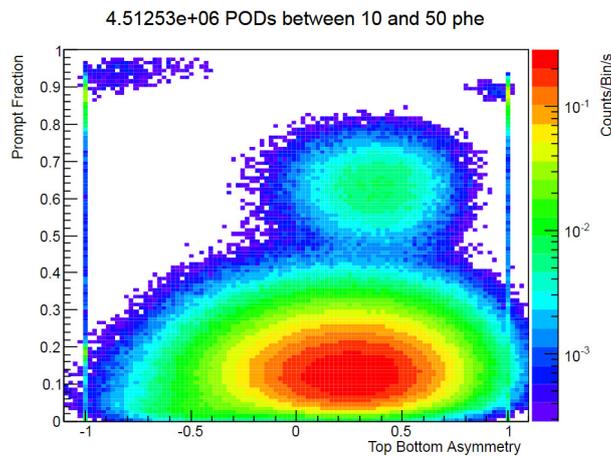


Figure 4.9: DPM plot of prompt fraction against top bottom asymmetry for pulses between 10 and 50 phe. The large population at the bottom are regular SEs, the spurious light shows as the population above this. The top left population are S1s.

4.3 Calibrations

4.3.1 Krypton-83m

^{83m}Kr is an internal source (i.e. one which disperses throughout the liquid xenon) and its deployment was described in section 3.2.8. The decay scheme is shown in Fig. 4.10. It de-excites in two stages, emitting a combination of conversion electrons, Auger electrons and a small fraction of γ -rays. Each decay adds up to a constant 41.5 keV electron recoil signal.

^{83m}Kr	
32.1 keV $t_{1/2} = 1.83 \text{ h}$	76% (30 keV IC) + (2 keV A) 9% (18 keV IC) + (10 keV A) + 2x(2 keV A) 15% (18 keV IC) + (12 keV A) + (2 keV A)
9.4 keV $t_{1/2} = 154 \text{ ns}$	95% (7.6 keV IC) + (1.8 keV A) 5% (9.4 γ)
^{83}Kr	

Figure 4.10: ^{83m}Kr decay scheme. It decays to ^{83}Kr via a short-lived intermediate state. IC is an internal conversion and A an Auger electron. [206]

The half-life of ^{83m}Kr is 1.83 hours, whereas that of the intermediate stage is only 154 ns. Therefore, the time separation between the S1s is a fast exponential and for this reason the two S1 pulses often overlap and are reconstructed as one pulse. On the other hand, they can also resolve into separate S1 pulses, as shown in Fig. 4.11. This separation in time for the S2s is not large enough to cause resolved pulses, as the S2s are much wider than S1. Although the two stages exist, ^{83m}Kr effectively gives a monoenergetic source that can be uniformly dispersed throughout the detector. An advantage of ^{83m}Kr over many other sources is that it can be deployed regularly, as the 1.83 hour half-life means that it decays away quickly so WIMP search data can be taken soon after. This means that any changes in detector conditions over time can be monitored with ^{83m}Kr (and in any case this energy scale is some 10 times higher than the WIMP region of interest).

The first application of ^{83m}Kr is to calibrate the energy scale of the detector. Energy is reconstructed from both the S1 and S2 signals using equation 3.10. ^{83m}Kr is one of many monoenergetic sources used to find g_1 and g_2 [207].

As this radioisotope disperses uniformly quickly (minutes), it provides a method for correcting position dependencies on the pulse area. The S1 position dependency is caused by variations in the light collection. The photon detection efficiency is lower near to the top of the detector, because it is further away from the bottom array where most of the S1 light is detected. The ^{83m}Kr S1 size is measured as a function of drift time, and parameterised. The S1 size is corrected to that at the centre of the detector (approximately the mean pulse area), so interactions near to the bottom have their pulse area decreased and vice-versa for those at the top. For S2s, the z dependence is due to losses of drifting electrons to electronegative impurities in the liquid. Therefore,

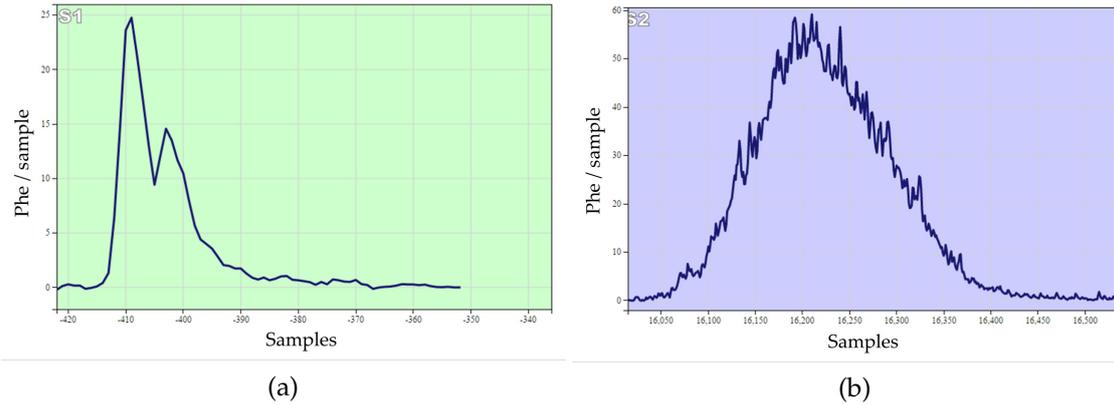


Figure 4.11: Example S1 (a) and S2 (b) from a ^{83m}Kr decay in LUX.

the ^{83m}Kr S2 size decreases with increasing drift time. An exponential is fitted to the S2 size as a function of drift time to measure the electron attenuation length and S2 sizes are corrected to that at the top of the detector. Regular ^{83m}Kr calibrations monitor purity throughout data taking.

There is also position dependence in the x - y plane. For S1 this is due to geometrical effects; S2s are affected by grid deflection (affects field and gas gap), detector levelling and variations in extraction field. The pulses are binned into a grid, and the areas are normalised to the centre of the detector.

The other correction from ^{83m}Kr is on the reconstructed position of interactions. There is some non-uniformity in the electric field near to the bottom of the detector, with an inwards radial component. This results in ionisation from events near to the bottom of the detector being mis-reconstructed further towards the centre since that is the position where the charge is eventually emitted into the gas. A look-up table is created as a function of radius, angle and drift time, and a correction to the radius is applied to each bin. This assumes the distribution of krypton is uniform throughout the detector.

The effects of the position corrections on the S1 and S2 pulse areas are shown in Fig. 4.12. For S1s, the mean value stays the same but the width of the distribution decreases. For S2 the mean shifts upwards and the relative width also decreases.

The correction to the event radius is shown in Fig. 4.13. Before the corrections near the bottom of the detector the events are slightly bunched towards the centre. Afterwards the events are more uniformly spread out and extend all the way to the detector wall.

4.3.2 Tritium

Tritium is a β^- emitter, with an endpoint of 18 keV. Therefore, it gives ERs within the WIMP search region and has many applications. It is an internal source so will be uniformly distributed throughout the detector and can provide much higher statistics in the fiducial volume than external sources. Tritium data is used to determine the

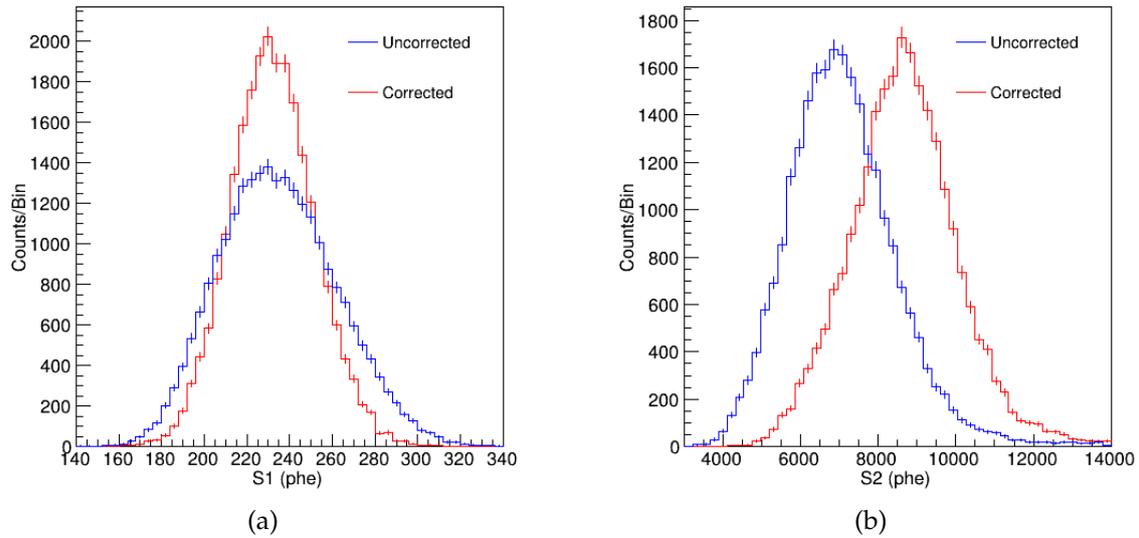


Figure 4.12: Position corrections on S1 and S2 pulse areas in ^{83m}Kr data. (a) S1 pulse area. (b) S2 pulse area.

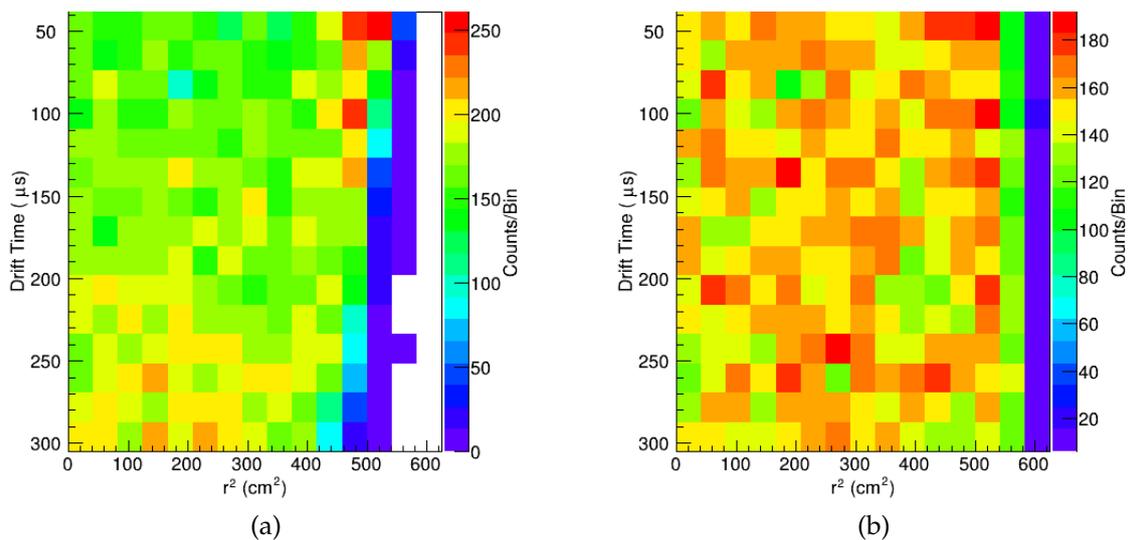


Figure 4.13: Effects of corrections to event positions from ^{83m}Kr calibration, (a) is before position corrections are applied and (b) is afterwards.

precise properties of the ER band, which characterises the background distribution and measures the discrimination. It is also used for measuring the charge and light yields, and the size of recombination fluctuations for electron recoils. For more details see Refs. [178] and [148].

In the first Run 3 analysis, the tritiated methane source was used to determine the pulse finding efficiency for ERs by comparing the true tritium spectrum to the measured spectrum in the detector. The efficiency can be measured for S1, S2, or reconstructed energy. The total rate of the tritium calibration source in the detector is not known precisely enough, and so an assumption that the pulse finding efficiency is

100% above some chosen point is required. Then, the true and measured spectra can be compared, and the efficiency loss is the ratio between them.

The first tritium calibrations for LUX were done in August 2013 towards the end of Run 3, with one small test injection of only 20 mBq, followed by a larger one of 800 mBq. Before this tests were performed with natural methane to verify that the LUX getter could remove it from the system. Approximately 1 ppm (g/g) of CH_4 was injected and was found to decrease exponentially with a time constant of 5.9 ± 0.1 hours [178]. This analysis is one of two independent analyses used to measure the ER detection efficiency from calibration data for the Run 3 WIMP search. The data was from the 800 mBq injection, which gave a total of 20 000 β decays in the active volume. The analysis was done with assistance from the University of Maryland group, who developed and deployed the source. It is the first time the ER detection efficiency can be directly measured in-situ.

Only a basic set of cuts are applied in this analysis, as it is important that the pulse finding efficiency is measured, rather than the efficiency of subsequent analysis cuts. Events with one S1 and S2 were selected within a fiducial volume with a drift time between 30 and 300 μs and radius of 17 cm. An S2 threshold is applied at 200 phe, as part of the pulse classification, so this is also included here. To remove high energy background events (the low energy background is essentially zero compared to the rate of tritium events), a maximum cut on S1 was applied of 140 phe (above the endpoint of the β^- spectrum).

The tritium β^- spectrum is well-known, both theoretically and experimentally. The theoretical spectrum used is shown in Fig. 4.14a, which gives the differential event rate as a function of energy.

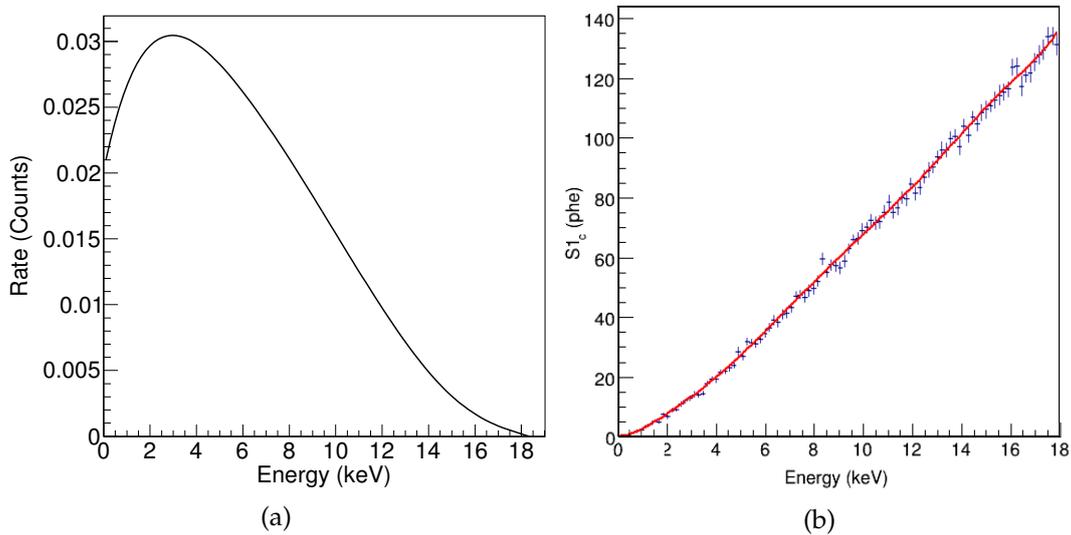


Figure 4.14: (a) Tritium β^- spectrum, from a Monte Carlo simulation of the three-body decay. The scale of the y axis is arbitrary. [208] (b) S1 pulse area from simulation of ERs up to 18 keV.

The Monte Carlo spectrum is known as a function of β^- energy. To measure the S1 detection efficiency, the spectrum must be found as a function of S1 too. The S1 response as a function of energy is found using LUXSim, the Geant4 [209] application developed for LUX simulations [210]. LUXSim incorporates the NEST package v0.98 [179] which combines measured scintillation and ionisation yields in liquid xenon from various experiments (pre-dating LUX at this point). ER events up to 18 keV were simulated (at evenly spaced points), recording only the S1 light, which greatly reduces computation time and allows higher statistics. The fiducial volume for the simulations was the same as that for the data. The simulation propagates the photons through the detector and includes the quantum efficiency of the PMTs and the the single photoelectron response for each PMT. This accounts also for the energy resolution of both channels. The mean S1 pulse area for each energy point was measured and parameterised by a polynomial fit, shown in Fig. 4.14b.

The fit is used to convert the energy spectrum of tritium into S1. As the Monte Carlo spectrum is a differential rate, the conversion is done using:

$$\frac{dN}{dphe} = \frac{dN}{dE} \frac{dE}{dphe}, \quad (4.1)$$

where dE/d_{phe} is given by Fig. 4.14b. Both the measured and the Monte Carlo spectrum are shown as a function of S1 in Fig. 4.15.

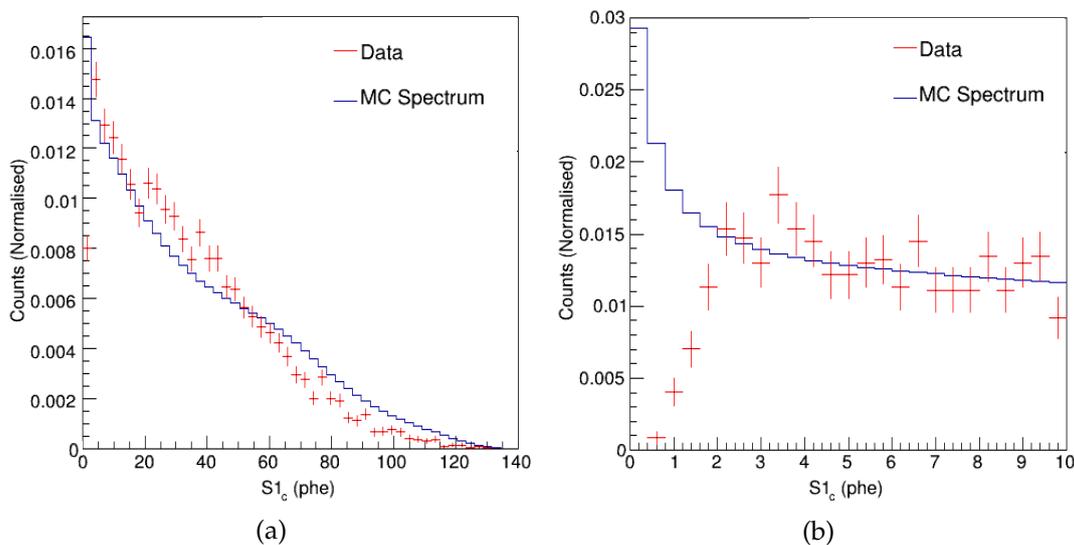


Figure 4.15: Tritium spectrum in S1 phe, measured from experimental data and from the Monte Carlo with the S1 response found from simulation. (a) Whole spectrum up to the maximum S1 phe. (b) Low part of spectrum showing efficiency losses.

The spectra are normalised to each other above 4 S1 phe, so a 100% pulse detection efficiency is assumed above this point. The overall spectrum shows some disagreement, but below 20 phe the agreement is within 20%. The Monte Carlo spectrum has a deficit between 20 and 50 phe, and an excess above that. In the low energy region between 4 and 10 phe the shapes of the two spectra agree well. The low energy part is

the most important for this analysis, as that is where the efficiency loss is expected to occur – this will mostly be due to requiring 2-fold coincidence requirement for S1.

The detection efficiency is then found by dividing the counts from the measured data by that of the Monte Carlo spectrum in each S1 bin. This analysis uses “golden” events, with both an S1 and an S2. Therefore, the efficiency should be regarded as the golden efficiency, the probability to find an S1 and S2, as a function of S1 phe. The efficiency is parameterised using an error function:

$$y = \frac{1}{2} \operatorname{erf} \left(\frac{x - a}{b} \right) + \frac{1}{2}. \quad (4.2)$$

The golden efficiency and fit are shown in Fig. 4.16.

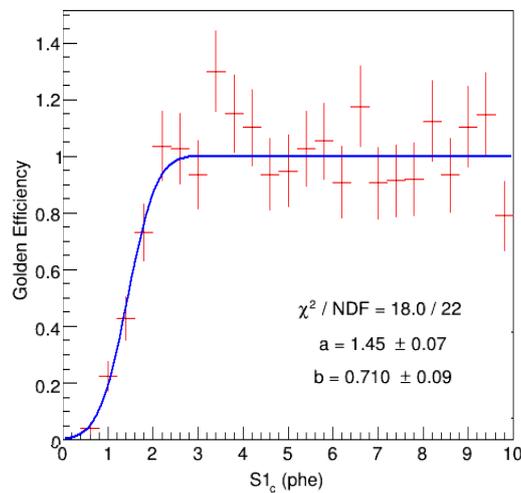


Figure 4.16: Golden detection efficiency from tritium data as a function of S1 phe.

The detection efficiency drops below 100% just below 2 phe. In the data processing, an S1 is required to have at least two detected phe in two PMTs within 100 ns. S1s with less than 2 phes can be detected as the area of two phes in separate PMTs can both fluctuate downwards.

Similarly, the golden efficiency can be found as a function of S2 signal. The S2 area is measured from the phe in the bottom array, denoted as $S2_b$, as was done for the WIMP search analysis. The initial method was to simulate interactions in the liquid xenon (recording only the electrons), and using the mean extraction efficiency and single electron size to convert this into $S2_b$ (for a particular energy). However, converting the Monte Carlo spectrum into $S2_b$ using only mean values of a resolution-smearred spectrum (as was done for the S1 spectrum) did not reproduce the data well. Going from electrons in the liquid to the detected S2 signal has several steps: drifting the electrons (with possible purity losses), extraction into the gas, generating the EL light and detecting the produced photons. The convolution of these processes results in an S2 distribution that is not Gaussian, particularly at low energies.

To account for the S2 distribution properly in simulation a different method was required. Instead of simulating particular energies, the actual tritium spectrum was

input into LUXSim. Again, only liquid electrons were simulated. Then, the extraction efficiency and mean number of phe per electron are applied using a toy Monte Carlo. The extraction efficiency is a binomial probability for emission into the gas, which was estimated to be 65% at the time of this analysis. For each extracted electron, the number of $S2_b$ is sampled from a Gaussian with a mean of 10.47 phe and standard deviation 4.5 phe, as measured from data. Losses of electrons to impurities are not included in the simulation as this is corrected for in the data. This process produces an $S2_b$ for each event, which are used to make the histogram for the Monte Carlo spectrum directly, as shown in Fig. 4.17a.

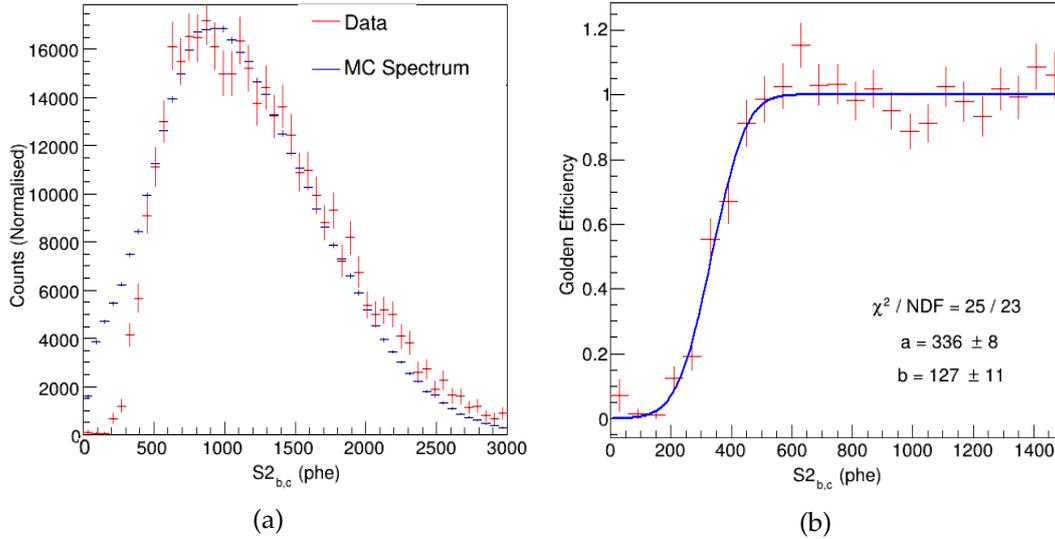


Figure 4.17: (a) Tritium spectrum in $S2_b$ from experimental data and simulations of the Monte Carlo spectrum. (b) Golden event detection efficiency from tritium data as a function of $S2_b$.

The two histograms are again normalised. There was some disagreement in the shape found at higher $S2$ phe, above 2000 phe. Therefore, the areas were normalised to each other between 600 and 2000 $S2_b$. The efficiency curve is found in the same way as for the $S1$, and is shown in Fig. 4.17b.

The efficiency reaches 100% at approximately 550 $S2_b$ phe (52.5 electrons). This is a large signal, above the 200 phe threshold. The NEST mean yield for a 0.8 keV ER is approximately 2 $S1$ phe and 500 $S2$ bottom phe, so the efficiency loss must be dominated by the $S1$ coincidence requirement.

The last way of presenting the efficiency is to keep the Monte Carlo spectrum in keV, and reconstruct the energy of events in the data using equation 3.10. g_1 , the photon detection efficiency had been estimated to be 14% at this point. For g_2 , the $S2_b$ was estimated as 10.47 phe and the extraction efficiency 65%. The reconstructed energy is shown in Fig. 4.18

There was significant disagreement between the two spectra, with the measured data giving less events in the tail of the distribution. A possible reason for the disagreement is if the values of g_1 and g_2 are not correct. Different values of the extraction

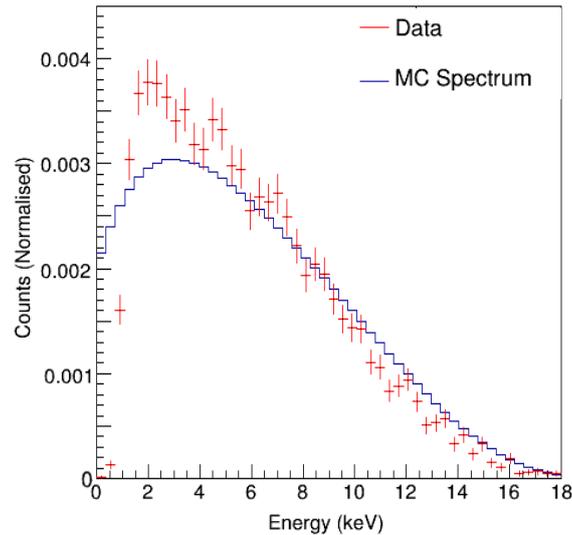


Figure 4.18: Tritium spectrum from data in reconstructed energy and the Monte Carlo tritium β^- spectrum.

efficiency were tried and the best fit of the data to the Monte Carlo spectrum was found to be with 56%. However, there was still an excess of events around 2 keV. One other effect could be due to the liquid purity. During the tritium injection it is possible that impurities could be introduced. This would mean that the z correction applied would not be using the current electron lifetime. This was checked by plotting the average position corrected $S2_b$ against drift time. This was found to be flat, which means that the electron lifetime was correct.

The tritium spectrum and energy scale were further investigated in the following months and reported in Refs. [207] and [178]. For the reanalysis of the Run 3 data the tritium beta spectrum was reproduced very well. The estimates of g_1 and g_2 were improved using monoenergetic calibration sources; the photon detection efficiency was $11.7 \pm 0.3\%$, the extraction efficiency $49.1 \pm 3.2\%$ and SE size was 12.1 ± 0.8 phd per extracted electron [178]. These values agreed with those found by fitting the Monte Carlo tritium spectrum and allowing g_1 and g_2 to float. In addition it was found that there was a probability for one VUV photon to emit two photoelectrons in the PMT that had to be accounted for in simulations. The finite detector resolution was applied to the Monte Carlo tritium spectrum, which was then compared to the data. After accounting for this the agreement between the measured tritium spectrum and the Monte Carlo spectrum convolved with the detector resolution was excellent. The final spectrum is shown in Fig. 4.19.

For the first Run 3 WIMP analysis the pulse finding efficiency was measured with several separate methods, the final curves are shown in Fig. 4.20. There was an independent analysis using the tritium spectrum and a similar one with an AmBe source (neutron emitter through (α, n) reaction). Also full NR simulations were created, with both S1 and S2 light and simulated waveforms for the event. The simulated data was processed through the data processing chain and the efficiency is found by comparing the output to the simulation truth. Good agreement was found between all of the different methods [185]. It was also necessary to determine the absolute efficiency, be-

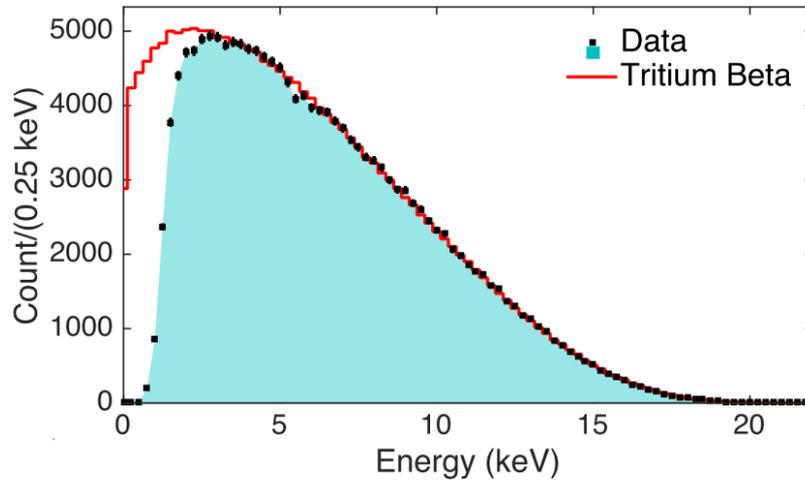


Figure 4.19: The measured tritium energy spectrum from LUX. The Monte Carlo tritium spectrum is convolved with the detector resolution. Figure reproduced from Ref. [178] with permission from the American Physical Society.

cause these methods had to assume 100% detection efficiency above a certain point, which may not be the case. This was estimated by visual inspection of waveforms from the AmBe NR calibration data to find golden events, which was then compared to the output from the data processing. The absolute efficiency was 98%, which was applied to all energies.

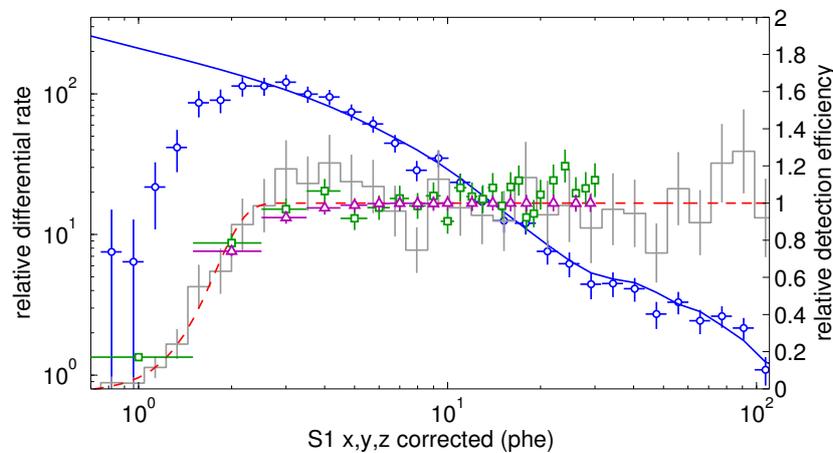


Figure 4.20: Efficiency curves from the Run 3 analysis. The blue points show the event rate (left axis) from AmBe (NR) data and the blue line is the corresponding simulation. The grey histogram shows the detection efficiency (right axis) from the AmBe, with the fit in red. The purple points show the efficiency from NR simulations, and the green points the ER efficiency measured from tritium (an independent analysis to the one presented here). Figure reproduced from Ref. [185], with permission from the American Physical Society.

The detection efficiency measured from calibration data confirms that the threshold in liquid xenon is very low. For ER, the efficiency reaches 50% at 1.24 keV [178]. For NR, the 50% point is at 4.3 keV and enabled the first WIMP search to go down to a mass of 2.4 GeV [185]. The tritium calibration allowed the first comprehensive measurements

of the scintillation and ionisation yields for ERs at low energy, along with recombination and its fluctuations. The NR efficiency from the simulations was applied to the WIMP signal model, and the measured tritium spectrum to the ER background model.

4.4 ^{85}Kr Background Analysis

^{85}Kr is one of the main backgrounds in LUX, being approximately 10% of the largest component, which is γ -rays from detector materials. In Run 3, the other backgrounds include the activated xenon isotope ^{127}Xe and ^{214}Pb (^{222}Rn decay chain). More details on the backgrounds are included with the Run 3 analysis in section 4.5.1. This analysis is a direct search for ^{85}Kr events in data to estimate the background. It provided an independent cross-check of the measurement from the sampling results for the Run 3 analysis.

4.4.1 Introduction

^{85}Kr makes up a small fraction of the total krypton in the atmosphere. It is a man-made radioisotope, originally produced in nuclear testing and during reprocessing of nuclear fuel. Xenon is obtained as a by-product of separation of air by fractional distillation. This process leaves some residual krypton content in the xenon, both being noble gases with reasonably similar boiling points, some of which is ^{85}Kr . The background comes from the main decay mode, a β^- decay with end point 687.4 keV and branching ratio of 99.563%. This gives some low energy β^- particles within the WIMP search region (0.9 to 5.3 keV). ^{85}Kr has a half-life of 10.76 years and therefore a significant amount will remain for the entire lifetime of an experiment.

Measurements of the $^{85}\text{Kr}/\text{Kr}$ ratio in the atmosphere are available, but there is some uncertainty and the results vary depending on season and location. Some examples include 1.3×10^{-11} [211], 1.7×10^{-11} [212], and 2.5×10^{-11} g/g [213]. In the commercial xenon used in LUX there were 130 ppb g/g $^{nat}\text{Kr}/\text{Xe}$. With a $^{85}\text{Kr}/\text{Kr}$ ratio of 2×10^{-11} the β decay from ^{85}Kr would give a background rate of 5 DRU [126], where 1 DRU is 1 event/kg/day/keV $_{ee}$. The target for the ^{85}Kr background was 0.2 mDRU, which is one quarter times the external γ -ray background predicted for a 100 kg fiducial volume. As krypton is a noble element it is not removed by chemical purifiers and must be removed before the start of the experiment. The level of krypton contamination required could not be achieved using commercial purification techniques. A method was developed using chromatographic separation, which reduced the krypton content to 4 ± 1 ppt g/g [200]. This processing was done off-site at CWRU before the underground deployment. This gives a background rate from ^{85}Kr of 0.19 mDRU, an acceptable rate for a WIMP search.

Currently, the estimate for the ^{85}Kr content in LUX comes from sampling measurements of krypton [191]. The sampling system was described in section 3.2.4. The solubility of the krypton in the xenon is an important factor, as the amount of krypton in the gaseous xenon sample may not be the same as that in the liquid xenon. This has been accounted for in the sampling results. Once the total krypton content is measured

the amount of ^{85}Kr is calculated using the ratio of $^{85}\text{Kr}/\text{Kr}$ as measured from the atmosphere. There are uncertainties in the $^{85}\text{Kr}/\text{Kr}$ ratio (approximately a factor 2) and the atmospheric measurement does not necessarily represent the krypton in this particular sample. Also, the background rate measured just after the chromatography does not account for any potential increases in the krypton content, for example if there was an air leak. Therefore, a direct search for the ^{85}Kr background is required that does not rely on the atmospheric $^{85}\text{Kr}/\text{Kr}$ ratio. We follow the method that was previously used in ZEPLIN-III [214] and XENON100 [151].

The analysis is possible due to another decay of ^{85}Kr , with a branching ratio of only 0.434%. This is a β^- decay with a 173 keV endpoint to an excited state of ^{85}Rb . After this the ^{85}Rb will de-excite, producing a 514 keV γ -ray. The half-life for this decay is $1.015 \mu\text{s}$ [215] — the decay scheme is shown in Fig. 4.21. In a dual-phase xenon detector this will produce a distinct signal, where there are two S1s separated by less than a few μs followed by their corresponding S2s. It is unlikely that there will be a random coincidence which could cause this particular signal with the correct energies. Even though this event is rare, it is possible to search for it in data due to the distinct signal. It is expected that there could be a total of approximately 8 $\beta - \gamma$ delayed coincidence events in 250 kg of xenon during 88.3 live days, based on the $^{85}\text{Kr}/\text{Kr}$ ratio measured from the atmosphere.

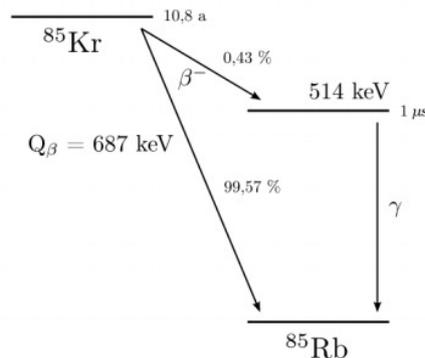


Figure 4.21: ^{85}Kr decay scheme [215].

Sampling measurements throughout Run 3 give a mean value of 0.17 ± 0.10 mDRU between 0.9 and 5.3 keV [126] [204]. Most of the systematic error on this comes from the $^{85}\text{Kr}/\text{Kr}$ ratio. If there is a sufficient number of $\beta - \gamma$ delayed coincidence events the statistical error could be small enough to give a measurement with a lower uncertainty. However, this is not likely to be the case for Run 3.

An illustration of the expected signal topology is shown in Fig. 4.22. The S1 from the β decay must occur before that from the γ -ray, and the time between them (t_{S1}) indicates the time between the two decays. The time between the S2s reflects instead the different vertical coordinates of the two interactions. Therefore, the S2s from the two decays can either overlap or be separated in drift time. In addition, the γ -ray can multiple scatter and give more than two S2s, which complicates the analysis.

This analysis relies on finding events with two S1s before any number of S2s. Note that we search for a mono-energetic γ -ray following a softer, variable-energy β particle.

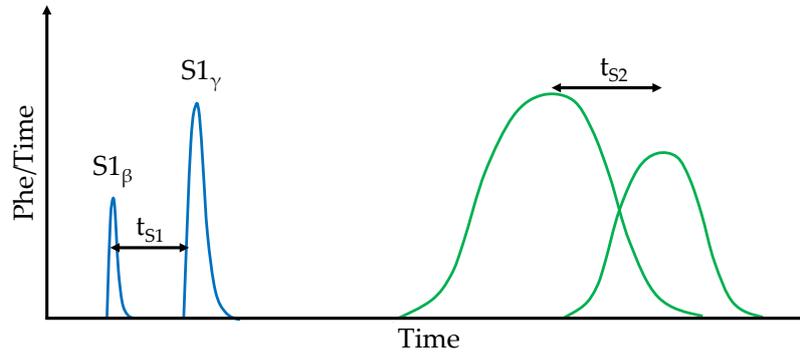


Figure 4.22: Illustration of event topology from a ^{85}Kr delayed coincidence. The time between the S1s indicates the time between the two decays. The S2 separation depends on where the γ -ray interacts, so they can either be separated or they can overlap if at the same height.

There are two versions, the first is an S1-only analysis to avoid the complication of reconstructing the combined energy from the multiple, possibly overlapping, S2s. This gives an initial estimate of the ^{85}Kr background. The second analysis does use the combined S1+S2 energy reconstruction which will improve the energy resolution over the S1-only analysis and give a better estimate of the ^{85}Kr content. For separate S2s, this requires reconstructing the energy of both the β^- and γ -ray. With only one S2 the total combined energy of the two decays is used.

The first step for this analysis is to calibrate the energy scale at >100 keV energies and resolution using background peaks in the WIMP search data. After other selection cuts, a search region in t_{S1} can be defined around the half-life of the intermediate state and the number of signal events counted.

4.4.2 S1 energy calibration

For the S1-only analysis the energy is estimated without using the S2. Therefore, equation 3.10 is not used. Instead a conversion between S1 phe and energy is required. The background spectrum in S1 phe was found. However, due to the low resolution in S1 alone it is difficult to find any peaks in the distribution.

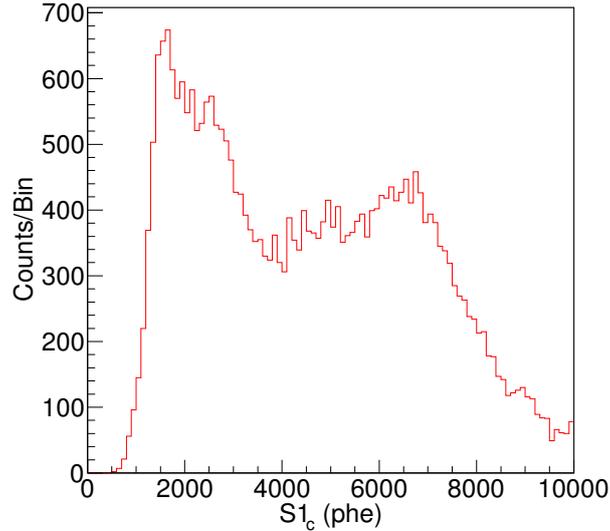


Figure 4.23: S1 spectrum of single scatter events.

To attempt to find a peak in S1 phe, a cut was applied in combined energy around the ^{214}Bi peak (as shown in section 4.4.3). However, this did not work, as the peak that was seen in S1 phe was not due to the actual γ line but an artifact of the combined energy cut. Using ^{137}Cs calibration data was also investigated but the line at 661 keV could not be identified in S1 only.

Therefore, the only way in which this could be calibrated was to use the ^{83m}Kr calibration data. Although the S2s from the two decays cannot be separated, the S1s can. Therefore, the 32.1 keV line from the first decay is used as a calibration point. It was difficult to find another good calibration point in S1 only.

Table 4.1: Summary of resolution for the peak from ^{83m}Kr in S1 (position corrected) phe.

Line	True Energy (keV)	S1 Area (phe)	σ (phe)
^{83m}Kr	32.1	218 ± 1	22.1 ± 0.9

For the signal peaks the mean S1 phe is taken to be the NEST mean yield for that energy. As there are no data points at high energy, extrapolating using equation 4.3 will not provide an accurate estimate. Therefore, the S1 only resolution is taken to be the same as that from the 32.1 keV ^{83m}Kr line. As the resolution is expected to improve with increasing S1 phe this is a very conservative approach. Taking 3σ around the signal mean should therefore give 100% acceptance.

Table 4.2: Summary of resolution in S1 for the ^{85}Kr signal.

True Energy (keV)	S1 Area (phe)	σ (phe)
173	909	79
514	2670	232

This over-estimate of the resolution will come at the cost of including more background events. Therefore, this preliminary S1 only result using only the ^{83m}Kr calibration could be improved. Nevertheless, this version of the analysis is only a check, and the combined energy analysis will always be the more sensitive measurement. The combined energy scale calibration in the next section works well, which confirms the S1 scale is also suitably calibrated.

4.4.3 S1 and S2 energy calibration

The ^{85}Kr events are selected based on their energy, so it is necessary to calibrate the energy scale and resolution in order to define the cuts. First the energy scale is calibrated from both the S1 and S2 signals. This is done separately for single scatter and for multiple scatter events. For single scatter events the energy is reconstructed using the standard formula (see equation 3.10). For multiple scatters the same formula is used but the S2 is the sum of all S2s in the event and therefore this is not strictly correct since the yield changes with γ -ray energy. Nevertheless, a reasonable energy reconstruction can be performed.

A cut on the S1 area in the top and bottom arrays is used to select γ -rays [216] and remove α -particle events. The value of these cuts is shown in Figure 4.24. The only other cuts are a fiducial volume and a “bad event area” cut. The bad event area is defined as the total event area minus the area of the S1 and S2s, and was also used in the standard Run 3 analysis. This is required to remove the delayed electron emission (‘electron trains’) following large S2 signals. The resulting background spectrum is shown in Fig. 4.25 for both single and multiple scatters.

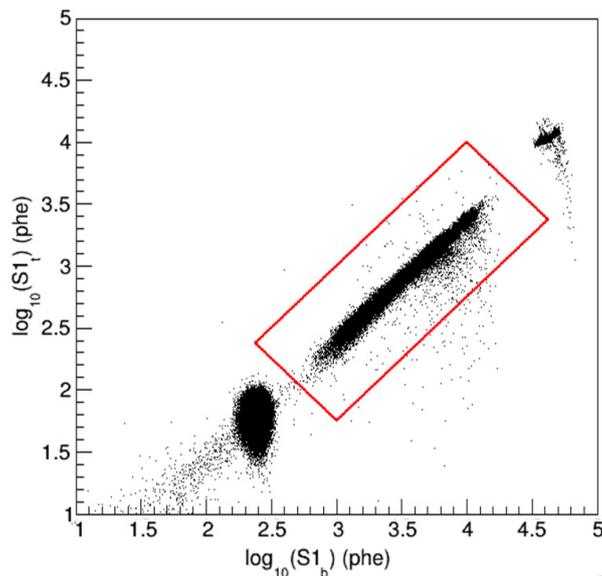


Figure 4.24: Selection cut for γ -rays, indicated by the red lines. The population of α events is at the top right and the events at the lower left are from residual ^{83m}Kr calibration.

There are several peaks that can be resolved in the energy spectra. The 511 keV positron annihilation line would be the ideal calibration for the 514 keV γ -ray; how-

Table 4.3: Cuts used to select events in addition to the cut shown in Figure 4.24.

Cut	Min	Max
Event Radius	0 cm	18 cm
Drift Time	38 μ s	305 μ s
Bad Event Area	0 phe	100 phe

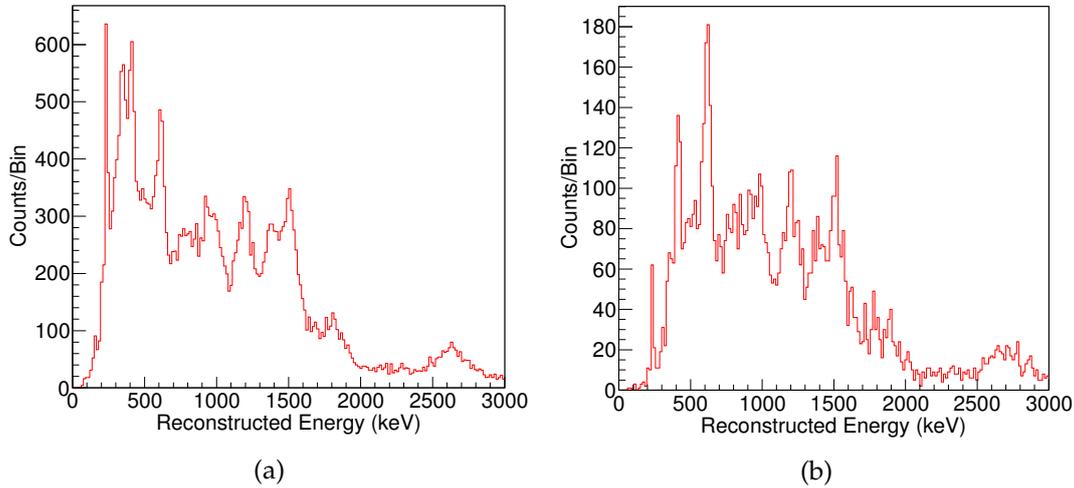


Figure 4.25: Energy spectrum for background γ -ray events. (a) Single scatters (one S1 and S2). (b) Multiple scatters (one S1 and two or more S2). Some of the peaks include ^{214}Bi (from ^{238}U chain) at 700 keV, ^{60}Co at 1170 keV and 1330 keV, ^{40}K at 1460 keV and a combination of ^{208}Tl and ^{60}Co at 2500 keV. There is some systematic offset, which is significant at high energy.

ever, this peak is not visible in either single or multiple scatters. Also note that the ^{83m}Kr calibration line at 41.6 keV is not useful here as it is a combination of various photons and electrons picking up different response yields. In addition, the second decay occurs very soon after the first (mean of 150 ns) and its response in scintillation and ionisation is affected by the presence of excitons from the first decay.

For the single scatters the peaks used for the energy scale calibration are ^{131m}Xe (164 keV), ^{129m}Xe (26 keV) and ^{214}Bi (609 keV) as they are all single γ -ray decays; which are shown in Figs. 4.26 and 4.27. The line from ^{127}Xe at 208 keV is an electron capture followed by γ -ray emission and the total spectral width has contributions from these two components. It is still fitted here because of its proximity to the other lines but not used in the energy calibration.

For the multiple scatters (specifically selecting events with more than one S2 pulse) the spectrum is sparser, and some peaks become difficult to resolve (see Fig. 4.25). The ^{131m}Xe peak and ^{127}Xe peaks cannot be seen. Therefore, the fit to the ^{129}Xe peak only has a single Gaussian plus exponential background. There are possibly some events from ^{127}Xe here, but this is assumed to be a negligible contribution. The fits to the peaks are shown in Figs. 4.26 to 4.29 and the results summarised in tables 4.4 and 4.5.

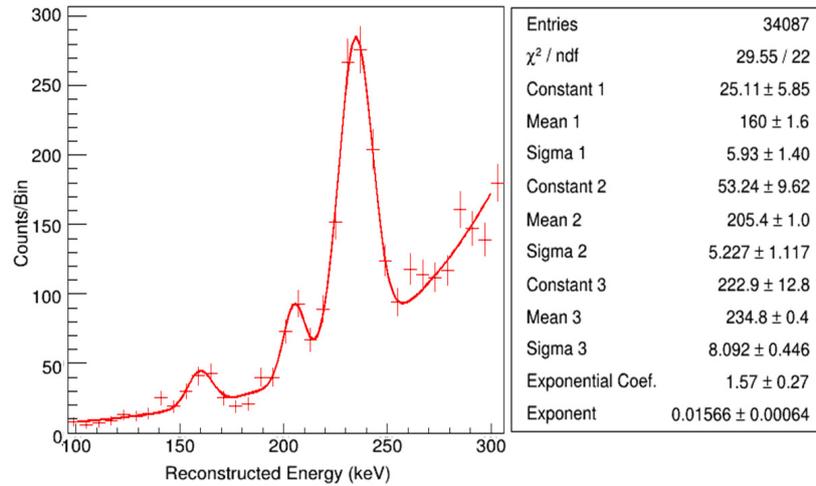


Figure 4.26: Fits to background from single scatter events. From left to right the peaks are ^{131m}Xe , ^{127}Xe and ^{129m}Xe . The fit is three Gaussians plus an exponential background.

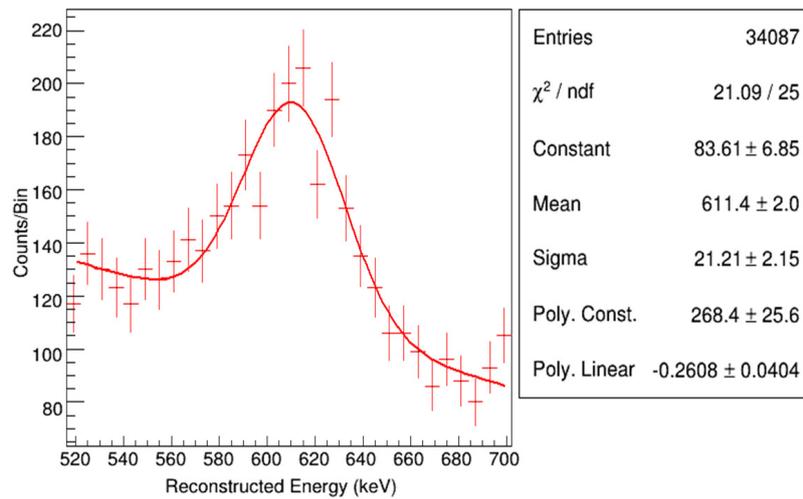


Figure 4.27: Fit to the ^{214}Bi peak from single scatter events. The background spectrum is taken to be linear over this energy range.

Table 4.4: Summary of reconstructed energies and widths for peaks in single scatter energy spectrum. The ^{127}Xe line is not used as a calibration point.

Line	True Energy (keV)	Reconstructed Energy (keV)	σ (keV)
^{131m}Xe	164	160 ± 2	6 ± 1
^{127}Xe	208	205 ± 1	5 ± 1
^{129m}Xe	236	234.8 ± 0.4	8.1 ± 0.4
^{214}Bi	609	611 ± 2	21 ± 2

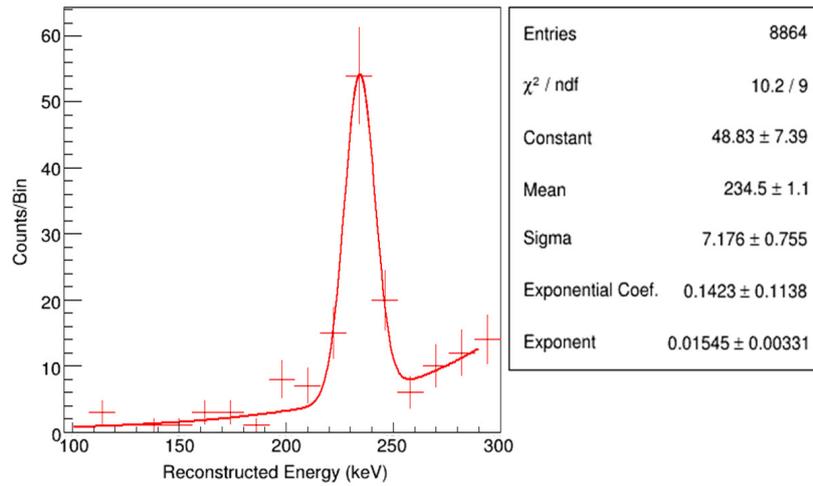


Figure 4.28: Fit to the ^{129}Xe peak for multiple scatters. The ^{131m}Xe and ^{127}Xe peaks cannot be seen here due to low statistics. The fit is a single Gaussian plus an exponential background.

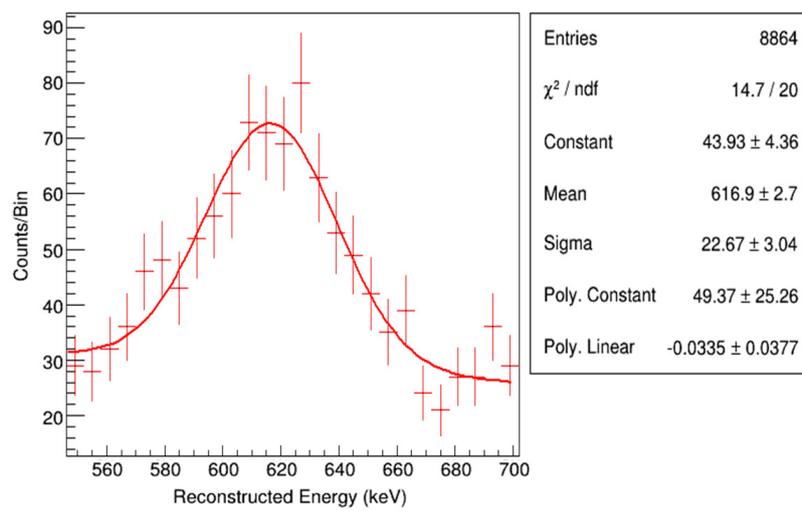


Figure 4.29: Fit to the ^{214}Bi peak from multiple scatter events. The background spectrum is taken to be linear over this energy range.

Table 4.5: Reconstructed energies and widths for peaks in multiple scatter energy spectrum.

Line	True Energy (keV)	Reconstructed Energy (keV)	σ (keV)
^{129m}Xe	236	235 ± 1	7.1 ± 0.8
^{214}Bi	609	617 ± 3	23 ± 3

The relative energy resolution (σ/μ) is fitted to the points from these peaks using the parameterisation:

$$\frac{\sigma}{\mu} = a + \frac{b}{E}. \quad (4.3)$$

Table 4.6: Energy resolution at the energies of the signal β endpoint and γ .

Energy (keV)	σ (keV)
173	6
514	15

It was found that the resolution for single and multiple scatters agrees within errors as shown in Fig. 4.3. The fit is to the two single scatter points from ^{131m}Xe and ^{129m}Xe , the standard calibration points. The ^{214}Bi line is then used to check that the extrapolation to higher energy is reasonable.

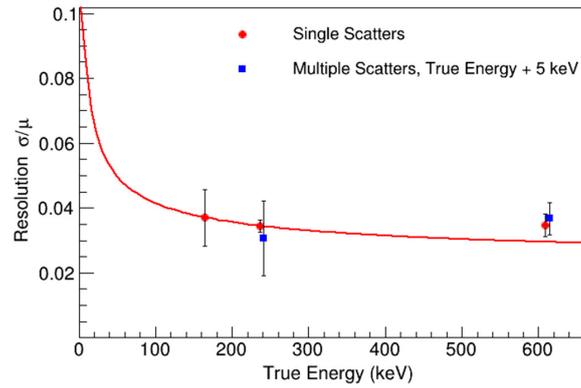


Figure 4.30: Fit of σ/μ for single and multiple scatters to the points at 164 keV and 209 keV. The result using equation 4.3 is $\sigma/\mu = 0.0214 + 0.200 \sqrt{E_T}$.

This fit was also checked against the resolution used in the simulation at low energies. It was found that the resolution in the simulation agreed very closely with the fit in Fig. 4.30 over the range 0 keV to 10 keV. The background spectrum has been well modelled with the simulation, so this is a good indication that the energy resolution used here is appropriate. There is still some systematic uncertainty due to the errors in the fits. This is taken to be 10% based on the size of the error bars from the single scatter points. The resolution at particular energies are then found from the fit in Figure 4.30.

4.4.4 ^{85}Kr search analysis

There are similarities between the S1-and the combined energy analyses. The main quantity from the $\beta - \gamma$ delayed coincidence analysis is the time delay between the S1s, t_{S1} . This is defined as the start time of the largest S1 pulse ($S1_L$) minus the start time of the smallest S1 pulse ($S1_S$). Only events with two S1s before any S2s are selected for the analysis. For S2s, a minimum (uncorrected) pulse area of 200 phe is required, otherwise the pulse is ignored. This is to remove single electrons, as in the Run 3 WIMP search analysis. In a histogram of t_{S1} , the signal events would give an exponential distribution with time constant equal to $\ln(2) T_{1/2}$.

The background for this search can be estimated by using the negative part of the time delay histogram, where the largest S1 came before the the smallest, which cannot be signal as the γ -ray is a much higher energy than the β^- . Counting the number of events with $t_{S1} < 0$ gives the background rate from random coincidences that pass

Table 4.7: Livetime losses in analysis.

	Time (days)
Total Run 3 WS Data	88.3
Datasets containing ^{83m}Kr	13.2
Detector Stability Cuts	3.2
Trigger Hold-off	1.8
Remaining live time	70.1

the selection cuts. A search region can then be defined to look for signal events. This requires a minimum value of t_{S1} because two pulses that are very close together will not be resolved by the pulse finder. The number of $\beta - \gamma$ coincidence decays in the search region are then counted.

The data used in this analysis is the Run 3 WIMP search data, but only the datasets which do not contain significant residual ^{83m}Kr from calibration. These were removed because the ^{83m}Kr has an energy within the range of the β from the delayed coincidence signal, and this was found to be give a significant source of background.

Detector stability cuts are also applied. This removes times of instability due to the grid voltage, detector pressure, liquid level or trigger rate. The livetime lost from these cuts is found after removing the ^{83m}Kr datasets. Then the livetime lost from the trigger hold-off is found from the remaining events. These inefficiencies are shown in Table 4.7.

4.4.5 S1-only analysis

Initially the analysis only used the S1s to estimate the energy. Due to the difficulty in finding an energy calibration in S1 only from the data the resolution used for the signal will be an over-estimate. For the maximum size of the $S1_S$ from the β decay it is only important that the cut is higher than the endpoint of the spectrum. The shape of the β spectrum means that there are only a few events near to the 173 keV endpoint, so as long as the cut is a small amount above this the signal acceptance will be very close to 100%. Therefore, the maximum cut on $S1_S$ is taken to be the number of S1 phe for a 173 keV β as given by NEST, plus 2σ , with σ estimated by taking the same energy resolution as the ^{83m}Kr peak at 32.1 keV, which is conservative.

The resolution at the 514 γ peak is estimated in the same way and the mean number of S1 phe is taken to be the NEST mean yield. The cut in $S1_L$ is taken to be 3σ on either side of this mean. This is again very conservative due to the uncertainty in the S1 resolution, but means that the signal acceptance is essentially 100%. This comes at the cost of a higher background rate. The values for the means and σ are summarised in Table 4.2.

First, the only cuts that were applied were the S1 area cuts and a fiducial cut of radius $r < 18$ cm and drift time $38 < t_d < 305$ μs . This encloses a mass of 118 kg, and is the same as was used in the Run 3 analysis. This was found to give a very large

number of candidate events. Plotting the time delay between the S1s (t_{S1}) with only these cuts showed some peaks in the distribution, as shown in Fig. 4.31.

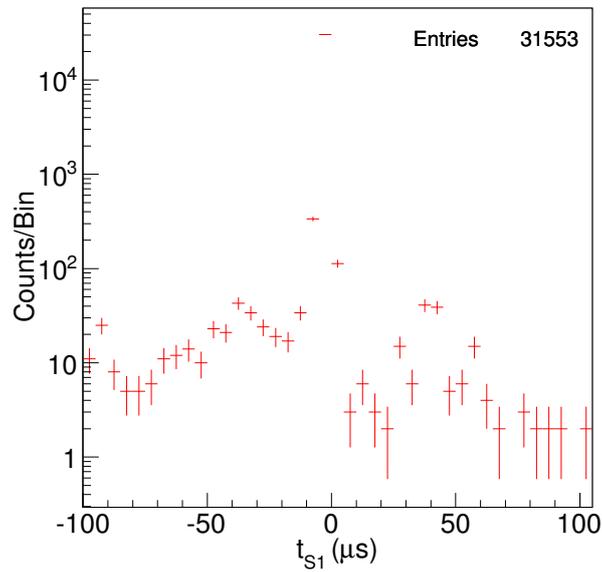


Figure 4.31: Time delay between the S1s after applying the S1 area and fiducial cuts.

The peaks around $-40 \mu\text{s}$ and between $-10 \mu\text{s}$ and $0 \mu\text{s}$ were investigated. All events in the population at $-10 \mu\text{s}$ to $0 \mu\text{s}$ were found to have small $S1_S$, below 200 phe. These were found to be due to the pulse finder splitting large S1s and finding small S1s in the tail; the $S1_S$ from this population is shown in Fig. 4.32a. The events were also mostly near the top of the detector as shown in Fig. 4.32b. Similarly the $-40 \mu\text{s}$ peak had $S1_S$ below approximately 50 phe, with more near to the top and bottom of the detector. These events also had a particular timing and hit pattern on the PMTs, which could be caused by afterpulsing after large S1s.

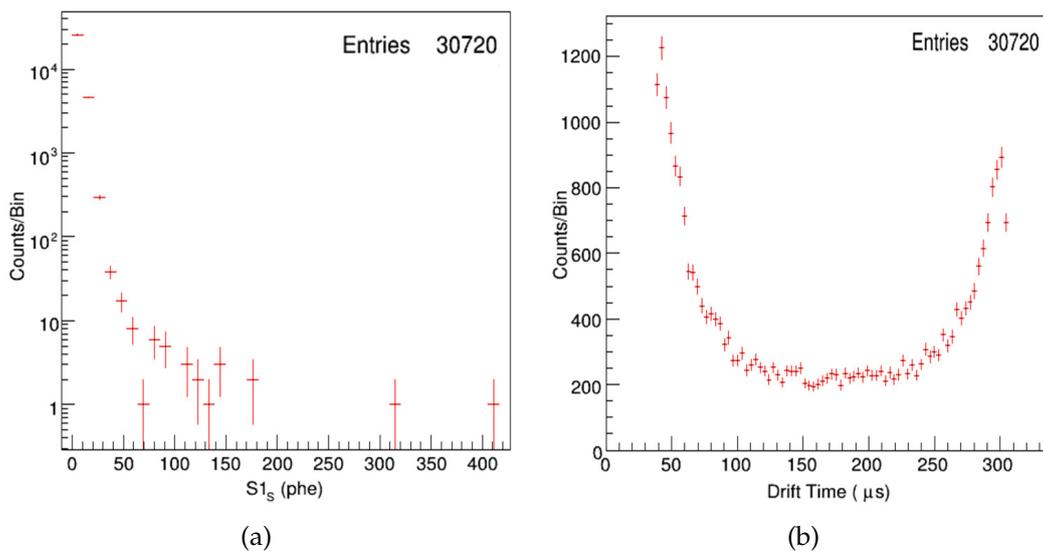


Figure 4.32: Events with t_{S1} between $-10 \mu\text{s}$ and $0 \mu\text{s}$, with the fiducial and S1 area cuts applied. (a) Area of $S1_S$ (corrected phe). (b) Drift time.

To remove the populations in the peaks of the t_{S1} distribution, a cut is placed at $S1_S > 202$ phe, which corresponds to a β with energy of 28 keV using the NEST mean yield. The signal acceptance is the fraction of the theoretical β^- spectrum above 28 keV. The cut on the $S1_S$ distribution is shown in Fig. 4.33, and was found to remove most peaks in the t_{S1} plot, with the resulting t_{S1} distribution shown in Fig. 4.34.

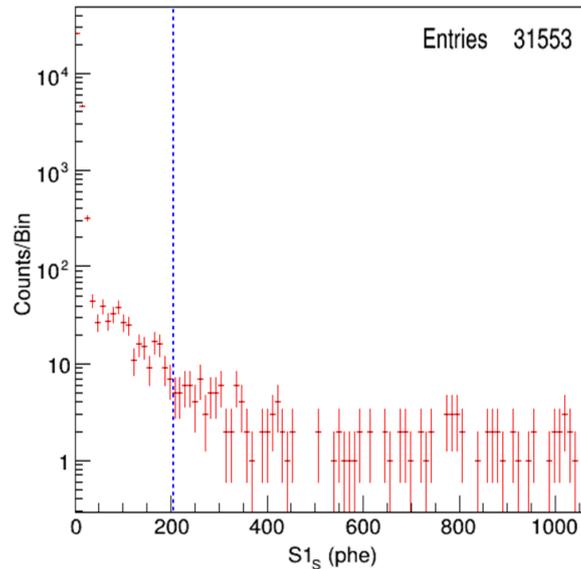


Figure 4.33: Area of $S1_S$ (corrected phe) after applying only the energy and fiducial cuts in the S1 only analysis.

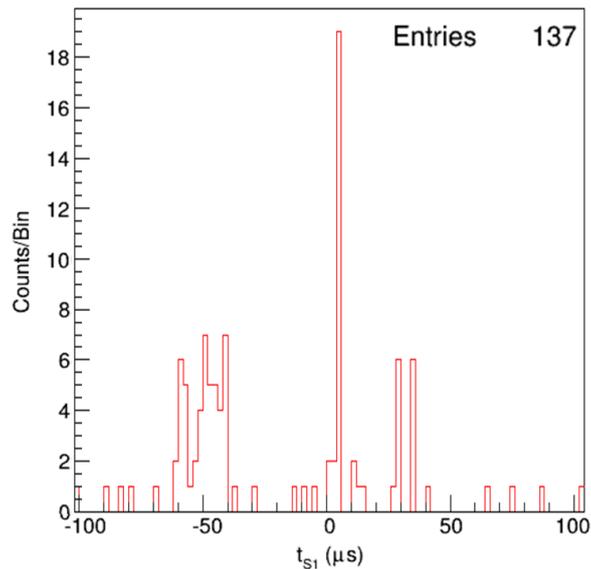


Figure 4.34: Time difference between the two S1s after applying the fiducial, S1 area, and $S1_S$ cuts

There was a remaining peak at a small negative t_{S1} and around $-50 \mu\text{s}$. After looking at the waveforms of these events it was found that the majority had a pulse between the two S1s that was not identified by the pulse classifier. To remove these

a cut is used on the maximum allowed area of pulses between the two S1s, a “bad S1 area” cut. There were also events where the $S1_S$ was two separate pulses clustered together by the pulse finder. Therefore, a cut is also applied on the width of the $S1_S$ using the RMS divided by the pulse area. The acceptance of this cut was confirmed to be very close to 100% by plotting the RMS/pulse area against pulse area for S1 pulses in single scatter events. The distributions before applying these cleanliness cuts are shown in Fig. 4.35.

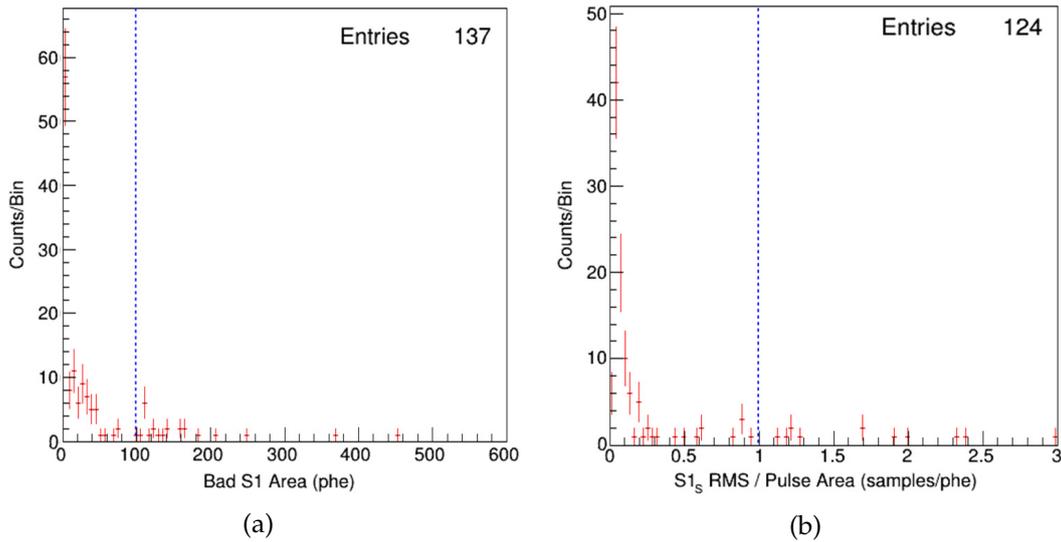


Figure 4.35: Distributions after the fiducial, S1 area and $S1_S$ cuts. Events to the right of the dotted line are cut in each figure. (a) Area of unidentified pulses between the two S1 pulses. (b) RMS width of the $S1_S$ divided by the pulse area, after also applying the cut bad S1 area cut.

Although this is an S1 only analysis, it is useful to use at least some information from the S2 with $\log_{10}(\Sigma S2_b / \Sigma S1)$, where Σ indicates taking the sum of all the S1s or S2s in the event. For a typical signal event, using the NEST mean yields gives this quantity to be 1.6. The plot of $\log_{10}(\Sigma S2_b / \Sigma S1)$ against the sum of the S1 areas in Fig. 4.36 showed a clear main population and some outliers. The signal acceptance of this cut will therefore be close to one.

All of the cuts in this initial analysis are intentionally chosen to be loose. The bad S1 area cut only cuts events where there is clearly a large pulse between the S1s that needs to be removed. The $S1_S$ pulse width cut is similarly chosen to be high. To check how likely a normal S1 is to fail this cut, the RMS pulse width divided by the pulse area is plotted for S1s from single scatters the background data. The fiducial and bad event area cuts are applied. Therefore, the remaining events should be clean single scatter events. Out of 34087 events, only 49 had a width/area greater than one. It is likely that some of these are two pulses clustered together. But, assuming that they are all true single S1s gives an estimate of the probability for a single S1 to fail this cut, which is 0.14%. The acceptance for this cut is therefore taken to be 1.0. Looking at Fig. 4.35a there are many events with a low bad S1 area. The main population of events becomes flat before the 100 phe cut. Therefore, the acceptance of this cut is taken to be very close to one.

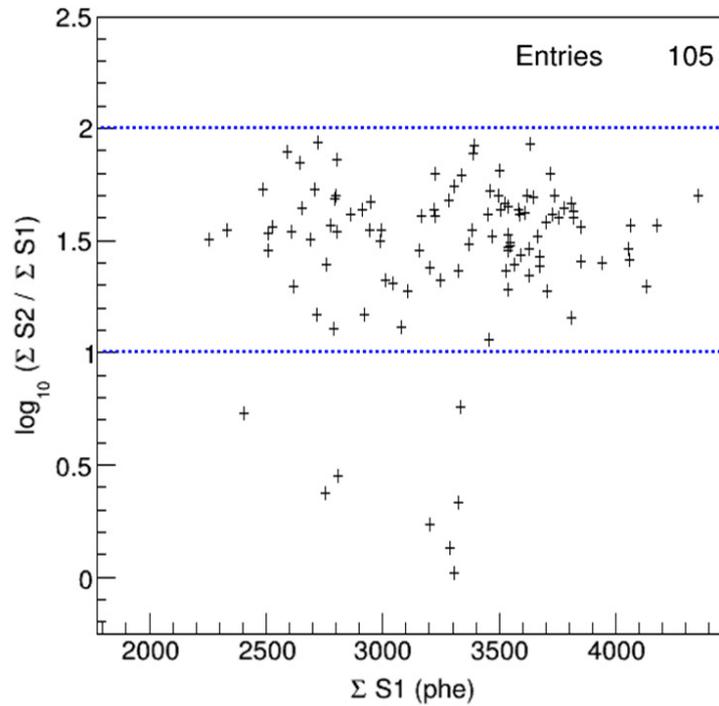


Figure 4.36: $\log_{10}(\Sigma S2_b/\Sigma S1)$ against S1 area after applying fiducial, S1 area, $S1_S$, bad S1 area and $S1_S$ pulse width cuts. The cut removes events that do not lie between the two lines which are typical of electron recoil interactions.

Table 4.8: Summary of cuts used in the S1 only analysis.

Cut	Minimum	Maximum	Acceptance	Comments
d_t	38 μ s	305 μ s	-	118 kg Fiducial
r	0 cm	18 cm	-	
$S1_L$	1974 phe	3366 phe	1	3σ cut around mean.
$S1_S$	202 phe	1067 phe	0.79	Max accepts 100%.
Bad S1 Area	0 phe	100 phe	$\simeq 1$	Area between the S1s.
$S1_S$ RMS / Area	0	1.0	1.0	
$\log_{10}(\Sigma S2_b/\Sigma S1)$	1.0	2.0	$\simeq 1$	

Table 4.9: Number of events remaining after applying each cut.

Cut	Events Remaining
Fiducial and Energy	31553
$S1_S$	137
Bad S1 Area	124
$S1_S$ RMS / Area	105
$\log_{10}(\Sigma S2/\Sigma S1)$	97

There is also a signal loss from the timing window used. This is chosen to be a minimum of $1\ \mu\text{s}$ and a maximum of $5\ \mu\text{s}$ which accepts 47% of the signal events ($T_{1/2} = 1.015\ \mu\text{s}$). With this loss and the other cuts, the total signal acceptance is 37%. The t_{S1} distribution after all cuts is shown in Fig. 4.37.

Calculating the total number of expected events using a nominal $^{85}\text{Kr}/\text{Kr}$ ratio of 2×10^{-11} and Kr/Xe ratio of 3.5×10^{-12} (from sampling results) gives an expected 1.2 events for the fiducial mass and livetime used in this analysis.

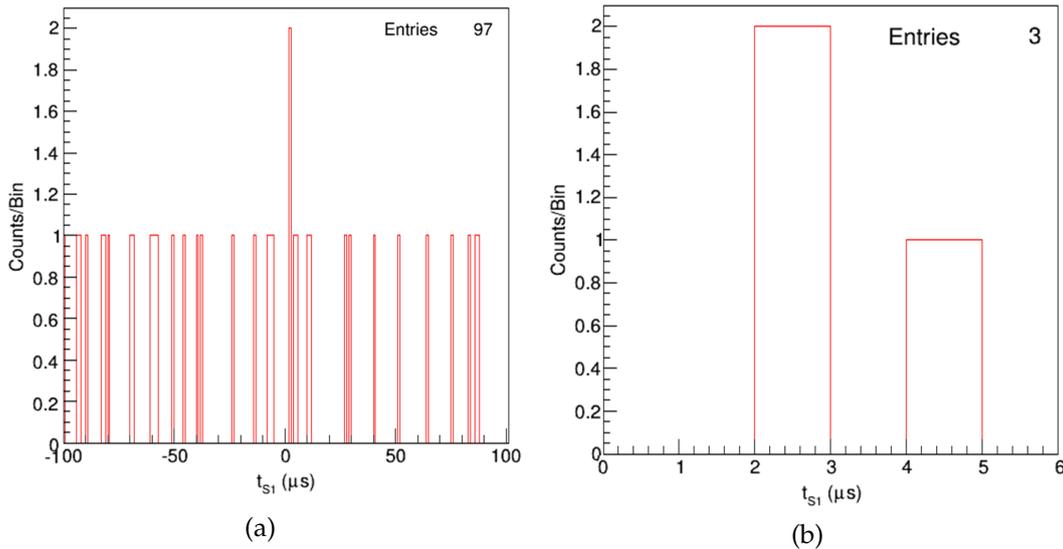


Figure 4.37: t_{S1} for events passing all cuts in the S1 only analysis. (a) Shows the whole range, (b) is only the signal region used in the analysis.

Using the region from $-100\ \mu\text{s}$ to $-1\ \mu\text{s}$ to estimate the background in the timing window of $1\text{--}5\ \mu\text{s}$ gives a prediction of 0.9 ± 0.2 events. There are 3 events observed in this region. A Feldman-Cousins test statistic returns a 90% CL interval of 0.2–6.5 delayed coincidence events. Using the branching ratio of 0.434% gives 47 to 1500 decays from the main $^{85}\text{Kr}\ \beta$. Over the energy range 0.9 to 5.3 keV this corresponds to a 90% confidence interval of 0.03 to 0.95 mDRU.

4.4.6 Combined energy analysis

A more robust analysis involves energy reconstruction from S1 and S2. The resolution in combined energy is much better than that in S1 alone and therefore more background events should be rejected by the energy cut.

For events where the S2s are resolved in time, the energy of the β and γ -ray can be reconstructed separately. For this analysis it is assumed that the β is a single scatter, as it is a relatively low energy electron with a low probability to Bremsstrahlung. The first step of this analysis is to decide which pulses are associated with the β and which with the γ -ray. The combination of pulses that is the most signal-like is chosen. For events with multiple S2s the energy and $\log_{10}(S2/S1)$ are calculated for all possible combinations of the S2s with the S1s (assuming that only one S2 is associated with the

$S1_S$ for the β). If there are more than two S2s the γ -ray is taken as a multiple scatter, so $\log_{10}(S2/S1)$ uses the sum of the S2 areas. Then, out of the combinations that pass a cut on $\log_{10}(S2/S1)$, the one which gives the closest reconstructed energy to the signal β and γ is selected. All reconstructed β energies which are less than the β endpoint are given equal weight.

Events with only one S2, where the β and γ -ray are at the same height, cannot have the two decays separated in this way. The energy reconstructed for these is also based on equation 3.10. In this case it is assumed that each of the decays was a single scatter. The two signal events deposit energy independently from each other, therefore the total S2 signal is the sum of the two individual decays, $S2_T = S2_\beta + S2_\gamma$. As the energy reconstruction formula is linear the total energy $E_T = E_\beta + E_\gamma$ can be written as:

$$E_T = W \left(\frac{S1_S + S1_L}{g_1} + \frac{S2_T}{g_2} \right). \quad (4.4)$$

It is possible that an event where the γ -ray multiple scatters appears as one S2, because the pulse finder clusters all of the S2s together. However, the same also occurs in the single scatter energy calibration. Therefore, the single scatter energy calibration should still be valid. These events also use the $\log_{10}(S2/S1)$ cut, with S1 being the sum of the $S1_S$ and $S1_L$.

It is important that, for each type of event, the overall signal acceptance is the same. It is not known how often the particles will single or multiple scatter in absence of a full simulation, so if the acceptance was different for each it would not be possible to find the total signal acceptance.

As the lower cut on $S1_S$ was found to be effective at removing the background in the S1-only analysis it is also applied here. The same populations of events were found after applying only the fiducial and energy cuts. In the events with only one S2, this cut is no longer independent of the energy cuts. Therefore, the total efficiency of the $S1_S$ cut, β and γ energy cuts for the multiple S2 events must be the same as that from the $S1_S$ cut and total energy cut for one S2 events.

To find the signal acceptance for events with multiple S2s the β spectrum is calculated and an energy resolution smearing is applied using a toy Monte Carlo. The relative width applied for a given energy is taken from equation 3.10 fitted to single scatter data. This gives the maximum reconstructed energy of the β , which was found to be 190 keV. This will not cut any signal events. Then the $S1_S$ cut acceptance is calculated from the smeared β spectrum below 28 keV (202 S1 phe). The acceptance for this was calculated to be 0.79, as in the previous analysis.

The γ -ray energy cut is chosen to be 2σ either side of the mean reconstructed energy. This was taken to be the same for single and multiple scatters. With this and the $S1_S$ cut, the total acceptance for the energy cuts is $0.79 \times 1 \times 0.95 = 0.75$.

For events with only one S2, the total energy is used as described by equation 4.4. The distribution is a Gaussian from the γ -ray convolved with the β spectrum. To calculate the total acceptance for the energy and $S1_S$ cut a toy Monte Carlo is used again. If an event would not pass the $S1_S$ cut it is removed (which has 0.79 signal acceptance),

to avoid counting the loss from the cut twice. The energy from the β and γ are added together for each to give the overall distribution in E_T . The area of this is integrated to give the points at which 2.5% of the events are removed at either side of the distribution, giving an overall 0.95 signal acceptance. Then the total acceptance for one S2 events is $0.79 \times 0.95 = 0.75$, as for events with multiple S2s.

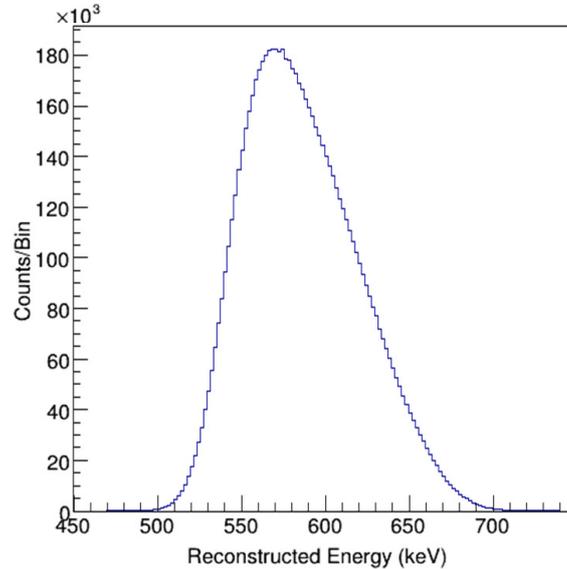


Figure 4.38: Energy spectrum found from the toy Monte Carlo for the total energy as described in equation 4.4. The cut on total energy is 550.7 to 677.0 keV.

The values used for the energy cuts are summarised in Table 4.10. The remaining cuts used are based on only the S1 pulses, and therefore do not depend on which combination of pulses is selected. The bad S1 area and $S1_S$ width cut are also applied here. After applying the energy, fiducial and $\log_{10}(S2/S1)$ cuts the time delay histogram presented in Fig 4.39 is obtained; this showed similar features to that from the previous analysis shown in Fig. 4.31. The peaks with negative t_{S1} were found to be populations with $S1_S$ less than 200 phe. Peaks on the positive side failed on the $S1_S$ width or the bad S1 area cut. The t_{S1} distribution before and after the S1 cuts is shown in Fig. 4.39.

There is additional information from the S1 pulses that can be used. As the cut on $\log_{10}(S2/S1)$ used here is only a loose selection, there are some events which pass the energy cut but do not have the expected S1 sizes. From the energy calibration for the S1-only analysis, it is expected that a 173 keV β would have a mean of 909 phe and a 514 keV γ -ray has 2670 phe. After applying all other cuts, there were events found with $S1_S$ much higher (1000s of phe) than expected and similarly for the $S1_L$. Therefore, a cut is also applied on the maximum area of the $S1_S$ of 1067 phe which is the mean sphe count for a 173 keV β (NEST mean yield) plus 2σ . An upper cut is placed on $S1_L$ of 3366 phe, which is the mean plus 3σ . Both are expected to have 100% signal acceptance, as the estimate of the resolution for the S1 is conservative. The $S1_S$ and $S1_L$ distributions with the position of these cuts are shown in Fig. 4.40. The values of all the cuts used are summarised in Table 4.11.

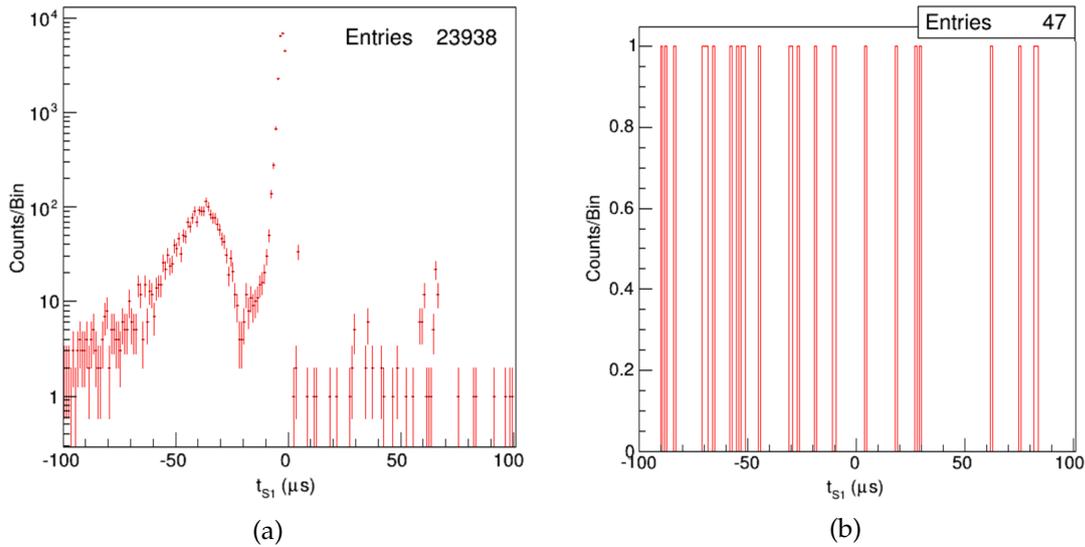


Figure 4.39: t_{S1} distribution, (a) with only the energy, fiducial and $\log_{10}(S2/S1)$ cuts, and (b) after also applying $S1_S$ minimum area, bad $S1$ area and $S1_S$ width cuts.

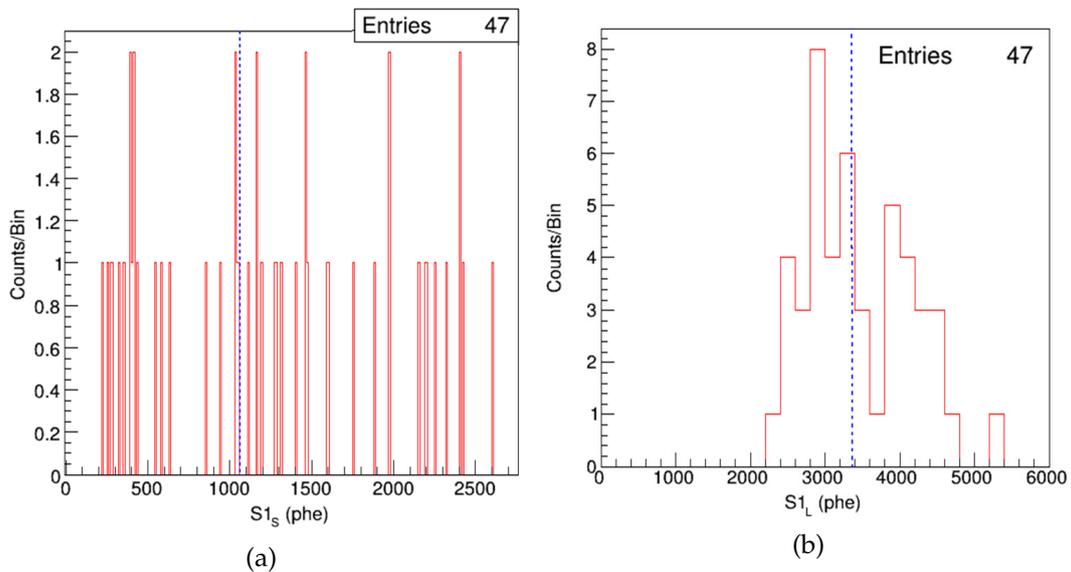


Figure 4.40: $S1$ sizes after applying the energy, fiducial, $\log_{10}(S2/S1)$, $S1_S$ minimum area, bad $S1$ area and $S1_S$ width cuts. (a) $S1_S$. (b) $S1_L$.

Table 4.10: Energy cuts used for different event types.

Cut	Minimum (keV)	Maximum (keV)	Comments
β	0	190	Always assumed to be a single scatter.
γ	482.9	545.1	2σ around mean
Total E	529.6	659.0	For events with one $S2$

Table 4.11: Summary of cuts used in the combined energy analysis. The energy, S_{1S} and $\log_{10}(S_2/S_1)$ cuts are all related, so are therefore applied at the same time.

Cut	Minimum	Maximum	Acceptance	Comments
Energy	-	-	0.95	See above
$\log_{10}(S_2/S_1)$	1.0	2.0	$\simeq 1$	
S_{1S}	202 phe	1067 phe	0.79	Signal loss is only from minimum cut.
S_{1L}	0 phe	3480 phe	1	3σ upper cut
Bad S_1 Area	0 phe	100 phe	$\simeq 1$	Area between the S_{1s}
S_{1S} RMS / Area	0	1.0 Samples/phe	1.0	Pulse width cut

Table 4.12: Number of events remaining after applying each cut. The fiducial, energy and $\log_{10}(S_2/S_1)$ cuts are applied at the same time.

Cut	Events Remaining
Fiducial, Energy and $\log_{10}(S_2/S_1)$	23938
S_{1S} Minimum	68
Bad S_1 Area	51
S_{1S} RMS / Area	47
S_{1S} Maximum	21
S_{1L} Maximum	4

The same timing window of 1 to 5 μs is chosen as before, which accepts 47% of signal events. The total efficiency for the signal is therefore 35%, combining the timing window with the other cuts. Calculating the expected number of signal events gives 1.1 $\beta - \gamma$ coincidences (using $^{85}\text{Kr}/\text{Kr} = 2 \times 10^{-11}$ and Kr/Xe ratio = 3.5×10^{-12}).

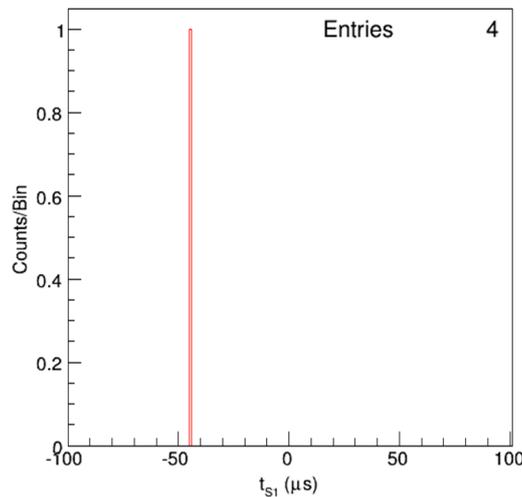


Figure 4.41: t_{S1} for events passing all of the cuts. Only one of the four events lies within the range shown, the others are at large positive t_{S1} .

Only one of the four events lies in the range $-100 \mu\text{s}$ to $100 \mu\text{s}$. Using the region from $-100 \mu\text{s}$ to $-1 \mu\text{s}$ to estimate the background gives 0.04 ± 0.04 events in the timing window of 1 μs and 5 μs . Using Feldman-Cousins statistics for 0 observed events and 0.04 background gives a 90% upper confidence limit of 2.4 $\beta - \gamma$ coincidence events. This corresponds to 553 decays from the main ^{85}Kr decay and the 90% upper CL between 0.9 and 5.3 keV is 0.26 mDRU.

4.4.7 Conclusions from ^{85}Kr background analysis

The delayed coincidence analysis provides a way to measure ^{85}Kr activity directly from data, which is independent from the gas sampling measurements and does not rely on the $^{85}\text{Kr}/\text{Kr}$ ratio measured from the atmosphere. Two different analyses were conducted, one using S1-only energy reconstruction and the other using combined (S1+S2) energy.

Both analyses set very useful limits on the ^{85}Kr content: the S1-only analysis sets an upper limit of 0.83 mDRU (90% CL) over the energy range 0.9–5.3 keV, whereas the analysis using the combined energy yields a more stringent upper limit of 0.26 mDRU. Both results are consistent with the measured ^{85}Kr background from the sampling of $0.17 \pm 0.10_{\text{sys}}$ mDRU.

Although this analysis does not currently improve the measurement it does confirm the result and gives confidence in the overall LUX background model. For Run 4 this analysis could improve as there will be more data and possibly more signal events. In 250 kg of xenon in 300 days there is expected to be 27 $\beta - \gamma$ coincidence decays (compared to 8 in 88 live days). An updated version of the combined energy analysis will be carried out for Run 4.

4.5 Run 3 WIMP Search Analysis

Run 3 enabled the first WIMP search with LUX, which took place in 2013. The world-leading results for spin independent WIMP scattering were presented in Ref. [185] and are summarised here. A total of 85.3 live days were selected for analysis, which excludes time lost to the trigger hold-off and DAQ deadtime. Using the slow control data, periods of detector instability were removed, when the liquid level, gas pressure or grid high voltage were outside of their usual operating ranges. The data were not blinded for this first analysis, and so only a basic set of cuts were made to minimise bias. All single scatter events with an S1 between 2–30 phe and S2 between 200–3300 phe were selected. The nuclear recoil energy range for the analysis is approximately 3–25 keV (corresponding approximately to 0.9–5.3 keV_{ee}). The lower limit of this region of interest is intentionally conservative, as this was the lowest energy at which direct NR yield measurements were available at the time [217] [218]. An additional cut was implemented based on the fraction of pulse areas in the event not accounted for by the S1 and S2 pulses ('bad area cut'), which was to remove events with high single electron rates that occur following large S2s. This cut corresponds to a live time loss of 0.8%. The final cut is the fiducial volume, which corresponded to drift times between 38 and 305 μs and a radius less than 18 cm from the centre of the detector (the liquid surface is at a drift time of 0 μs and the cathode is at 328 μs ; this fiducial volume extends between 6 and 46 cm below the surface). Table 4.13 shows the number of events remaining after each analysis cut. Figure 4.42 shows the positions of events passing all cuts except the fiducial volume.

The calibrations were described previously in section 4.3. A combination of calibration sources was used to find the NR and ER detection efficiency. The tritium

Table 4.13: Events remaining after each analysis cut for Run 3.

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
Single electron background	19 796
Fiducial volume	160

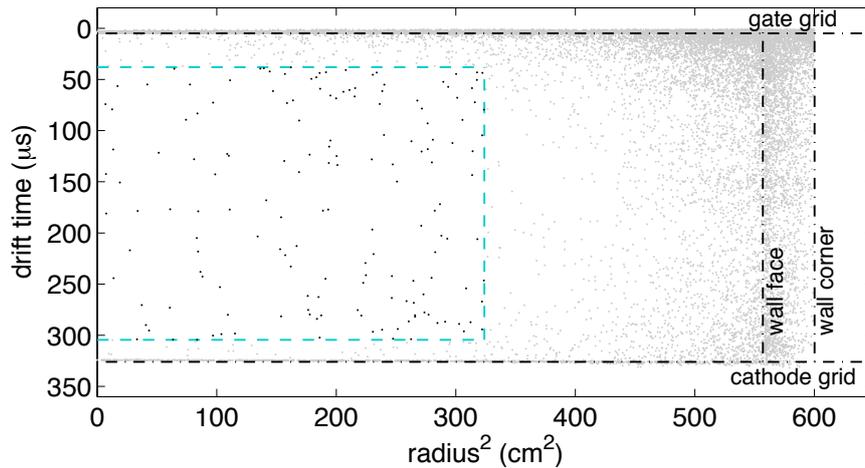


Figure 4.42: Location of events after all analysis cuts except the fiducial volume are applied. The blue dotted line shows the fiducial volume used in this analysis. Figure reproduced from Ref. [185] with permission from the American Physical Society.

calibration determines the ER band. For the NR band the external neutron sources AmBe (α, n) and ^{252}Cf (spontaneous fission) were used to characterise the detector response. Figure 4.43 shows the ER and NR calibration results within the fiducial volume. These neutrons cause an NR spectrum by elastic scattering which mimic the signature of relatively heavy WIMPs. The average discrimination for ERs was 99.6% in the WIMP search region. The fiducial mass determined from the tritium calibration was 118.3 ± 6.5 kg.

The WIMP signal spectrum is modelled with a standard isothermal Maxwellian velocity distribution [219] with $v_0 = 220$ km/s, $v_{esc} = 544$ km/s, $\rho_0 = 0.3$ GeV/c³, average Earth velocity of 245 km/s and a Helm form factor [220] [122].

The analysis uses a Profile Likelihood Ratio (PLR) on an event-by-event basis to calculate the observed limit. Four observables are used: S1, S2, radius and depth. Signal and background have different distributions in each of these observables. The PLR method is described in more detail in relation to the reanalysis of the 2013 LUX data in section 5.2. The backgrounds included for the PLR analysis are described in the following section.

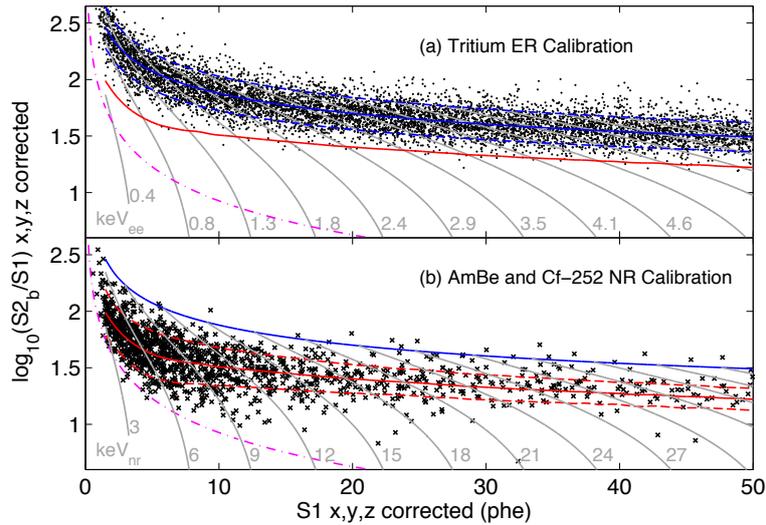


Figure 4.43: NR and ER calibrations within the fiducial volume. The top panel shows the ER calibration from tritium data. The bottom panel is NR, with data from AmBe and ^{252}Cf external neutron sources. Blue lines show the ER band mean, and red the NR band mean. The dashed red and blue lines show $\pm 1.28\sigma$ contours (10% tails on either side). The grey lines are contours of constant energy, and the magenta lines show the approximate location of the minimum S2 cut. Figure reproduced from Ref. [185], with permission from the American Physical Society.

4.5.1 Backgrounds

Backgrounds for our WIMP search come from NRs or ERs in the active xenon. These can come from radioactivity outside of the xenon (external backgrounds), or inside the liquid xenon (internal backgrounds). More details can be found in Refs. [126] and [216].

External γ -rays come from contaminants within the materials used to construct the detector and are the dominant background in LUX. Materials undergo a comprehensive screening programme, which involves counting the activity from samples of each in a detector sensitive to low activities of γ -rays [126] [187] [186]. LUX screening took place at the Soudan Low Background Counting Facility (SOLO) and at the Berkeley Oroville facility. The largest contribution comes from the PMTs, despite these being a specially-developed low background model. With the activities of all the materials, the background can be simulated to obtain the probability density functions (PDF) in each of the observable quantities used in the analysis (radius, drift time, S1 and S2). For high energy γ -rays the simulation was compared with the observed spectrum in data. The activity for each isotope in the simulation was allowed to float, and the best fit values were found to agree well with the screening measurements [126]. The PDF at low energy is naturally more concentrated near to the edges of the TPC. Within the chosen fiducial volume the total background rate from γ -rays was $1.8 \pm 0.2_{\text{stat}} \pm 0.3_{\text{sys}}$ mDRU.

The next largest contribution is the internal background from cosmogenic activation of the xenon. The peaks from the activated xenon isotopes were measured at the start of the run, soon after LUX was deployed underground allowing the rates to be measured. Only one of the activated xenon isotopes, ^{127}Xe , is a significant background

for the Run 3 WIMP search. This decays via electron capture with a half-life of 36 days. If the electron capture is from the L or M shells ($\sim 15\%$ probability) the energy of the resulting cascade is 5.2 keV or 1.1 keV respectively. The L shell cascade is just at the upper edge of the WIMP search region; approximately half of the events are accepted. It is assumed all of the lower energy cascades are accepted. This is usually accompanied by a high energy γ -ray from the decay of the excited daughter nucleus, ^{127}I . However, if the ^{127}I was in the 375 keV excited state (47% probability), it can decay to ground via a single γ -ray (17%). This γ -ray has a 2.6 cm mean free path in xenon and sometimes escapes from the active volume. Therefore, although the ^{127}Xe is uniformly distributed, this background is mostly concentrated near to the walls. The estimated average rate for the entire run was $0.5 \pm 0.1_{\text{sys}}$ mDRU.

Another background is the β^- source from the ^{222}Rn decay chain. Backgrounds from this chain were identified via α decay signatures which allows the rates of each isotope to be characterised. The ^{222}Rn daughters that produce ER backgrounds are ^{214}Pb and ^{214}Bi . The latter of these has a ^{214}Po daughter with a short half-life (160 μs), which has a 90% probability of being recorded within the same event as its parent decay and can be vetoed. Therefore, the most significant contribution from ^{222}Rn is from ^{214}Pb , a 'naked' β (no γ -rays) with $E_{\text{max}} = 1019$ keV and branching ratio = 11%. This is constrained by measuring the event rate between ^{127}Xe peaks at 238 and 410 keV, within a 30 kg fiducial volume. In this region other backgrounds from the walls are suppressed, so it is assumed that all activity is from the ^{214}Pb decay. The lower bound is obtained from the decay rates of other isotopes in the chain. The total contribution from ^{214}Pb is 0.11 to 0.22 mDRU.

The last background included in the PLR is ^{85}Kr , which was described in section 4.4. This produces an ER background from a 687 keV endpoint β^- decay. The measured rate from the sampling system used in the background PDF was 0.13 ± 0.07 mDRU. The direct analysis from section 4.4 sets a 90% upper limit of 0.26 mDRU, which confirms this is small compared to the dominant γ -ray background. A source that was not included in the Run 3 analysis was from wall events. These are from radon daughters plated out on the inner PTFE walls, which can be mis-reconstructed towards the centre of the detector. At the time of this analysis there was no model available for this background. The fiducial volume was chosen so that the expectation for leakage into the fiducial from the walls was essentially zero. The wall events are described further in section 5.3.1, where a PDF was developed for the reanalysis.

External backgrounds include radioactivity from the cavern walls and air, and cosmic ray induced events. Cosmic ray muons can create a background by generating neutrons in the water shield or surrounding rock [124], but are mitigated by the Davis cavern being 1.5 km underground. γ -rays and neutrons from radioactivity in the rock were effectively shielded by the water tank and were also found to be negligible [126].

NR backgrounds dominated by neutron elastic scatters are estimated by searching for multiple scatter events, which gave a conservative 90% upper limit of 0.37 events [185] for single-scatter neutron events (after studying the expected scattering multiplicity by Monte Carlo). The predicted background from simulations is only 0.06 events; this was treated as signal in the PLR analysis.

The total predicted ER background was $2.6 \pm 0.2_{stat} \pm 0.4_{sys}$ mDRU, which agrees with the observed $3.1 \pm 0.2_{stat}$ mDRU. With 99.6% discrimination this gives a total of 0.64 ± 0.14 events in the WIMP search region [185].

4.5.2 Results

The total number of observed events after all cuts was 160, and those are shown in Fig 4.44. There are no events below the NR band median, and the number of events and their distribution were consistent with the expectation from the ER background: the p -value for the null hypothesis was 0.35. The 90% confidence limit is shown in Fig 4.45.

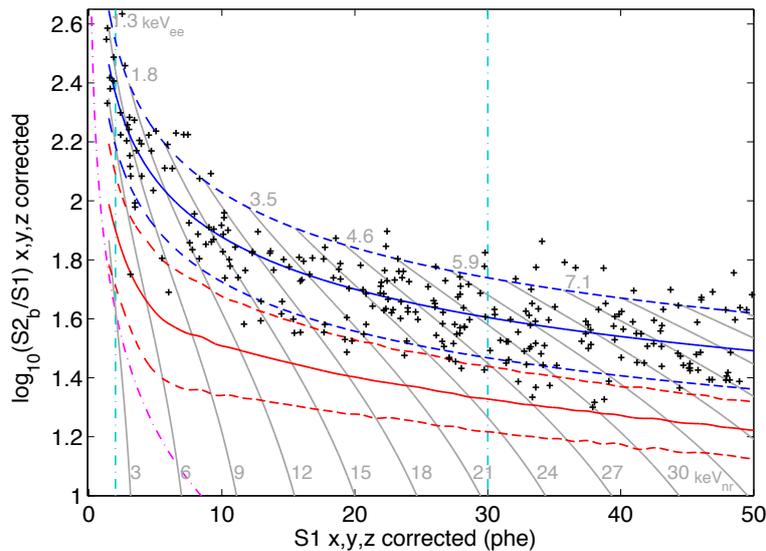


Figure 4.44: $\log_{10}(S_2/S_1)$ against S_1 for events that pass all analysis cuts in the LUX Run 3 exposure. The blue and red lines show the ER and NR bands respectively. The magenta line is the approximate location of the minimum 200 S_2 phe cut. The S_1 cut is indicated by the vertical cyan lines. Figure reproduced from Ref. [185], with permission from the American Physical Society.

The LUX result shown in Fig. 4.45 was the most stringent direct limit at the time of publication, improving over previous direct searches. The minimum excluded spin-independent cross section was $7.6 \times 10^{-46} \text{ cm}^2$ for a 33 GeV WIMP mass. LUX achieved exceptional technical performance including a high light collection efficiency (due to a high-reflectivity design), low and well-understood background rate, and high ER discrimination. The sensitivity for small WIMP masses is driven by the low energy threshold, which was achieved despite conservative assumptions for the detector response to nuclear recoils. In particular, previous detection claims at low masses from CoGeNT [225], CDMS II [227], CRESST II [228] and DAMA/LIBRA [132] [134] are ruled out as a SI WIMP signal. The tension between the LUX limit and the possible detections cannot be resolved with different astrophysical models [229] [230].

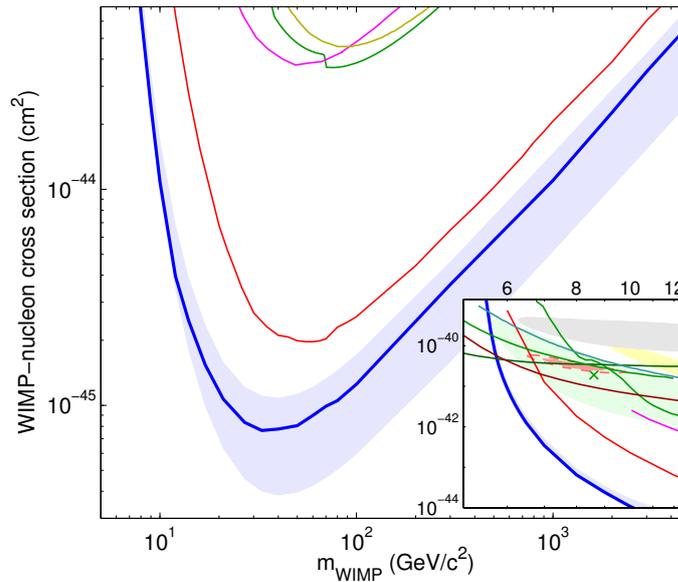


Figure 4.45: LUX 90% CL upper limits on the elastic SI WIMP-nucleon scattering cross section (blue line) and $\pm 1\sigma$ sensitivity prediction from repeated MC trials. The other limits shown in the main panel are: Edelweiss II (dark yellow) [221], CDMS II (green)[222], ZEPLIN-III (magenta) [223], XENON10 S2-only analysis (brown)[224], XENON100 (red) [151]. The inset for low masses shows the allowed region from the CoGeNT annual modulation signal [225] (light red), exclusion limit from the low threshold CDMS II analysis [226] (green line), 95% CL allowed region from the CDMS II silicon run [227] (green, with centroid x), 90% allowed region from CRESST II [228] (yellow), and the DAMA/LIBRA allowed region [132] as interpreted in Ref. [134] (grey). Figure reproduced from Ref. [185], with permission from the American Physical Society.

4.6 Conclusions

The first WIMP search run of LUX was very successful. It employed extensive calibrations, including a new internal tritium source. In section 4.3.2 this was used to measure the detection efficiency of low energy electron recoils. The background rates were low and the data were in good agreement with the expectations. A method for determining the ^{85}Kr background directly from data was presented in section 4.4, which confirmed the measured rate from the sampling system and that this background was sub-dominant.

The result was world-leading and excluded more regions of the WIMP parameter space. The non-blind analysis was conservative, using only a minimal set of data quality cuts, making the excluded region very robust. It ruled out many possible detections of low mass WIMPs, which resulted in the low mass WIMP hypothesis losing favour.

Chapter 5

Spin-Dependent Limits from the 2015 LUX Run 3 Reanalysis

At the end of Run 3, LUX published world leading limits on the spin-independent (SI) WIMP-nucleon cross section for WIMP masses above 5.7 GeV from 1.1×10^4 kg·days of exposure [185]. The minimum excluded cross section was 7.6×10^{-46} cm² at 33 GeV/c². A reanalysis of the same data was motivated by various improvements to the calibration, background model, position reconstruction and event selection. Also, an extra 10 live days of WIMP search data were added. The SI result was the most stringent direct limit for masses above 4.0 GeV/c², with a minimum excluded cross section of 5.6×10^{-46} cm² at 33 GeV/c². The reanalysis also motivated calculation of the first spin-dependent (SD) results from the LUX collaboration. This chapter presents the SD limits from this reanalysis, which were published in Ref. [231].

Firstly the nature of the SD interaction is described, including the details of the interaction, nuclear spin structure and calculation of the differential event rate for SD WIMP scattering off of a xenon target. The next section describes the profile likelihood ratio (PLR) method and its implementation for LUX. The remaining sections discuss the models for the PLR, the effects of nuisance parameters, the updates for the reanalysis and the final results.

5.1 Spin-Dependent Interaction

The exact composition of the WIMP and its interactions are not known. Besides the standard SI interaction, which can occur via Higgs exchange, a SD interaction is also possible. The latter arises from an axial-vector interaction between the WIMP and the nucleus and could be mediated by the Z^0 boson.

The WIMP velocity relative to the Earth, v , is small, which means the WIMPs are non-relativistic. Although the interaction is a WIMP-quark interaction, the whole nucleus is involved in this low energy regime. Direct detection experiments measure the WIMP-nucleus cross section.

The SD interaction is an axial-vector to axial-vector coupling, which implies exchange of a vector particle. It is a coupling between the spin of the WIMP and that of

the quark. The Lagrangian which describes this interaction is [232]:

$$\begin{aligned}\mathcal{L}^{SD} &= \frac{G_F}{\sqrt{2}} \int d^3\mathbf{r} j^\mu(\mathbf{r}) J_\mu^A(\mathbf{r}) \\ &= -\frac{G_F}{\sqrt{2}} d^3\mathbf{r} \tilde{\chi} \gamma \gamma_5 \chi \sum_q A_q \tilde{\psi}_q \gamma \gamma_5 \psi_q.\end{aligned}\tag{5.1}$$

In this equation, G_F is the Fermi constant, χ is the neutralino field, A_q are the WIMP-quark coupling constants, ψ_q the fields of the quarks. $J_\mu^A(\mathbf{r})$ and $j^\mu(\mathbf{r})$ are the hadronic and leptonic currents of the WIMP, respectively.

An axial-vector coupling of WIMPs arises in some beyond the standard model theories. In supersymmetry, Little Higgs theories, and universal extra dimensions, the WIMP candidate has both scalar and axial-vector couplings. Which of the two interactions is preferred depends on the exact details of the theory. However, there are possible scenarios where the SI interaction is suppressed leaving a predominantly SD interaction. This can occur in some regions of the SUSY parameter space [233], Little Higgs theories [234] and universal extra dimensions [235].

Equation 5.1 can be used to derive the differential cross section for SD WIMP elastic scattering from nuclei. Its detailed derivation can be found in Ref. [236]. The final result is:

$$\frac{d\sigma}{dq^2} = \frac{8G_F^2}{(2J+1)v^2} S_A(q),\tag{5.2}$$

where J is the total nuclear spin and $S_A(q)$ is the spin structure function for momentum transfer q . $S_A(q)$ is analogous to the form factor in the SI case; it describes the spin distribution within the nucleus and is the subject of the following section. Afterwards, the calculation of the differential event rate for SD WIMP-nucleus scattering is described.

5.1.1 Nuclear spin structure

The nuclear spin can be described using a nuclear shell model, similar to that for atomic electrons, but with distinct energy levels for the protons and neutrons [237]. This is because protons and neutrons are spin half fermions that obey the Pauli exclusion principle. Although one may expect many interactions between the nucleons in the dense nucleus, this cannot occur if all nearby states are occupied. For each nucleon type, each particle occupies the lowest available state and no two particles can have the exact same quantum numbers. The spin quantum number for protons and neutrons has two possible values, $\pm 1/2$. Two nucleons of the same type can occupy the same energy state with different spin quantum numbers, so the total spin of each pair is zero. Only when a nucleus has an odd number of protons and/or neutrons is there a significant nuclear spin.

There are several shell model calculations available for xenon isotopes, but the general method is the same. Firstly, the total inter-nucleon potential must be calculated with each nucleon moving within this effective potential. After this, the nuclear interactions must be specified. A Hamiltonian for the potential and interactions is then

constructed. This must then be diagonalised to find the eigenvalues, which are the energy levels of the nuclear states. The end result of the calculation is the spin structure function for the isotope.

The existing calculations differ in various ways, and there are limitations. The calculation of the nuclear potential can be different, and which nuclear forces are accounted for can vary. Earlier calculations were limited to only two-body interactions (1-body currents), but now calculations are available that include three-body interactions (2-body currents). When diagonalising the matrix for the Hamiltonian, all calculations have computational limitations. Therefore, the model space must be restricted as shown in Fig. 5.1. For heavy elements such as xenon, only the outer shells are included on top of an inert core of full shells. Nucleons in the states above the core are allowed to move, within a number of allowed configurations. There are truncations in both the possible configurations and the allowed states.

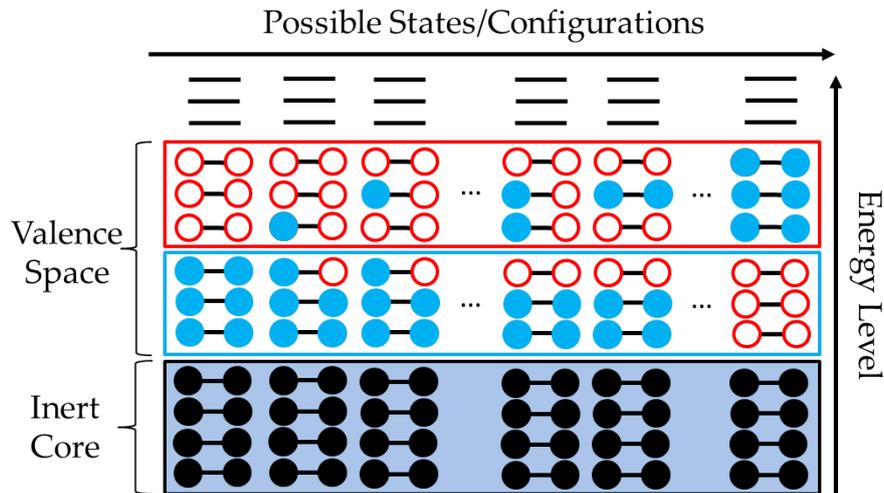


Figure 5.1: Energy levels and truncations in the nuclear shell model. The inert core represents complete shells, and nucleons here cannot be excited into higher states. The valence space contains nucleons that are allowed to move between states. The first column shows the ground state, where no nucleons are excited. The next two show possible configurations with one excited nucleon. Then there are two examples for two excited nucleons. This can continue until all of the nucleons in the valence space are excited, as shown in the final column. Truncations are implemented in two ways: one is to limit the number of allowed states, represented by the lines above the valence space. The other is to restrict how many ways the nucleons can be arranged within the valence space.

In order to find a WIMP-nucleon cross section instead of the measured WIMP-nucleus cross section an assumption of values of coupling constants to protons and neutrons is required. This leads to the standard proton-only or neutron-only SD WIMP interaction. Typically, the spin structure functions are presented in terms of this assumption. For leading order interactions (1-body currents), in the proton-only case the WIMP can only couple to protons and vice-versa for neutron-only. Therefore, for elements where there is an odd number of neutrons (such as ^{129}Xe and ^{131}Xe) there will be a significantly higher neutron-only sensitivity than proton-only. This is discussed further in the following section.

It is important to compare the different available structure function calculations. As computational power has increased the accuracy of the shell model calculations has improved due to less truncations of the model. Here some of the available calculations are described.

The calculation from Ressel and Dean (RD) [238] has been used for many SD results from other experiments. There are two similar potentials considered here, Bonn A [239] and Nijmegen [240]. Only leading order nuclear interactions are included. Due to computational limits both the model space and the number of allowed configurations were truncated. The model space is the $3s_{1/2}$ $2d_{3/2}$ $2d_{5/2}$ $1g_{7/2}$ $1h_{11/2}$ levels, on top of an inert ^{100}Sn core. The result was compared to experimental data via the nuclear magnetic moment. The agreement between the two potentials was good, with both the calculated magnetic moment and spin expectation values being similar. The agreement with the experimental magnetic moment improved over previous calculations, but there were still some differences.

Another calculation by Toivanen *et al.* [241] came after the RD calculation. The model space is the same, but the number of allowed configurations of nucleons is larger. A different nuclear potential was used, the Bonn-CD (Charge Dependent) potential [242], which again includes only the leading order currents. The results reproduced the experimentally measured nuclear magnetic moment better than the RD calculation. Also, the energy levels were compared to data. For WIMP-nucleus scattering the lowest levels are the most important. There was disagreement in the position and ordering of the energy levels between experimental measurements and this calculation.

Structure functions from Fitzpatrick *et al.* are available in Ref. [243]. This uses an effective field theory treatment of the WIMP-nucleus interaction, which considers all possible operators. This study restricts the operators to those that are Hermitian, spin one or less exchange particle, and up to second order in momentum exchange. After also removing any T-violating operators, those that remain are constructed from combinations of the WIMP spin, nucleon spin, WIMP-nucleon velocity, and momentum transfer. The relevant experimental quantities are the nuclear responses arising from the operators. There are five responses (that only depend on the nuclear physics): the standard SI interaction, the transverse and longitudinal SD components, angular momentum dependent, and spin-orbit. The form factor for any of the effective field theory operators can be written as linear combinations of these nuclear responses. The standard SD form factor is calculated from a combination of the longitudinal and transverse components and used in a shell model calculation. For the odd-neutron xenon isotopes the model space was the same as in the previous calculations, but quite severe truncations were made to the allowed configurations, and only leading-order interactions were included. There was no direct comparison to data in this work, which was considered “exploratory” in nature. It is important that an effective field theory treatment can be used to evaluate the spin structure function. With this treatment limits can be set on the other non-standard interactions, but in this work only the SD response was considered.

The last two structure function calculations considered are from Menendez [244] and Klos [236], which also use a chiral effective field theory treatment of the nuclear

forces. As well as the 1-body currents, this includes the leading order long-range 2-body currents (3 body interactions). The interaction potential [245] [246] has been successfully tested in other applications. The model space is the same as the previous works with an inert ^{100}Sn core, but with fewer restrictions in the allowed configurations. 2-body currents represent interactions between a WIMP and two nucleons and have a significant effect. Even for “proton-only” coupling, a 3-body interaction can occur between the WIMP, a proton, and a neutron. For isotopes with an odd number of neutrons and even number of protons (such as ^{129}Xe , ^{131}Xe), this significantly enhances the proton-only spin structure function as the neutron carrying most of the nuclear spin is still involved in the interaction. Reference [236] is an update to Ref. [244] where an assumption of low momentum transfers in the nuclear currents was removed. There is a theoretical uncertainty due to the inclusion of 2-body currents given in the paper. The energy states of ^{129}Xe and ^{131}Xe were reproduced better than in previous works, with the ordering of the lower states being correct.

A comparison between a selection of the other structure functions at low momentum transfer is shown in Fig. 5.2. The Klos calculation is shown with only 1-body currents and with 2-body currents.

The Klos structure functions were compared to the Fitzpatrick calculation in Ref. [247] at the 1-body level. Reasonable agreement was found, but there were differences. In Ref. [243] the pseudoscalar contribution was taken to be a separate response whereas in Ref. [236] it is included in the SD response. In Fig. 5.2 the Klos 1b currents are with this included, which is why these two structure functions look very different for the proton-only response. The Klos 1b structure functions without the pseudoscalar contribution agree reasonably well with those from Fitzpatrick. There are still some differences, which are attributed to the more severe truncations used in the Fitzpatrick calculation.

For the neutron-only structure factors there is reasonable agreement between the calculations. For ^{131}Xe the Toivanen structure function is smaller, and for ^{129}Xe the RD structure function comes out slightly higher. The effect of including 2-body currents is small for neutron-only. On average, for ^{129}Xe (^{131}Xe) the 2-body structure function is 19% (9%) smaller than the 1-body structure function (up to $u = 1$). However, there is a significant enhancement for proton-only: for ^{129}Xe (^{131}Xe) the 2-body structure function is 5.0 (3.7) times larger than the 1-body structure function (up to $u = 1$). Also, for proton-only there is disagreement even between the different 1-body calculations.

In summary, the structure functions that will be used in this analysis are the ones from Klos *et al.* [236]. It uses the largest model space available, with the least truncation to the allowed states. It includes the 2-body currents which are important, particularly for the proton-only interaction in even-proton isotopes. Also, the interaction and nuclear potential have been successfully used in other applications. The following section discusses how the spin structure function is implemented into the calculation of the differential event rate for WIMP scattering.

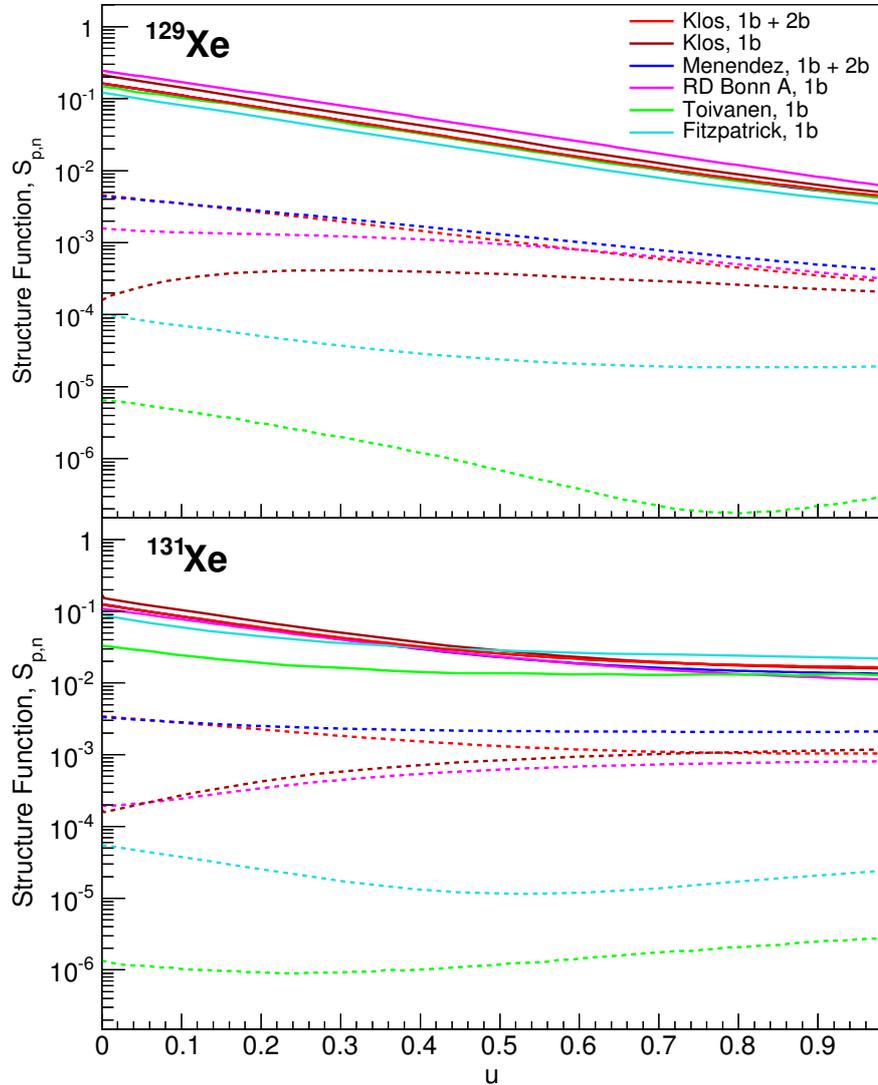


Figure 5.2: Comparison of structure factors from different nuclear shell model calculations for low momentum transfer. The top panel shows the ^{129}Xe structure factors and the bottom is for ^{131}Xe . Solid lines are neutron-only and dotted lines proton-only. The independent variable u is a function of the momentum transfer q , $u = q^2 b^2 / 2$, where b is the harmonic oscillator length (2.2853 fm for ^{129}Xe and 2.2905 fm for ^{131}Xe). The maximum shown is $u = 1$ which corresponds to a recoil energy of 62.5 keV (and u is a linear function of recoil energy). Included are calculations from Klos [236], Menendez [244], Ressel and Dean [238], Toivanen [241], and Fitzpatrick [243]. Note that where the neutron-only curve from Menendez cannot be seen, it is the same as the Klos structure function. The structure factor from Menendez including only 1-body (1b) currents is not shown as it is the same as the Klos calculation for this range of momentum transfer.

5.1.2 SD WIMP recoil spectrum

The WIMP signal spectrum depends on the WIMP mass and cross section, which are unknown. The rate also depends on the local velocity distribution of WIMPs and the nuclear physics of the interaction. The difference between the SI and SD signal model is the interaction and nuclear physics. Here the calculation of the differential event rate is described, with details given for the various inputs.

In direct detection experiments the measured quantity is the total WIMP-nucleus elastic cross-section, σ :

$$\sigma = \int_0^{q_{max}^2} \frac{d\sigma}{dq^2} dq^2, \quad (5.3)$$

with the maximum momentum transfer given by;

$$q_{max}^2 = 4\mu_A^2 v^2. \quad (5.4)$$

Here, $\mu_A = m_\chi m_A / (m_\chi + m_A)$, is the WIMP-nucleus reduced mass and v the relative WIMP-target velocity. The differential SD WIMP-nucleon cross section is obtained from the Lagrangian (Equation 5.1), and is reproduced here [236]:

$$\frac{d\sigma}{dq^2} = \frac{8G_F^2}{(2J+1)v^2} S_A(q). \quad (5.2)$$

It follows from equations 5.3 and 5.2 that the cross section at zero momentum transfer σ_A is:

$$\sigma_A = \frac{32G_F^2 \mu_A^2}{(2J+1)} S_A(0). \quad (5.5)$$

The standard treatment involves writing the differential WIMP-nucleus cross section in terms of the WIMP-nucleon cross section, to allow the comparison between detectors of different target nuclei and with collider searches. This requires a relationship between σ_A and the WIMP-nucleon cross section. The first step is to rewrite equation 5.2 as:

$$\frac{d\sigma}{dq^2} = \frac{8G_F^2}{(2J+1)v^2} S_A(0) \frac{S_A(q)}{S_A(0)} = \frac{\sigma_A}{4\mu_A^2 v^2} \frac{S_A(q)}{S_A(0)}. \quad (5.6)$$

For the SI case the relationship between the WIMP-nucleon cross section and σ_A is relatively simple. It gives an A^2 enhancement, therefore favoring heavier target nuclei [248]:

$$\sigma_A = A^2 \left(\frac{\mu_A}{\mu_n} \right)^2 \sigma_n^{SI}, \quad (5.7)$$

where μ_n is the WIMP-nucleon reduced mass (assuming equal masses of protons and neutrons) and σ_n^{SI} the SI WIMP-nucleon cross-section.

This enhancement is not present for SD due to the cancellations of the spin between pairs of nucleons, and most of the contribution to the SD sensitivity comes from un-

paired nucleons. $S_A(q)$ can be decomposed into the isoscalar and isovector parts:

$$S_A(q) = a_0^2 S_{00}(q) + a_0 a_1 S_{01}(q) + a_1^2 S_{11}(q), \quad (5.8)$$

with isoscalar and isovector couplings a_0 and a_1 respectively. In order to proceed, two particular choices of these couplings are made ¹:

$$\text{Proton-only:} \quad a_0 = a_1 = 1 \quad (5.9)$$

$$\text{Neutron-only:} \quad a_0 = -a_1 = 1. \quad (5.10)$$

In these cases the structure factor becomes:

$$\begin{aligned} S_p(q) &= S_{00}(q) + S_{01}(q) + S_{11}(q) \\ S_n(q) &= S_{00}(q) - S_{01}(q) + S_{11}(q), \end{aligned} \quad (5.11)$$

for proton-only and neutron-only respectively. The WIMP-neutron/proton couplings are defined as:

$$\begin{aligned} a_0 &= a_p + a_n \\ a_1 &= a_p - a_n. \end{aligned} \quad (5.12)$$

Although using the 2-body currents means that the definition of neutron- or proton-only is no longer strictly true (protons can still be involved in neutron-only and vice-versa), this choice is still useful for setting limits. Later, regions can be excluded in the a_p - a_n plane to remove the model dependence introduced by this choice. For zero momentum transfer we have [236]:

$$S_A(0) = \frac{(2J+1)(J+1)}{4\pi J} |(a_0 + a'_1)\langle S_p \rangle + (a_0 - a'_1)\langle S_n \rangle|^2. \quad (5.13)$$

where $\langle S_{p,n} \rangle$ are the spin expectation values of protons or neutrons in the nucleus; $a'_1 = a_1(1 + \delta a_1(0))$, where the δ term includes the effects of the two-body currents. The couplings are set to either the neutron-only or proton-only cases as noted in footnote 1. Using equations 5.5 and 5.13 the expression for σ_A becomes:

$$\begin{aligned} \sigma_A &= \frac{8G_F^2 \mu_A^2 (J+1)}{\pi J} |(a_0 + a'_1)\langle S_p \rangle + (a_0 - a'_1)\langle S_n \rangle|^2 \\ &= \frac{8G_F^2 \mu_A^2 (J+1)}{\pi J} |\langle S_p \rangle + a'_1 \langle S_p \rangle + \langle S_n \rangle - a'_1 \langle S_n \rangle|^2, \end{aligned} \quad (5.14)$$

¹As with the SI case, we assume that only one interaction occurs, whereas, in reality this need not be the case.

where for the second line, $a_0 = 1$ has been used. It is analogous to do the same as for σ_A , but for the case of an individual proton or neutron outside of the nucleus:

$$\sigma_p = 8G_F^2 \mu_p^2 \frac{(J_{ip} + 1)}{\pi J_{ip}} |\langle S_{ip} \rangle + a'_1 \langle S_{ip} \rangle + \langle S_{in} \rangle - a'_1 \langle S_{in} \rangle|^2. \quad (5.15)$$

The subscripts ip, in refer to the values for an individual proton or neutron outside of the nucleus and the known values can be inserted:

$$\begin{aligned} \sigma_p &= 8G_F^2 \mu_p^2 \frac{(\frac{1}{2} + 1)}{\pi \frac{1}{2}} \left| \frac{1}{2} + \frac{1}{2} a'_1 + \frac{1}{2} - \frac{1}{2} a'_1 \right|^2 \\ &= \frac{8G_F^2 \mu_p^2}{\pi} \times 3. \end{aligned} \quad (5.16)$$

Equations 5.14 and 5.16 give the relationship between the two cross-sections:

$$\sigma_A = \frac{\sigma_p \mu_A^2 (J + 1)}{3\mu_p^2 J} |(a_0 + a'_1) \langle S_p \rangle + (a_0 - a'_1) \langle S_n \rangle|^2. \quad (5.17)$$

It is also equivalent to use equation 5.13 to write:

$$\frac{(J + 1)}{J} |(a_0 + a'_1) \langle S_p \rangle + (a_0 - a'_1) \langle S_n \rangle|^2 = \frac{4\pi S_A(0)}{(2J + 1)}, \quad (5.18)$$

which gives:

$$\sigma_A = \frac{4\pi \mu_A^2}{3\mu_p^2} \frac{S_A(0)}{(2J + 1)} \sigma_p. \quad (5.19)$$

Equation 5.19 allows us to leave only $S_A(q)$ in the equation for the differential cross section. Note that the result is the same for σ_n , as the a'_1 terms still cancel out in equation 5.16. Substituting into equation 5.6 for σ_A gives:

$$\frac{d\sigma}{dq^2} = \frac{\pi \sigma_{p,n}}{3\mu_{p,n}^2 (2J + 1) v^2} S_A(q). \quad (5.20)$$

This is the expression for the differential cross section in terms of either the WIMP-proton or WIMP-nucleon cross section. Note that $S_A(q)$ reduces to either $S_p(q)$ or $S_n(q)$ in this case. Alternatively this can be written as $d\sigma/dE$ using:

$$\begin{aligned} \frac{d\sigma}{dq^2} &= \frac{d\sigma}{dE} \frac{dE}{dq^2} \\ &= \frac{d\sigma}{dE} \frac{d}{dq^2} \left(\frac{q^2}{2m_A} \right) \\ &= \frac{1}{2m_A} \frac{d\sigma}{dE}. \end{aligned} \quad (5.21)$$

The quantity used in the calculation is the differential event rate, dR/dE , where R is the event rate per unit mass. It is related to the differential WIMP-nucleon cross-section by [51]:

$$dR = \frac{\rho_0}{m_\chi m_A} v f(v) \frac{d\sigma}{dq^2} dq^2 dv. \quad (5.22)$$

Then, replacing dq^2 with dE and using equation 5.20 results in :

$$\frac{dR}{dE} = \frac{2\pi\rho_0\sigma_p}{3m_\chi\mu_{p,n}^2(2J+1)} S_A(q) \int \frac{f(v)}{v} dv. \quad (5.23)$$

This rate is per unit detector mass, but only some isotopes are sensitive to the SD interaction. The total differential event rate is a sum of those from each isotope. Naturally, each isotope has different spin structure function. The contribution should be multiplied by a factor:

$$f_A N_T, \quad (5.24)$$

where f_A is the fractional isotopic abundance of isotope with mass number A , and N_T is the number of target atoms per kg. Summing the dR/dE from different isotopes gives:

$$\frac{dR_T}{dE} = \sum_A \frac{2\pi\rho_0 f_A N_T}{3m_\chi\mu_{p,n}^2(2J+1)} \sigma_{p,n} S_A(q) \int \frac{f(v)}{v} dv. \quad (5.25)$$

This is the final equation used in the statistical analysis of the experimental observation for the signal model.

The last part of this equation is the integral over the WIMP velocity distribution in the galaxy, which must be evaluated. This depends on the astrophysical parameters of dark matter. The local dark matter is taken to be a standard Maxwellian velocity distribution, with the velocity of the sun in the galactic rest frame $v_0 = 220$ km/s, a galactic escape velocity of 544 km/s [249] and a local WIMP density of 0.3 GeV/cm³. The average Earth velocity is 245 km/s (during Run 3 data taking). An analytical form of the velocity integral from Ref. [250] is used.

For SI scattering the differential event rate is given by [248]:

$$\frac{dR}{dE} = \frac{\rho\sigma_A |F(q)|^2}{2m_\chi\mu_A} \int \frac{f(v)}{v} dv, \quad (5.26)$$

and the SI WIMP-nucleus cross section σ_A is replaced with the WIMP-nucleon cross section using equation 5.7 (assuming the interaction is completely SI). The standard Helm form factor is used for the SI signals [122]. In Fig. 5.3 the event rates per unit cross section from SI scattering, SD neutron-only, and SD proton-only are compared. It is also useful to compare the recoil spectra. To see the differences the spectra are normalised to an area of one, which is shown in Fig. 5.4. In Fig. 5.5 the recoil spectrum is shown for neutron-only and proton-only for different WIMP masses.

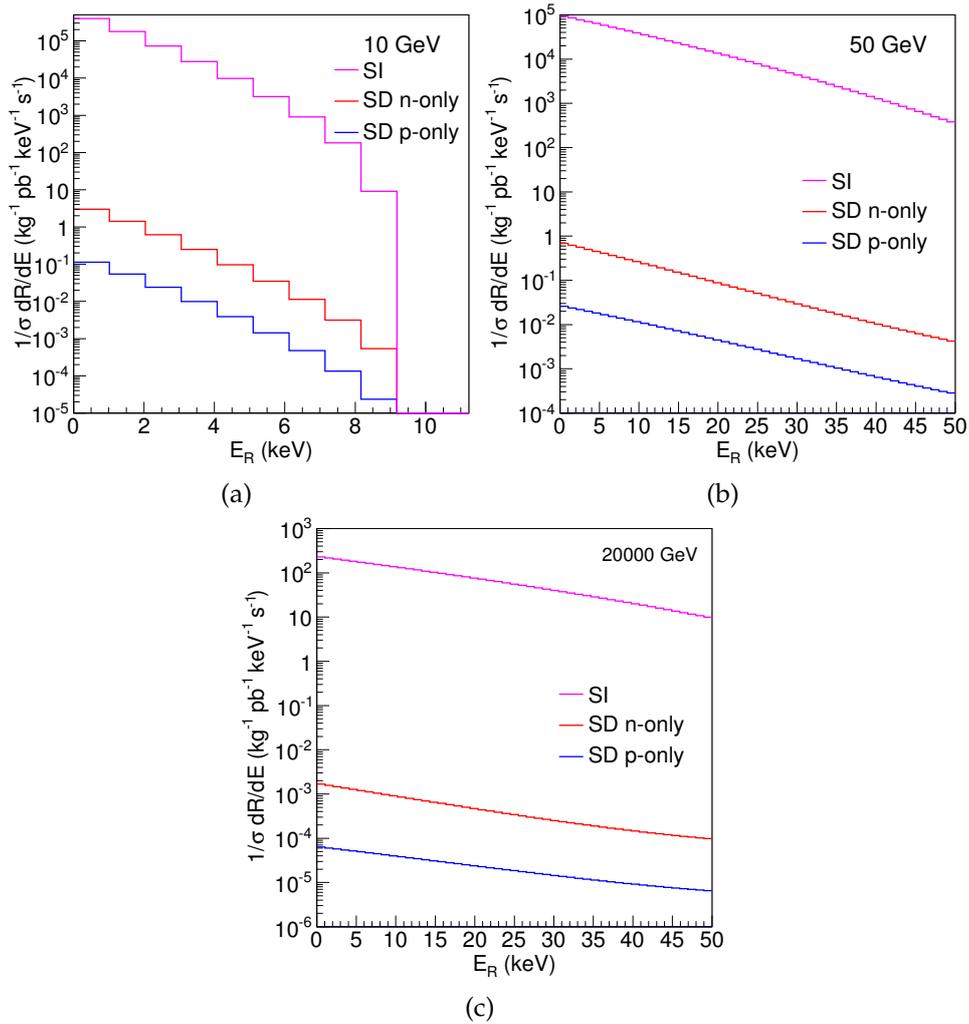


Figure 5.3: Differential event rate per unit cross section for SI and SD WIMP scattering for WIMP masses: (a) 10 GeV, (b) 50 GeV and (c) 20 000 GeV.

The coherent enhancement of the SI scattering in xenon can clearly be seen in Fig. 5.3. The neutron-only SD spectrum is approximately five orders of magnitude below the SI curve, and the proton-only SD curve lies another two below that.

Small differences in the shape can be seen in the normalised spectra. At low masses (such as 5.4a) the SD spectrum has a slightly smaller fraction of events at low recoil energies than SI. The SD proton-only produces a harder recoil spectrum than SI at all masses. The effect is more prominent at higher masses; the fraction of events in the upper half of the spectrum (> 25 keV) at 50 GeV is 11% for proton-only and 8% for SI, and at 20 TeV it is 28% for proton-only and 20% for SI. SD neutron-only is very similar to SI, with a slightly larger fraction of events being towards the centre of the recoil spectrum than SI. The normalised spectra in Fig. 5.5 show the effects of the WIMP mass. For small WIMP masses the spectrum is very peaked at low energy.

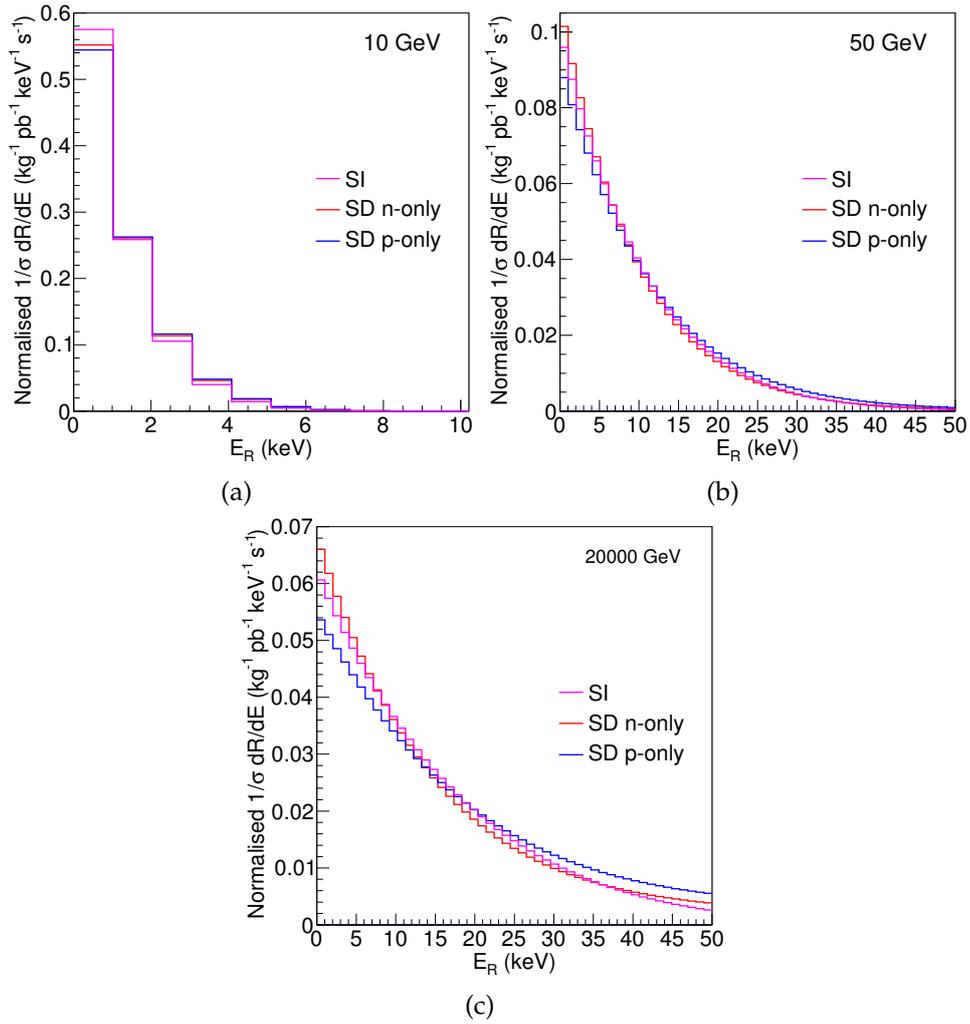


Figure 5.4: Differential event rate per unit cross section for SI and SD WIMP scattering for WIMP masses, with the area normalised to one: (a) 10 GeV, (b) 50 GeV and (c) 20 000 GeV.

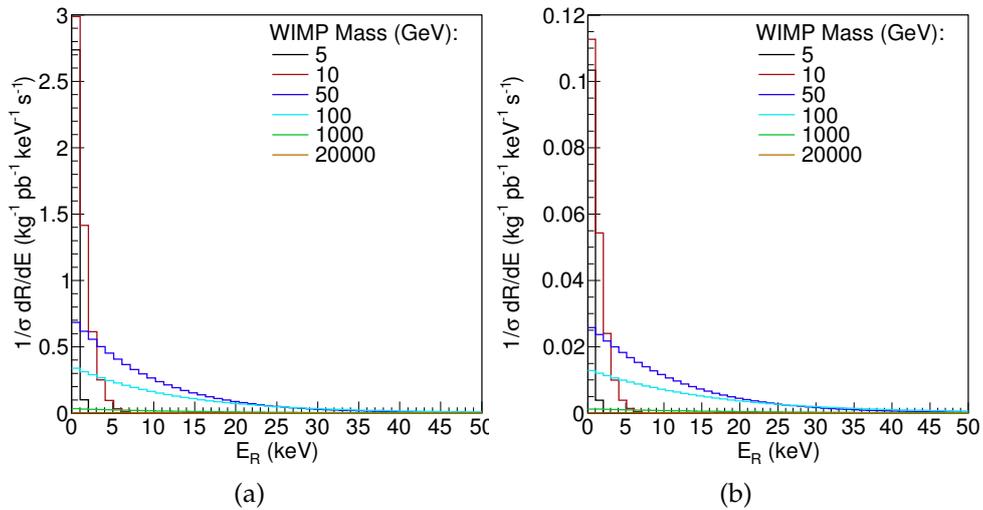


Figure 5.5: Differential event rate per unit cross section for SD WIMPs. (a) Neutron-only (b) Proton-only.

5.2 Signal Estimation

Now that the theoretical signal spectrum has been calculated, we present the statistical treatment to extract that spectrum from a dataset also containing background. In principle, a simple counting experiment could be performed by defining a signal region and seeing if there is a statistically significant excess over the background. Typically, the more powerful Profile Likelihood Ratio (PLR) is used.

The PLR method allows comparison between an observed result (either experimental or simulated data) and a model. It tests whether the model can be rejected given that data. The models are multi-dimensional probability density functions (PDFs), that describe the probability of each outcome (i.e. the observables). Each PDF may depend on several parameters; some are to be measured (such as the cross section in this instance) and are called a parameter of interest (POI), denoted σ . The PDF can also depend on other parameters which may not be of interest. These are called nuisance parameters and allow the inclusion of systematic uncertainties, denoted as a vector θ .

In general there are two models, one is the model to be tested (the alternative hypothesis) and the other is a baseline model to test this against (the null hypothesis). In the case of the PLR for a WIMP search, the hypothesis under test is that there is a WIMP signal with some non-zero cross section (the parameter of interest). The model for this is the signal plus background (S+B). This is compared to the null hypothesis that there is no WIMP signal, a background only (B-only) model. More detail on the construction of the models is given in the following section. By scanning across test values of the parameter of interest for each WIMP mass, cross sections can be tested, and possibly excluded.

In the following section an “event” refers to one sample from the PDF, i.e. obtaining one value for each observable. For example, in the WIMP search case we use four observables (S_1, S_2, r, z) and an event contains a sample from each of these parameters, derived from data or from Monte Carlo. A trial refers to some number of events sampled from a PDF, i.e. a dataset. The likelihood function compares a particular trial with a model. The likelihood is the probability of obtaining our dataset given a specific model. So a high likelihood means that the outcome is very typical for that model, whereas a lower likelihood indicates that outcome is more uncommon. The likelihood is a function of the model parameters σ and θ : $\mathcal{L}(\sigma, \theta)$.

For a trial of many repeated samples, the overall likelihood is calculated from the probabilities of observing each event in the model under test. This is known as the extended unbinned likelihood [251] for a dataset to contain N observed events:

$$\mathcal{L}(\sigma, \theta) = \frac{\mu^N e^{-\mu}}{N!} \prod_{i=1}^N P(x_i). \quad (5.27)$$

Here, x represents the observables from the events and x_i is the outcome for one particular event. The likelihood is the product of the probability of observing each event, multiplied with the Poisson probability to observe that number of events or, equivalently, a particular cross section.

Now it is useful to state the entire form of the profile likelihood ratio and subsequently discuss how it is used:

$$\lambda = \frac{\mathcal{L}(\sigma_t, \hat{\theta})}{\mathcal{L}(\hat{\sigma}, \hat{\theta})}. \quad (5.28)$$

For a particular trial two likelihoods are calculated. The outcome is compared to the model for the null hypothesis (S+B). In the numerator, the parameter of interest is fixed to a particular trial value, σ_t . The nuisance parameters denoted as $\hat{\theta}$ are those which maximise the likelihood of observing that data from the S+B model with σ_t fixed. The denominator is also a maximum likelihood, but σ and θ are not fixed: it is whichever set of parameters that give the highest probability of observing that trial outcome. The model parameters for this global maximum likelihood are denoted $\hat{\sigma}$ and $\hat{\theta}$.

Typically, the test statistic used is not the PLR λ , but instead:

$$q = -2 \ln \lambda(\sigma). \quad (5.29)$$

When the likelihood ratio is one, q is zero. Because the denominator is the global maximum likelihood for that observation, the ratio can never go above one. If the likelihood ratio is smaller then q increases. A large q indicates that there is a model that is more likely to produce the observed result than the model with σ_t fixed. In other words, observing that result from the PDF with σ_t is unlikely. A small q means the model with σ_t is almost as likely to produce the observed result as the best fit model. Equivalently, obtaining that trial outcome from the model with those parameters is quite likely.

The test involves scanning over values of σ_t . For each σ_t the q value of the observed data is determined, which includes the two fits of the model parameters. As there is only one measured dataset (typically), simulated MC datasets must be generated in order to proceed. One is from the S+B model with σ_t . The nuisance parameters have some estimated value and uncertainty. In order to account for this the MC trials profile the nuisance parameters; they change between different trials so that a sample of the possible values and combinations are included. For each MC trial q is evaluated, which requires maximising the likelihood as described above. The other MC data are generated in the same way but from a B-only model ($\sigma = 0$). Three examples of the test statistic distribution are shown in Fig. 5.6 – for a small, intermediate, and large σ_t .

For the S+B model, the numerator of equation 5.28 is the likelihood of observing a MC dataset which was created from a model with σ_t , from a maximum likelihood fit of the S+B model with σ constrained to σ_t . Therefore, most of the time the likelihood in the numerator is going to be similar to that in the denominator and gives many small values of q . Sometimes, if the trial outcome is a fluctuation away from the mean, the numerator maximum likelihood will be reduced. In this case the denominator has no restriction on σ in the fit so the likelihood does not decrease and q is larger. This is true for all σ_t , so the distributions for the S+B do not change much as the test POI is increased.

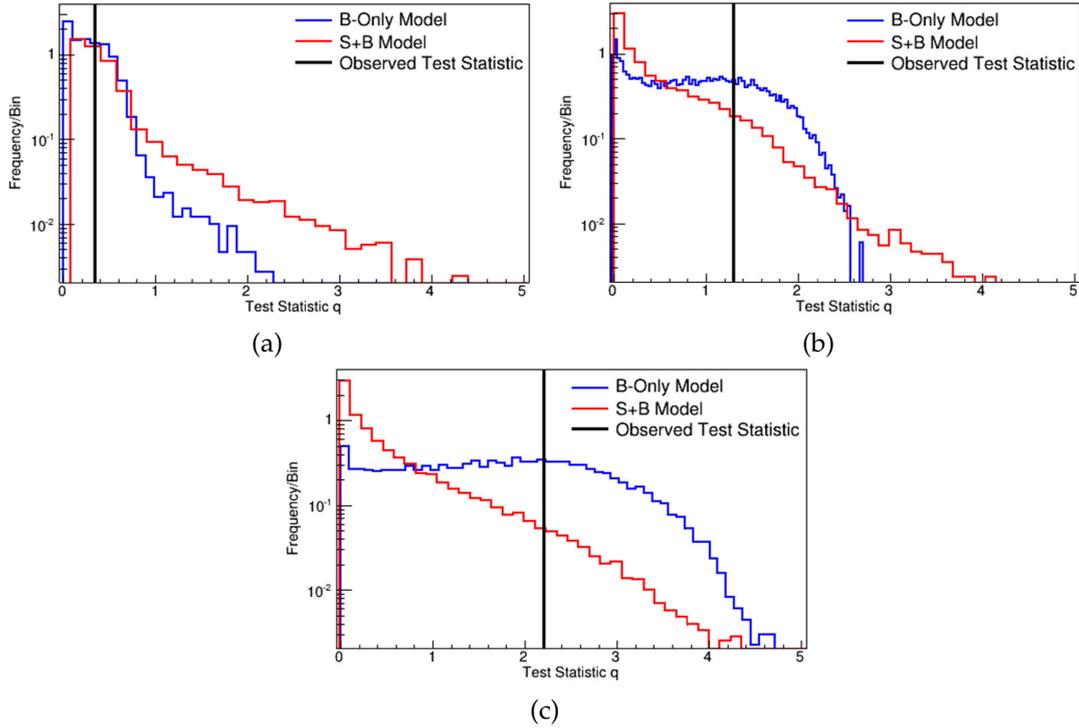


Figure 5.6: Test statistic distributions from the PLR test. The blue histograms are MC trials from a background-only model, red are from a signal-plus-background model. The black line is the observed test statistic from data. The test cross section is increasing from low in (a), to a high in (c).

For the B-only model, the MC is always generated with $\sigma = 0$. The comparison in the numerator of equation 5.28 is the maximum likelihood of observing a MC dataset with $\sigma = 0$ to the S+B model with σ_t . For a small σ_t , there is still a reasonable probability of observing the B-only MC trials. As σ_t increases, the probability of observing those trials from the B-only model becomes smaller. In the denominator because σ is not constrained it can always give a high likelihood of observing the B-only trials, so the q values for the B-only trials become larger as σ_t increases.

These distributions are used to determine confidence intervals on the POI in the following way. A p-value is defined as the probability for a given q and σ_t , that if the experiment was repeated a result equally or less probable would be obtained, assuming the S+B model under test is true. The smaller this is, the more certain we can be that the model being tested is false. A large p-value indicates that there is no strong evidence against the model. In terms of Figure 5.6, the p-value is the integral of the S+B histogram from a point q to infinity. In Figure 5.6a, the observed q is in the middle of the S+B distribution, so the probability of observing q from the S+B model is clearly high, and consequently the assumed cross section cannot be excluded. Looking at Figure 5.6c, the q for the observed data is now in the tail of the S+B distribution. The probability of observing that value from the S+B model under test is low, so that cross section can be excluded. We exclude a model at 90% CL (with a particular cross section) when the p-value is less than 10%, i.e. when there at most a 10% chance that the model under test could give that observed result. However, affirming that the S+B

model is not excluded is not stating that it is true. As can be seen in Figure 5.6a, the B-only model is also quite likely to give that very same observation.

Going to higher cross section means that signal and background are easier to distinguish. A typical plot of the p-value for a particular WIMP mass is shown in Figure 5.7, where the solid black line is that from the observed data.

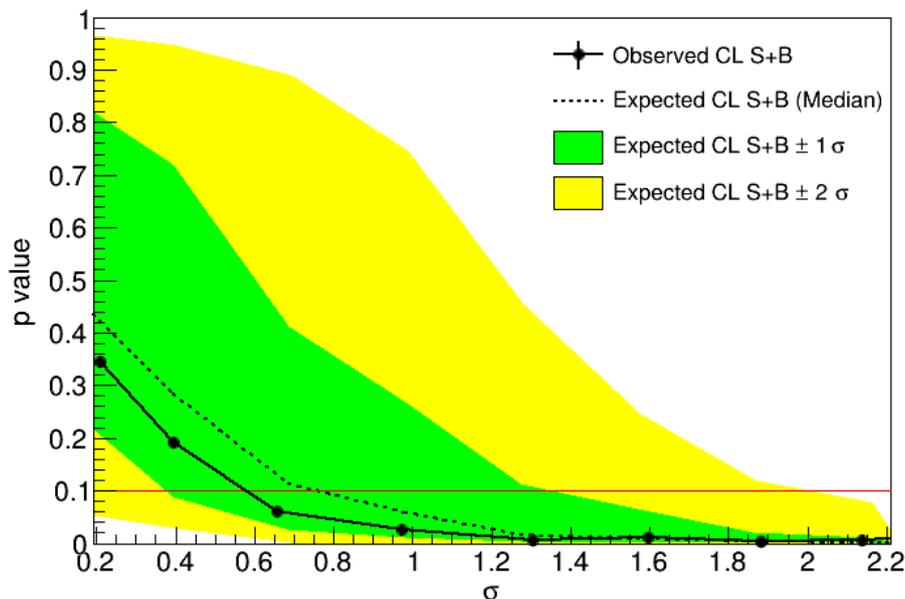


Figure 5.7: p-values from a hypothesis scan in the PLR. The solid black line is the observed p-value, with the dotted line being that expected from the median of the background-only model. The 1σ and 2σ bands on the expected p-value are shown in green and yellow, respectively. The red line shows the 10% point, below which a cross section can be excluded at 90% CL.

The bands in Fig. 5.7 are also determined from the plots like those in Fig. 5.6. The mean is the expected result if the background-only hypothesis was true, corresponding to the p-value for the median of the background-only distribution (integrating the S+B histogram from the median of the B-only histogram). The upper and lower 1σ bands are defined by the points where 84% and 16% of the B-only trials have a larger q . Similarly for 2σ it is 97.7% and 2.3%.

In the LUX analysis, the PLR is implemented using RooStats [252] and RooFit [253], with a modified version of the hypothesis test inverter.

5.3 Models for the LUX PLR

The models are input to the PLR in the form of PDFs for both signal and background. Four observables are used in the Run 3 PLR: S_1 , S_2 , height, and radius, a typical set for WIMP searches in liquid xenon TPCs. A PDF for each observable is required.

For both the signal and the background model the method is similar. Firstly, the energy spectrum of the events must be found (ie. dR/dE) which describes the energy

deposits in the liquid xenon. For the signal model, the WIMP recoil spectrum was described in section 5.1.2, with the final result given in equation 5.25. In the LUX implementation of the PLR the recoil spectrum is stored as a histogram of dR/dE per unit cross section. The various background sources are distributed differently in the detector, and these are independently simulated to produce the energy deposits in the liquid xenon. For both the signal and the background model the detector response to the energy deposits is simulated to produce PDFs in the required observables.

5.3.1 Backgrounds and constraints

Unlike the signal model, the background model has several distinct components, dominated by the leakage of electron recoils into the signal region, as described in section 4.5.1. This section describes updates for the Run 3 reanalysis. The background model is common to both the SI and SD results. Background rates are estimated by screening of detector materials, gas sampling results, or directly from data sidebands. In the PLR, each background is given a constraint based on these results. When fitting to the data these are allowed to float within their constraints. The final constraints on each background are reported from Ref. [148].

The ^{85}Kr and ^{222}Rn induced electron recoil backgrounds were described previously. These are uniformly dispersed in the xenon and for the PLR these β^- backgrounds are grouped together with a total of 55 ± 22 predicted events for the 96-day search data. The cosmogenically produced xenon isotope ^{127}Xe is included in the background model with a total of 91 ± 27 events. The ^{127}Xe decay is usually accompanied by a γ -ray allowing the events to be vetoed. It can only contribute to the background if the γ -ray escapes the detector (see section 4.5.1). Therefore, ^{127}Xe backgrounds are concentrated near to the detector walls.

Since the first Run 3 analysis, an excess of ER events was identified at the bottom of the detector, that was not accounted for by the existing screening results [148]. These were salient between 500 and 1500 keV and uniform in x - y . The exact origin is not known with certainty, but the energy spectrum could be reproduced. This required a combination of heavily downscattered ^{238}U -chain, ^{232}Th -chain, and ^{60}Co gamma-rays from below the PMTs, originating within the large copper γ shield. There are two scaling parameters in the PLR for the γ -ray background, one for materials at the bottom of the detector (including the new source), and one for the rest of the γ -rays. The bottom γ -ray activity is estimated to be 172 ± 74 events (105 of which are attributed to the unidentified source) and the remainder is 247 ± 106 . This background will be mostly near to the walls, with some events reaching the fiducial volume.

In addition to the extra γ -ray background, a weak, low energy peak at 2.8 keV was identified due to the significantly improved energy resolution [148]. This energy is consistent with electron capture decays by ^{37}Ar in the fiducial volume. In the PLR this is an unconstrained fit parameter on the number of decays. Measurements of the ^{37}Ar concentration in the lab air and sampling measurements are planned in the future to give an upper limit on this background.

The last addition for the reanalysis was modelling of the “wall backgrounds”. These are ER and NR events originating from the walls of the TPC that can leak into the signal region. In the original analysis we opted to reduce the fiducial radius to eliminate this spurious population. For the reanalysis, a data-driven method was employed to model this background [254]. In general, events near to the wall can undergo charge loss, resulting in a reduced S2 signal. There are a variety of sources, many being from decay products of the ^{222}Rn chain. The α decay of ^{210}Pb was found to be the most significant contribution to the WIMP search background. It mimics a low energy NR as the two decay products ($\alpha + ^{206}\text{Pb}$) are emitted back-to-back, but one is not visible as it goes into the PTFE wall. Other backgrounds come from ER events that undergo charge loss, lowering them into the NR band. These events can leak inwards to the fiducial volume due to misreconstruction of the position, which is more likely for very small S2s. The model is derived in two parts, one is to select events that must come from the walls, ^{210}Po and ^{206}Pb α decays where the charge loss was small. These are used to parameterise the radial distribution of the wall background events, in bins of S2 (as the distribution is S2 dependent). The other observables, S1, S2 and height are not independent of each other, so are modelled non-parametrically from the data using a multi-dimensional Kernel Density Estimate PDF of a sample of wall events. The selected events are those which are reconstructed beyond the walls of the detector, which only has a small contamination from regular ERs. The estimated background from this model for the reanalysis was 24 ± 7 events.

The total expected neutron background is 0.08 ± 0.01 NR events, and coherent neutrino-nucleus scattering from ^8B solar neutrinos contributes 0.10 events. Random coincidences of isolated S1 and S2s adds another 1.1 events. These backgrounds are small compared to the total predicted ER background of 565 events and are not included in the background model.

5.3.2 Detector response simulations

Up until now, for the signal model we have the differential recoil spectrum of events and for the background model the activity of each source. The next step is to progress from these to PDFs in the four observables, S1, S2 radius, and height.

For the backgrounds, PDFs for each source are required. For the wall event model, the PDF is already obtained directly from data sidebands. For the others the PDF comes from Monte Carlo simulations. The first step is to simulate energy deposits in the xenon, arising from the various backgrounds. This is done using LUXSim [179], which incorporates Geant4 [209] to simulate radioactive decays within detector components or within the xenon itself. For the signal model, the energy deposits are calculated using equation 5.25, for a particular WIMP mass and cross section.

After obtaining the energy deposits in the xenon, the S1 and S2 signals must be generated. This is done using NEST [179], and there are differences in the implementation for the signal and background. In general, NEST is based on a physical model and tuned to calibration data for both NRs and ERs. For ER, the relationship between the

energy, S1 and S2 is:

$$E = W \left(\frac{S1_c}{g_1} + \frac{S2_c}{g_2} \right), \quad (5.30)$$

where $W = 13.7$ eV is the average energy required to create one quantum (photon or electron) in the liquid xenon, g_1 is the S1 photon detection efficiency and g_2 is the product of the electron extraction efficiency and the mean single electron response. Equation 5.30 can be rewritten as:

$$\frac{S1}{E} = \frac{g_1}{W} - \frac{S2}{E} \frac{g_1}{g_2}. \quad (5.31)$$

This describes a straight line in a plot of S1/E against S2/E, known as a ‘‘Doke plot’’ [172]. The y intercept and gradient are:

$$\begin{aligned} b &= \frac{g_1}{W}, \\ m &= \frac{g_1}{g_2} = \frac{bW}{g_2}. \end{aligned} \quad (5.32)$$

Each monoenergetic calibration source or background ER peaks can be plotted to determine g_1 and g_2 . In LUX, ^{82}Kr (41.55 keV), ^{131}Xe (163 keV), ^{127}Xe (207, 237 and 410 keV), ^{129}Xe (237 keV), ^{214}Bi (609 keV), and ^{137}Cs (662 keV) were used to produce the Doke plot. This will be presented in Ref. [255]. The resulting g_1 and g_2 are presented in Table 5.1.

In addition, a low-energy ER calibration was performed with the tritium source, described in section 3.2.8, dispersed throughout the xenon [178]. This is a β^- decay with a 18.6 keV end point generating strictly single-site interactions in the liquid xenon. In order to test the validity of the ER simulations, a full simulation of the tritium spectrum was produced, see Fig. 4.19. The S1 and S2 spectra are tuned to the observed data, and the obtained g_1 and g_2 were consistent within errors to those obtained from line sources via the Doke plot.

With the calibration against the ER data, full simulations were carried out for all of the background sources. This involved producing the S1 and S2 signals for the energy deposits. Complete simulated events were created, including the waveforms. This allows the simulated data to be processed and analysed in the same way as experimental data. With this the PDFs for each background in each observable can be obtained.

For the signal PDF, the NR response was calibrated using a DD neutron generator [256]. The model used for NRs is as follows:

$$\begin{aligned} n_e &= \mathcal{L}(E) \frac{E}{W} \left(\frac{1}{1 + N_{ex}/N_i} \right) (1 - r) + \frac{C n_{ph}}{f_l} (1 - f_l) \\ n_{ph} &= \left(\mathcal{L}(E) \frac{E}{W} - n_e \right) f_l, \end{aligned} \quad (5.33)$$

which gives the number of electrons (for S2) and photons (for S1); E is the energy deposited in the medium, N_{ex} and N_i are the number of electron excitations and ionisations respectively; r is the Thomas-Imel box (TIB) model recombination factor, which describes the fraction of ionisation electrons that recombine with Xe^+ ions, and so instead of producing an electron they produce a photon [257]. f_l is the biexcitonic quenching parameter, which is the suppression of photons as two excitons collide and de-excite via the process $Xe^* + Xe^* \rightarrow Xe + Xe^+ + e^-$ [258]. The electron can recombine with an ion and results in one photon being produced from two excitons. Some of this is regained as charge from Penning effects, where the collision produces an electron, which is represented by the term C . The Lindhart factor \mathcal{L} represents the fraction of energy that goes to atomic electrons, with the rest lost as atomic motion.

The parameter r depends on N_i and a field-dependent parameter ζ [257]:

$$r = 1 - \frac{\ln(1 + N_i \zeta)}{N_i \zeta}. \quad (5.34)$$

For L , there are two possible forms, using either the Lindhart model [259] [175] or Bezrukov model [260]:

$$\begin{aligned} \text{Lindhart: } \mathcal{L}(E) &= \frac{kg(\epsilon)}{1 + kg(\epsilon)}, \\ \text{Bezrukov: } \mathcal{L}(E) &= \alpha \frac{s_e}{s_e + s_n}. \end{aligned} \quad (5.35)$$

For the Lindhart model, k is a proportionality constant between the electronic stopping power and the velocity of the recoiling nucleus. The value of k has been calculated as either 0.17 [259] or 0.11 [258] with experimental results lying between the two [175]. The parameter $g(\epsilon)$ is proportional to the ratio of electronic to nuclear stopping power and it is parameterised as a function of the quantity ϵ [122]:

$$\begin{aligned} \epsilon &= 11.5(E/\text{keV})Z^{-7/3}, \\ g(\epsilon) &= 3\epsilon^{0.15} + 0.7\epsilon^{0.6} + \epsilon. \end{aligned} \quad (5.36)$$

The Bezrukov model depends on the nuclear and electronic stopping power, s_n and s_e . The constant α accounts for the cascade of collisions generated from a single nuclear recoil. The nuclear and electronic stopping powers are [260]:

$$\begin{aligned} s_n(\epsilon Z) &= \frac{\ln(1 + 1.1383\epsilon Z)}{2(\epsilon Z + 0.01321\epsilon_Z^{0.21226} + 0.19593\epsilon_Z^{0.5})}, \\ s_e &= 0.166\epsilon^{0.5} \end{aligned} \quad (5.37)$$

where ϵ is defined by equation 5.36 and $\epsilon_Z = 1.068\epsilon$. Finally, f_l for the biexcitonic quenching is parameterised by [261]:

$$f_l = \frac{1}{1 + As_e}. \quad (5.38)$$

In these equations there are only five free parameters in the model: the ratio N_{ex}/N_i , the TIB parameter ζ , the biexcitonic quenching factors C and A , and either k or α depending on whether the Lindhart or Bezrukov model is used for $\mathcal{L}(E)$.

The DD calibration data are used to determine the free parameters of the model. These calibration data, derived from a mono-energetic and collimated neutron beam, allow the recoil energy to be reconstructed from the kinematics of multiple scatters [256]. Therefore, it provides an independent measurement of the recoil energy that does not rely on the size of the charge or light signals, and can be used to measure their yields as a function of energy down to 1.1 keV. A fit was carried out simultaneously on the charge and light yield using Markov Chain Monte Carlo. The data were found to be consistent with both the Lindhart and the Bezrukov $\mathcal{L}(E)$, with Lindhart giving the best goodness of fit. The model parameters and their errors are shown in Table 5.1 [255].

Table 5.1: Best fit nuisance parameters. The values of g_1 and g_2 are determined from monoenergetic calibration sources. The other parameters are from the fit of the NR model to data; values are from Ref. [255].

Parameter	Best fit value
g_1	0.117 ± 0.003
g_2	0.095 ± 0.006
N_{ex}/N_i	0.48 ± 0.07
TIB ζ	0.067 ± 0.008
BEQ A	13.2 ± 2.3
Lindhart k	0.17 ± 0.06
BEQ C	0.11 ± 0.06

For the signal model the PDFs are generated using FastNEST. This is a parameterised version of the full NEST simulations, which only produces the number of S1 and S2 photons (not the waveforms). FastNEST uses the NR model and parameters described above. This is to allow variation of the nuisance parameters, which is described in the next section. Some examples of the signal PDF at different WIMP masses are given in Fig. 5.8.

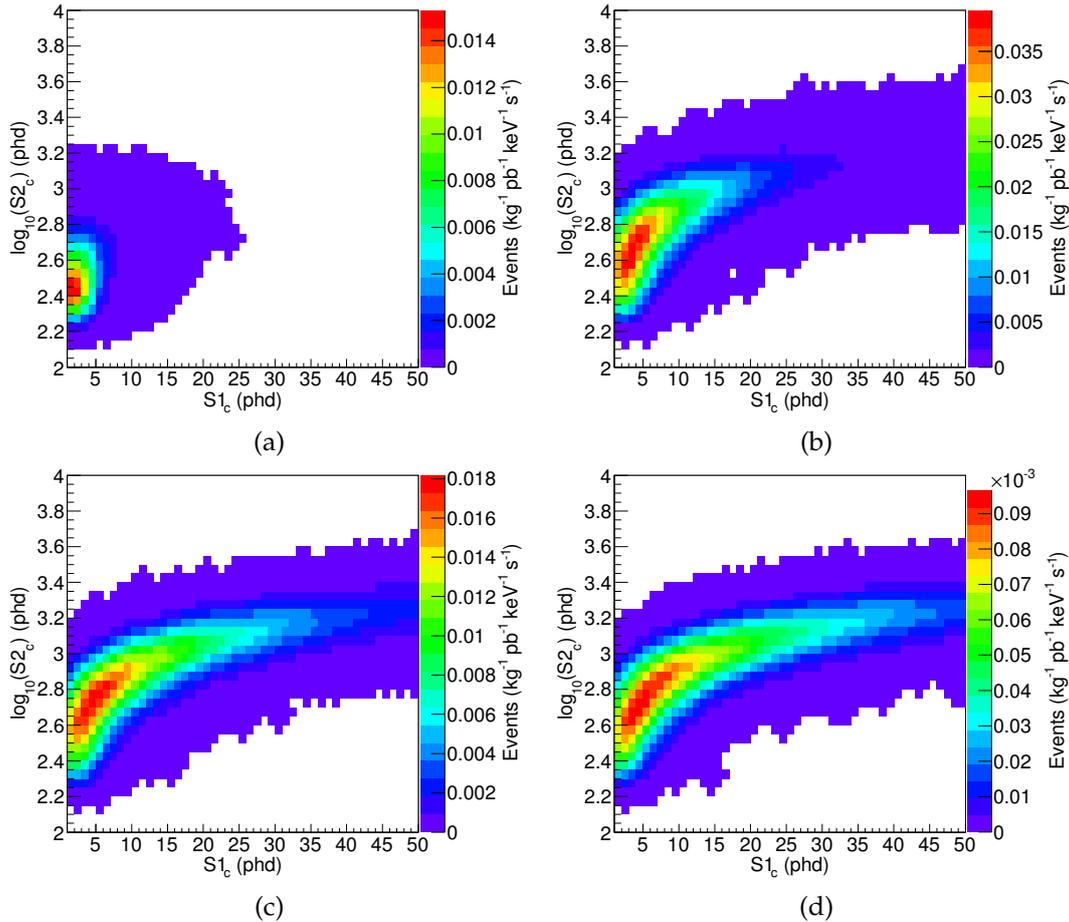


Figure 5.8: Signal PDFs for neutron-only SD WIMP interactions in $\log_{10}(S2)$ against $S1$ (within the 1 to 50 phd WIMP search region). The $S1$ and $S2$ sizes are given as phd, the estimated number of detected photons, which is described in section 5.5. WIMP mass is (a) 10 GeV, (b) 33 GeV, (c) 100 GeV, and (d) 20 000 GeV.

5.4 Implementation of Nuisance Parameters in the PLR

Nuisance parameters are uncertain parameters in the model that affect the outcome, but are not the main aim of the measurement. In the background model, the number of events of each of the five types of background PDF is considered a nuisance parameter. There are the charge and light yields g_1 and g_2 , along with the five parameters used in the NEST model for the signal: N_{ex}/N_i , ζ , C , A , and k (for Lindhart model). There are two extra possible nuisance parameters in the SD analysis over the SI due to uncertainties in the spin structure function of ^{129}Xe and ^{131}Xe .

For the background, the PDFs are normalised. When fitting the background to a given dataset the total number of events from each background is allowed to float within their constraints from the measurements described previously, although the model is static (the shape of the PDF is not allowed to change). The PDFs are determined from full LUXSim simulations as mentioned above.

For the signal model, in order to include the parameters from the model as nuisance parameters in the PLR, they must be allowed to change in the fit, within their errors from the fit to the DD calibration data. This is why end-to-end simulations are not used and FastNEST is incorporated into the analysis instead. There are two ways in which signal model nuisance parameters can be included in the LUX PLR. The first way is used when it affects the shape of the PDF. Each time the value of a nuisance parameter is changed between MC trials, the signal PDF is regenerated which takes approximately one minute of CPU time. For the case of the uncertainty in the spin structure functions, this affects the WIMP recoil spectrum. So, if those change both the recoil spectrum and the PDF in the observables have to be regenerated. Having multiple nuisance parameters greatly increases the parameter space to scan over when running the MC trails, so the PDF must be regenerated more often. Alternatively, if the nuisance parameter only significantly affects the overall number of events, it can be parameterised which avoids having to regenerate the PDF during the calculation. Before running the statistical test, FastNEST produces the number of events per unit cross section as a function of the nuisance parameter.

With this method, a nuisance parameter can be used to either modify the shape or as a scaling factor. Also, it is not possible to include all nuisance parameters in the shape, because it greatly increases the computation time. Therefore, the effects of the nuisance parameters must be quantified to determine which ones can be neglected and, for the significant ones, whether they affect the shape or scaling of the PDF.

The effects of different nuisance parameters on the signal model are compared using the PDFs in S1 and S2. First, a number of signal PDFs are generated with the nuisance parameters at their nominal value for a range of WIMP masses. Then, each nuisance parameter in turn is changed to its $\pm 1\sigma$ value, and another sample of signal PDFs is generated. The mean PDF from each set of trials is compared. The number of events is compared to the PDFs with nominal nuisance parameter values to see if there is an effect:

$$\Delta N = N_v / N_n. \quad (5.39)$$

N_v is the number of events with the nuisance parameter at $\pm 1\sigma$ and N_n is with the nominal value. To assess the impact on the shape of the PDF a χ^2 test can be used to quantify the differences in each bin:

$$\chi^2 = \sum_i \frac{(N_{ni} - N_{vi})^2}{\sigma_{vi}}. \quad (5.40)$$

N_{ni} is the number of events in the nominal histogram for bin i , and N_{vi} is the same for the histogram with a varied nuisance parameter. σ_{vi} is the standard deviation of bin i in the varied histogram. Repeated trials are necessary to determine σ_i .

Example plots of the varied histograms subtracted from the nominal histogram are shown for the g_2 and TIB model parameter variation in Figs. 5.9 and 5.10. A total of 50 PDFs were produced and averaged with both the nominal and varied nuisance parameter.

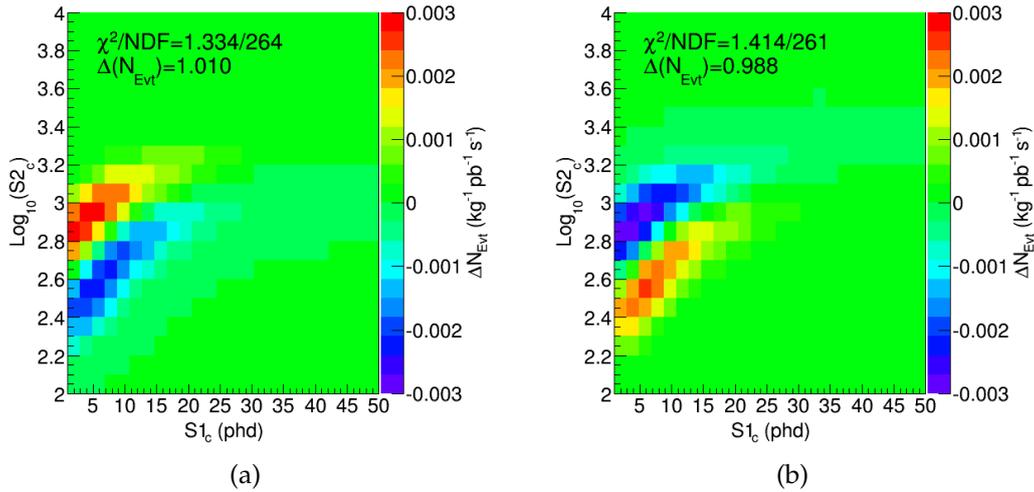


Figure 5.9: Difference between the $\log_{10}(S2)$ vs $S1$ histogram with g_2 increased (a) or decreased (b) by 1σ , for a 33 GeV WIMP mass in the 1–50 phd WIMP search region. The histogram with the nominal nuisance parameter value is subtracted from the one with the increased or decreased nuisance parameter.

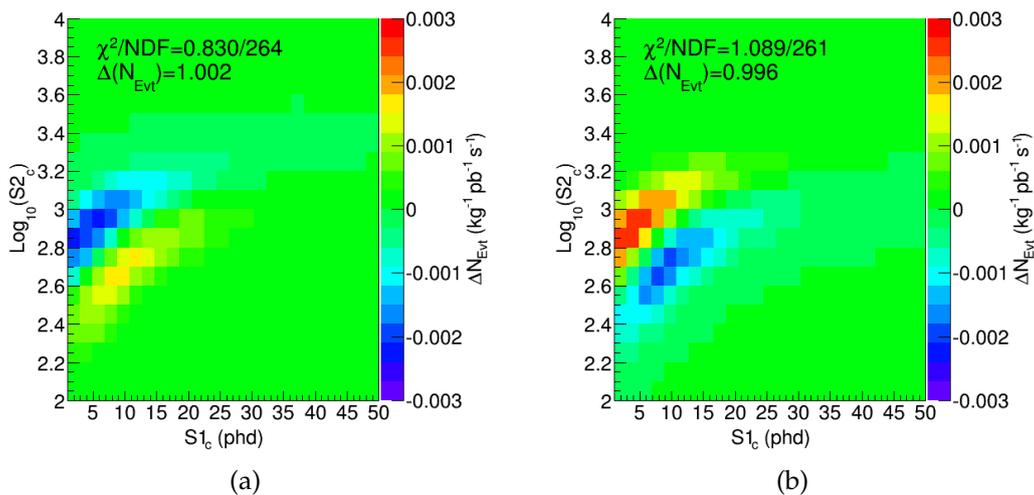


Figure 5.10: Difference between the $\log_{10}(S2)$ vs $S1$ histogram with the TIB parameter increased (a) or decreased (b) by 1σ , for a 33 GeV WIMP mass in the 1–50 phd WIMP search region. The histogram with the nominal nuisance parameter value is subtracted from the one with the increased or decreased nuisance parameter.

Table 5.2: Effect of increasing or decreasing nuisance parameters to their $\pm 1\sigma$ values on the number of events, ΔN_{Eff} , for SD neutron-only interactions.

m_χ (GeV)	5	10	17	33	50	100	1000	20000
No change	1.001	0.998	1.000	1.000	1.000	1.000	1.001	1.004
	0.994	1.000	0.999	1.000	1.000	0.999	1.000	1.000
Biexcitonic C	1.024	1.016	1.007	1.007	1.001	1.000	1.001	1.002
	1.004	0.996	0.997	0.999	0.999	0.999	0.998	1.000
Biexcitonic A	0.985	0.985	0.993	0.997	0.998	1.005	1.011	1.012
	1.024	1.013	1.010	1.004	1.000	0.993	0.987	0.988
g_1	1.034	1.020	1.017	1.007	1.001	0.998	0.995	0.997
	0.978	1.016	0.988	0.994	0.997	1.000	1.002	1.003
g_2	1.281	1.068	1.022	1.010	1.006	1.004	1.004	1.004
	0.758	0.923	0.970	0.988	0.995	0.994	0.993	0.995
Lindhart k	1.162	1.059	1.027	1.012	1.007	1.000	0.998	0.999
	0.876	0.938	0.972	0.990	0.993	0.998	1.002	1.001
N_{ex}/N_i	0.970	1.002	1.001	1.003	0.999	1.000	0.999	1.000
	1.042	1.017	0.993	0.995	0.997	0.999	0.997	0.999
TIB	1.003	1.010	1.003	1.002	1.000	0.997	0.998	0.997
	1.017	0.994	1.003	0.996	0.999	1.000	1.001	1.003
$^{129}\text{Xe } S_A$	1.087	1.073	1.074	1.073	1.073	1.070	1.070	1.073
	0.934	0.922	0.925	0.927	0.928	0.929	0.932	0.930
$^{131}\text{Xe } S_A$	1.036	1.025	1.022	1.022	1.021	1.020	1.023	1.023
	0.983	0.977	0.980	0.977	0.977	0.977	0.976	0.977

Table 5.3: Effect of increasing or decreasing nuisance parameters to their $\pm 1\sigma$ values on the number of events, ΔN_{Est} , for SD proton-only interactions.

m_χ (GeV)	5	10	17	33	50	100	1000	20000
No change	0.995	0.997	0.999	0.999	1.000	1.000	1.000	1.000
	1.002	1.001	1.001	1.000	1.000	1.000	1.000	1.000
Biexcitonic C	1.009	1.010	1.004	1.002	1.006	1.001	1.001	1.001
	0.996	0.998	0.988	0.999	0.999	1.000	1.000	1.001
Biexcitonic A	0.979	0.985	0.990	0.997	1.003	1.012	1.018	1.018
	1.027	1.017	1.001	1.003	0.995	0.987	0.981	0.981
g_1	1.040	1.022	1.010	1.005	1.001	0.996	0.993	0.993
	0.969	0.979	0.979	0.995	0.999	1.004	1.008	1.011
g_2	1.266	1.065	1.023	1.008	1.006	1.005	1.003	1.003
	0.754	0.927	0.967	0.990	0.992	0.995	0.996	0.996
Lindhart k	1.139	1.056	1.026	1.009	1.004	0.998	0.995	0.995
	0.876	0.943	0.964	0.989	0.994	1.002	1.006	1.008
N_{ex}/N_i	0.955	0.996	1.002	1.003	1.001	1.002	1.001	1.000
	1.025	0.996	0.984	0.996	0.997	0.999	1.000	1.001
TIB	0.005	1.003	1.004	1.002	1.000	0.998	0.996	0.996
	1.014	0.992	0.986	0.998	0.999	1.003	1.005	1.007
$^{129}\text{Xe } S_A$	1.443	1.468	1.404	1.377	1.362	1.348	1.339	1.338
	0.564	0.581	0.591	0.623	0.638	0.653	0.661	0.662
$^{131}\text{Xe } S_A$	1.115	1.125	1.129	1.129	1.125	1.125	1.125	1.124
	0.876	0.876	0.865	0.873	0.874	0.875	0.875	0.876

Table 5.4: Effect of increasing or decreasing nuisance parameters to their $\pm 1\sigma$ values on the shape of the PDF for neutron-only. The numbers are the χ^2 values as defined by equation 5.40. The degrees of freedom (DOF) represent the number of non-zero bins in the histograms (such as Fig. 5.9), and are an average value between the $\pm 1\sigma$ results.

m_χ (GeV)	5	10	17	33	50	100	1000	20000
D.O.F.	46	74	128	263	286	295	297	295
No change	0.001	0.000	0.005	0.000	0.001	0.001	0.000	0.001
Biexcitonic C +	0.001	0.025	0.043	0.14*2	0.278	0.384	0.478	0.467
Biexcitonic C -	0.002	0.019	0.047	0.171	0.303	0.429	0.483	0.482
Biexcitonic A +	0.000	0.027	0.076	0.172	0.251	0.279	0.329	0.305
Biexcitonic A -	0.001	0.020	0.105	0.192	0.260	0.306	0.314	0.319
g_1 +	0.002	0.019	0.066	0.080	0.071	0.074	0.072	0.084
g_1 -	0.003	0.027	0.052	0.078	0.081	0.087	0.105	0.088
g_2 +	0.127	0.785	1.033	1.334	1.540	1.527	1.692	1.595
g_2 -	0.160	0.740	1.124	1.414	1.701	1.750	1.740	1.795
Lindhart k +	0.010	0.040	0.050	0.039	0.024	0.021	0.011	0.015
Lindhart k -	0.014	0.036	0.067	0.043	0.030	0.019	0.014	0.015
N_{ex}/N_i +	0.037	0.304	0.401	0.436	0.397	0.380	0.398	0.384
N_{ex}/N_i -	0.024	0.272	0.569	0.510	0.481	0.437	0.493	0.449
TIB +	0.010	0.151	0.410	0.830	1.058	1.127	1.290	1.249
TIB -	0.002	0.155	0.464	1.089	1.282	1.488	1.638	1.559
$^{129}\text{Xe } S_A$ +	0.001	0.002	0.009	0.001	0.003	0.002	0.003	0.007
$^{129}\text{Xe } S_A$ -	0.002	0.001	0.001	0.001	0.002	0.003	0.002	0.004
$^{131}\text{Xe } S_A$ +	0.001	0.001	0.008	0.001	0.005	0.003	0.001	0.001
$^{131}\text{Xe } S_A$ -	0.003	0.001	0.001	0.002	0.001	0.001	0.006	0.001

Table 5.5: Effect of increasing or decreasing nuisance parameters to their $\pm 1\sigma$ values on the shape of the PDF for proton-only. The numbers are the χ^2 values as defined by equation 5.40. The degrees of freedom (DOF) represent the number of non-zero bins in the histograms (such as Fig. 5.9), and are an average value between the $\pm 1\sigma$ results.

m_χ (GeV)	5	10	17	33	50	100	1000	20000
D.O.F.	47	77	129	278	289	296	300	299
No change	0.000	0.000	0.012	0.001	0.001	0.000	0.001	0.000
Biexcitonic C +	0.001	0.025	0.054	0.196	0.311	0.497	0.527	0.560
Biexcitonic C -	0.000	0.019	0.039	0.204	0.358	0.507	0.607	0.633
Biexcitonic A +	0.001	0.017	0.088	0.210	0.267	0.328	0.334	0.361
Biexcitonic A -	0.000	0.018	0.052	0.216	0.271	0.329	0.378	0.384
g_1 +	0.001	0.018	0.055	0.073	0.079	0.077	0.081	0.080
g_1 -	0.001	0.021	0.106	0.084	0.102	0.078	0.078	0.077
g_2 +	0.121	0.683	1.050	1.396	1.548	1.671	1.632	1.661
g_2 -	0.165	0.741	1.090	1.521	1.681	1.802	1.846	1.875
Lindhart k +	0.008	0.042	0.054	0.038	0.024	0.014	0.010	0.009
Lindhart k -	0.015	0.037	0.064	0.043	0.034	0.013	0.009	0.012
N_{ex}/N_i +	0.030	0.302	0.421	0.445	0.414	0.418	0.420	0.370
N_{ex}/N_i -	0.031	0.392	0.650	0.488	0.506	0.403	0.401	0.428
TIB +	0.005	0.149	0.445	0.925	1.189	1.287	1.331	1.349
TIB -	0.003	0.183	0.703	1.115	1.473	1.552	1.621	1.716
$^{129}\text{Xe } S_A$ +	0.000	0.016	0.006	0.044	0.074	0.105	0.122	0.119
$^{129}\text{Xe } S_A$ -	0.000	0.003	0.046	0.205	0.333	0.440	0.492	0.479
$^{131}\text{Xe } S_A$ +	0.001	0.000	0.010	0.003	0.001	0.002	0.002	0.002
$^{131}\text{Xe } S_A$ -	0.000	0.002	0.016	0.002	0.003	0.003	0.003	0.003

For the shape of the S2 vs S1 distribution, the most significant nuisance parameter is g_2 . This effect can be seen in Fig. 5.9; an increase in g_2 moves events upwards and vice versa. The TIB parameter is also significant, but slightly less so than g_2 . Ideally, this would be included too, but with two shape nuisance parameters the computation time is too long. For the total number of events, in neutron-only the Lindhart k had the largest effect at small masses, and the ^{129}Xe structure function was the most important at higher masses. The Lindhart k was also the largest effect for the SI calculation. However, for proton-only the nuisance parameter for both the ^{129}Xe and the ^{131}Xe structure functions were the most significant. ^{131}Xe has a similar effect to the Lindhart k and ^{129}Xe is much larger. For the PLR calculation g_2 was included as a nuisance parameter in the shape and the Lindhart k factor for the normalisation, the same as those for SI. The structure function nuisance parameters were not included for the SD calculation.

It is important to note that the numbers listed above are for when a nuisance parameter is fixed to its $\pm 1\sigma$ value. When calculating the result, a range of values are sampled for each nuisance parameter between the minimum and maximum (with a Gaussian constraint). So the numbers in the tables do not represent a shift in the result and the overall effect is relatively small. The results are calculated without any nuisance parameters in section 5.6 to quantify the final effect.

5.5 LUX Data Analysis

The Run 3 reanalysis followed the same methodology as the first analysis [185], but with several improvements. The changes and main points are noted here, and more details can be found in Refs. [148] and [255].

The event selection starts with those with a single S1 followed by a single S2, which we term ‘golden events’. Events with extra S2s before the S1 are allowed, along with events with extra S1s after the S2 – although these are rare. If there are multiple S2s after the S1 (multiple scatter), or multiple S1s before the S2, then the event is rejected. The S2 is required to be larger than the S1. This selection is the same as in the original analysis.

The depth of events is determined via the drift time between the S1 and S2. The (x, y) position is determined via the hit pattern of the S2 on the top PMT array. As with the first analysis, this uses the Mercury algorithm [205] but with updated light response functions for the PMTs.

A fiducial cut is defined to avoid backgrounds from the edges of the detector. At the top and bottom the backgrounds are primarily from the PMT arrays. At the sides the wall background determines the fiducial radius. In the first analysis there was no rigorous wall background model, which motivated a conservative 18 cm fiducial radius (the average TPC radius is 25 cm). With a good wall model, the number of background events can be estimated allowing a larger fiducial of 20 cm. The cuts for the top and bottom of the detector remained the same as 38–305 μs (43.5–3.5 cm above the cathode).

The fiducial mass is determined using the tritium calibration [178] by taking the ratio of the number of tritium events in the fiducial volume to the number of events between the cathode and the gate. The fiducial mass is then that number multiplied by the known total mass of active xenon. This method does not require precise knowledge of the cold detector dimensions, and avoids systematic errors from the drift time or position reconstruction. The measured fiducial mass was 145.4 ± 1.3 kg. The larger fiducial radius gives an extra 27.4 kg over the first analysis.

There is a roughly 20% probability that a single photon that hits the PMT photocathode results in the production of two photoelectrons [262]. To avoid this a digital “spike-counting” method was developed [148] to calibrate the PMT waveforms, which gives the pulse area as the number of detected photons, denoted phd. A spike is an extrusion above threshold in a single PMT channel. Two samples where the probability of a single PMT receiving two photons was low were used to estimate the mean waveform area of single spikes; small S1s (below 10 phd) and SE pulses. The pulse area is estimated using both spike counting and the calibrated pulse areas. A look-up table was created of the most likely number of phd as a function of pulse area and counted spikes. This was done via a MC simulation to estimate the pile-up. For small signals the estimated phd is driven primarily by the spike counting. For S1s above 20 keV ER equivalent energy and all S2s pileup is dominant and the phd estimate is from the pulse area alone [148]. Calibrating with the spike counting method means that the wavelength-dependent double-photoelectron emission is accounted for.

Discrimination between NR and ER events is estimated from high statistics tritium calibration, by counting the number of events that fall below the NR band median. This gave an average discrimination of $99.80\% \pm 0.03\%_{stat} \pm 0.10\%_{sys}$ for events in the fiducial volume with an S1 smaller than 50 phd. Note that this discrimination is merely a figure of merit, it is not used in the statistical analysis.

Previously, only light in the bottom PMT array was used in the S2 pulse area due to the more uniform response in this array and less PMT saturation. However, it was found that using both arrays reduced the leakage of ER events into the NR band due to the higher photon statistics (hence the improved discrimination over that in the first analysis). The reduced leakage outweighed the degradation due to non-uniformities in the response of the top PMT array, so in this analysis both arrays are used to determine the S2 pulse area.

A data quality cut was used to exclude times with high rates of single electrons (SE). High energy events give a tail of SE due to delayed emission of electrons from the liquid surface. This can occur for several ms after a large S2. It can produce a background due to faked S2s from overlap of SEs or upwards fluctuations in single electron size, with random coincidences from S1s. The cut to eliminate these events is based on the total pulse area not contained in S1s and S2s. Up to a signal area of 630 phd the non-signal area must be below 80 phd, past which the allowed non-signal area slowly increases. This cut excludes less than 1% of NR events.

The lower S2 threshold is determined by the ability to remove events arising from the SE tails after large events. Below 150 phd the rate of fake S2s from electrons following large pulses is higher. There is also a very low probability for events with these small S2 signals to have a valid S1.

The lower S1 threshold is determined by a two-PMT coincidence requirement. After applying position corrections to normalise S1s to the detector centre, S1s smaller than 2 phd are possible. The upper S1 threshold was also increased to 50 phd for the reanalysis, because the ^{127}Xe activation line is now included in the background model. Previously this limited the range to 30 phe.

The total WIMP search data was a total of 105 days. More live time was added for the reanalysis. The slow control was used to monitor detector stability of several parameters: liquid level, grid voltages and currents, outer vacuum pressure, detector pressure, and circulation flow rate. Periods where any of these parameters went out of the specified bounds were excluded from the data. 2.5% was lost to the detector instability cuts, 1.8% to the trigger hold-off, and 1.6% to DAQ dead-time or lost files. The total live time was 95 days, an extra 10 live days over the previous analysis.

The detection efficiency for NR events is determined by simulations following NR calibration through elastic neutron scattering. The simulation model for NR is fit to the DD calibration data, as described for the background model. Then mono-energetic NR events are simulated and processed in the same way as experimental data. As it is a simulation, the true number of events is known so the detection efficiency at each energy can be determined. The low energy cut-off is 1.1 keV, which is the starting point of the signal model. This was 3.0 keV in the original analysis. At this point the NR detection efficiency is 0.3%, rising to 50% at 3.3 keV. In addition, a hand scan of AmBe calibration data was performed. Single scatter events were selected visually using only the waveforms. Out of these, $97.5 \pm 1.7\%$ were identified after the data was processed, and this is applied as a flat acceptance to all energies.

5.6 Spin-Dependent Results

The fiducial exposure was 1.4×10^4 kg·days. A total of 591 events passed all analysis cuts, which was consistent with the predicted ER background, as shown in Fig. 5.11. The observed limit was found to be almost a 2σ downward fluctuation from the expected sensitivity (background-only hypothesis). If this were used, the measured limit would be more constraining than an equivalent result from a background-free experiment. The statistical test excludes a region where there is little or no sensitivity to a signal. Therefore, to avoid benefiting from such a fluctuation a power constrained limit is applied, which fixes the result to the expected (median) limit from background trials [263].

The resulting limits on the SD WIMP-nucleon cross sections are shown in Fig. 5.12. The LUX limit at low mass was aided by the low energy ER and NR calibrations from tritium and a DD neutron generator. These calibrations demonstrated the sensitivity of xenon to small recoil signals which allowed reduced analysis thresholds. The neutron-only limits from LUX are the most stringent from a direct search and exclude new regions of the WIMP parameter space. At low masses, the indirect limits from CMS are better but these are not sensitive to higher masses due to the kinematic limitations of producing a heavy WIMP at a collider. Also, there are uncertainties in the theoretical treatment used to translate the results into the same plane as the direct detection

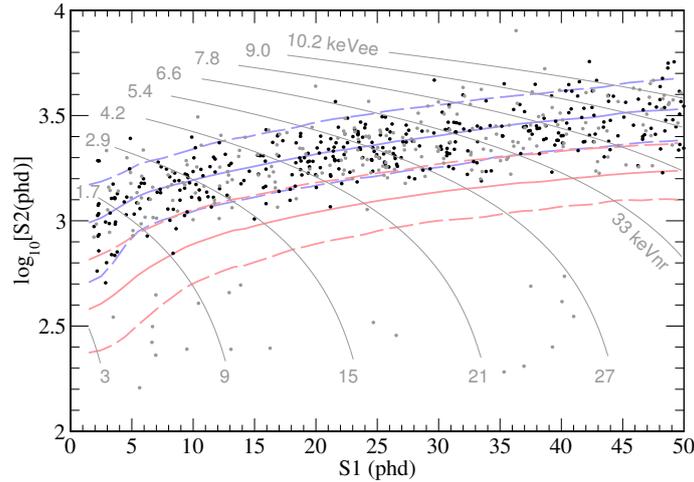


Figure 5.11: Observed events in the LUX 2013 data, from an exposure of 95 live days and 145 kg fiducial mass. The black points are events within the 18 cm fiducial, and the grey points are those between 18–20 cm. The blue lines represent the bands from an ER source with a flat energy spectrum. The red bands are an example from a 50 GeV WIMP with a spin-independent interaction. The dotted lines are the 10th and 90th percentile for S2 at a given S1, and the solid lines are the median. These bands are illustrative, and are not used in the statistical analysis. The gray lines are the linear combined S1+S2 energy scale, for the nuclear (keVnr) and electronic (keVee) energy scales indicated [264]. Figure reproduced from Ref. [148], with permission from the American Physical Society.

experiments. The proton-only limits are not as sensitive as those from the PICO experiments, which have enhanced sensitivity due to the unpaired proton of the fluorine nuclei in the C_3F_8 or CF_3I target. However, with the updated structure functions including two-body currents there is significant sensitivity to proton-only in xenon. LUX now excludes the DAMA allowed region in both channels assuming the standard halo model. For proton-only, indirect searches for dark matter annihilations in the sun from IceCube are the most stringent, as for the collider limits, there are uncertainties in the conversion to this plane.

The presented limits are for a particular choice of the WIMP couplings to protons and neutrons (a_p and a_n), and so it is model dependent. The method described in Ref. [271] allows limits to be set in the plane of a_p against a_n using:

$$\sum_A \left(\frac{a_p}{\sqrt{\sigma_p^A}} \pm \frac{a_n}{\sqrt{\sigma_n^A}} \right)^2 > \frac{\pi}{24G_F^2 \mu_p^2}. \quad (5.41)$$

This requires the limits from each isotope. The cross section $\sigma_{p,n}^A$ is either the proton- or neutron-only limit assuming all events came from isotope with mass number A . To include results from other experiments where only one of the proton- or neutron-only limits is presented (which is typical if an experiment is much more sensitive in one channel), the limits in the a_p - a_n plane can be found following the method in Ref. [272]. This is a good approximation if $a_p \gg a_n$ or vice-versa. The excluded regions for some particular WIMP masses are shown in Fig. 5.13.

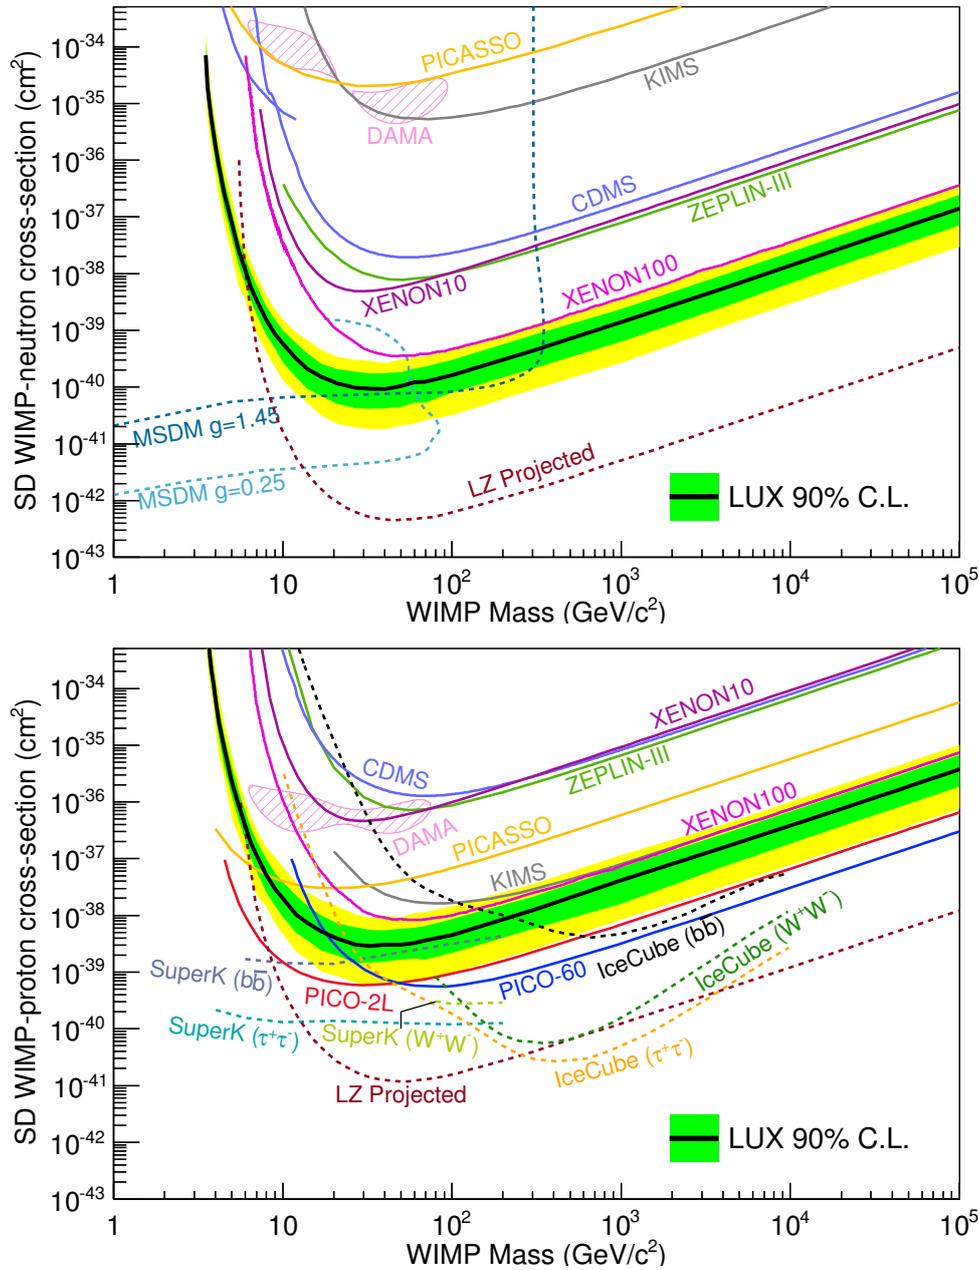


Figure 5.12: Upper limits on the WIMP-neutron (top) and WIMP-proton (bottom) elastic SD cross sections at 90% CL. The observed limit is shown in black with the $\pm 1\sigma$ ($\pm 2\sigma$) band from simulated background-only trials in green (yellow). Also shown are the 90% CL from: CDMS [226], KIMS [135], [265], PICASSO [266], PICO-2L [143], PICO-60 [144], XENON10 [267], XENON100 [268], and ZEPLIN-III [223], [269]. The DAMA allowed region at 3σ as interpreted in [134] without ion channeling is the shaded areas. Three indirect limits from IceCube [121] and SuperK [270] are shown. Collider limits from CMS mono-jet searches are shown, assuming the MSDM model with two coupling scenarios [77]. The projected sensitivity for the LZ experiment is shown for an exposure of 5.6×10^6 kg·days [197]. Reproduced from Ref. [231] with permission from the American Physical Society.

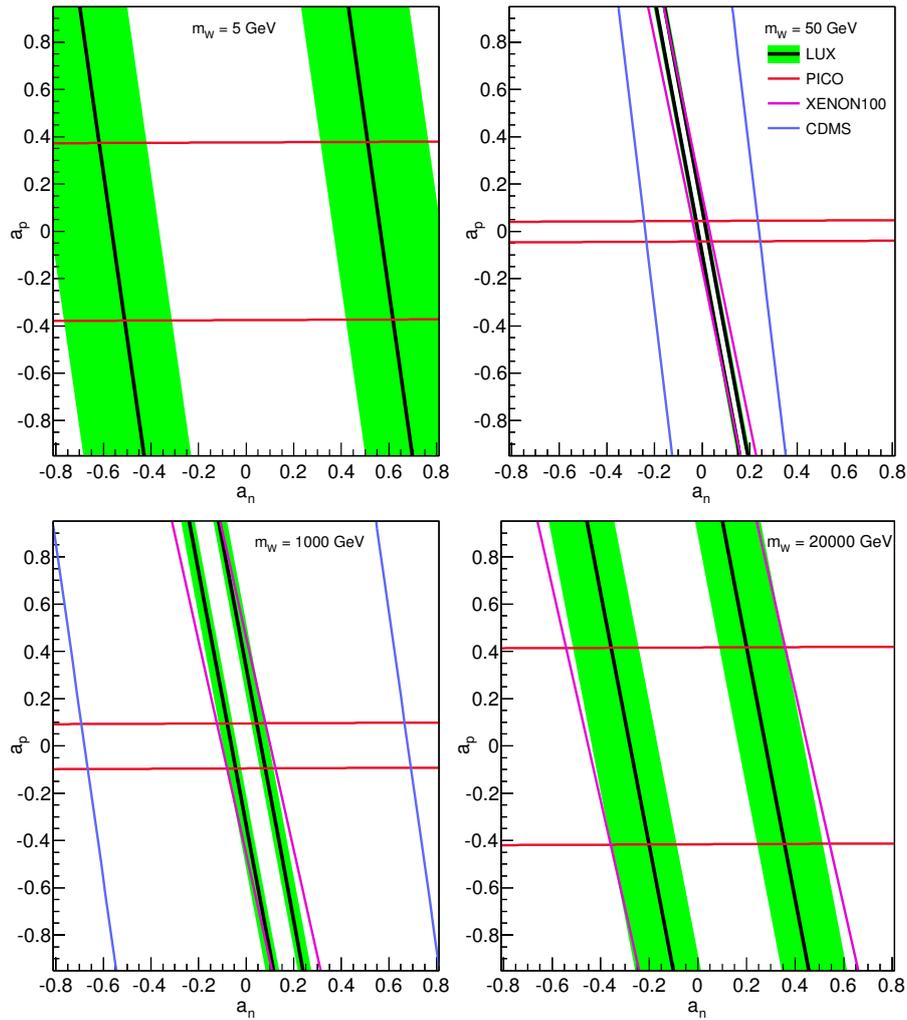


Figure 5.13: Constraints on the effective WIMP couplings to protons and neutrons, a_p and a_n , at 90% CL for various WIMP masses (5, 50, 1000, 20000 GeV). The LUX limit is shown in black with the $\pm 1\sigma$ band from simulated background-only trials in green. Also shown are CDMS [273], PICO-2L (for 5 GeV) [143], PICO-60 (for other masses) [144] and XENON100 [268], where the constraints have been inferred from the limits on $\sigma_{p,n}$ using the method in Ref. [272]. Reproduced from Ref. [231] with permission from the American Physical Society.

For each experiment the allowed region is within an elongated ellipse. This result improves the constraint on a_n over previous results. These excluded regions allow constraints on models with coupling to both protons and neutrons. The LUX and XENON100 ellipses are at slightly different angles due to differences in the assumed spin structure functions and the energy scale used in the analysis. There was also a slightly different isotopic abundance in XENON100, due to the addition of isotopically modified xenon. These excluded regions also show the complementarity between the detectors employing different target material.

Now we move on to discuss some technical points from the results. As discussed previously, the final LUX results used the Klos structure functions [236] including g_2 variation as a nuisance parameter in the shape of the signal model, and the Lindhart k

as a nuisance parameter in the number of events. However, for low WIMP masses the g_2 variation had to be removed. At low WIMP masses there is a higher probability that the recoil does not produce an observable S1. This means that to generate the signal PDF more Monte Carlo trials are required to have an acceptable number of observable events in the histogram. The extra computation time required is compounded by including the g_2 variation, which requires the signal PDF to be regenerated each time the nuisance parameter is varied. Therefore, to keep an acceptable computational time, the g_2 variation is removed below 4.5 GeV for neutron-only and below 5.5 GeV for proton-only.

To interpret the dark matter annihilation searches from the sun as a limit on the scattering cross section, IceCube [121] used the method from Ref [274] and similarly from SuperK [270]. LHC searches for dark matter particles produced with a single jet or single photon [80] were translated to a limit on the scattering cross sections by the treatment in Ref. [77]. This uses a Minimal Simplified Dark Matter (MSDM) model, which allows the relevant kinematic properties of the collider searches to be describes with only a few parameters. Limits on the SD WIMP-nucleon cross sections were presented in Ref. [77] for the case where the coupling of the mediator to quarks and the dark matter particle are equal ($g = g_q = g_{DM}$). The limits included in Fig. 5.12 were those with the lowest and highest couplings presented in Ref. [77].

Lastly, it is useful to quantify the effects of nuisance parameters and structure functions on the LUX exclusion curves. In Fig. 5.14 the cross section limits are reproduced without any nuisance parameter variation. The results using the Klos structure functions are compared to those with the Ressel and Dean structure functions [238] in Fig. 5.15.

The effect of the nuisance parameters is noticeable, but the difference is much smaller than the 1σ error band on the limits. Therefore, including only the most significant nuisance parameters is justified, as even these do not have a large effect. At most the limit with nuisance parameter variation was 23% higher than without.

Using the different structure function has a larger effect, with the neutron-only from Klos on average 1.9 times higher than the Ressel and Dean limit. For proton-only, the limit using the Klos structure functions is better than from Ressel and Dean by on average a factor 0.4. These differences are mostly due to the inclusion of two-body currents in the Klos structure functions, which enhances the proton-only limit and degrades the neutron-only. The limits calculated with the Ressel and Dean structure functions also allow a direct comparison to previous results that used this structure function.

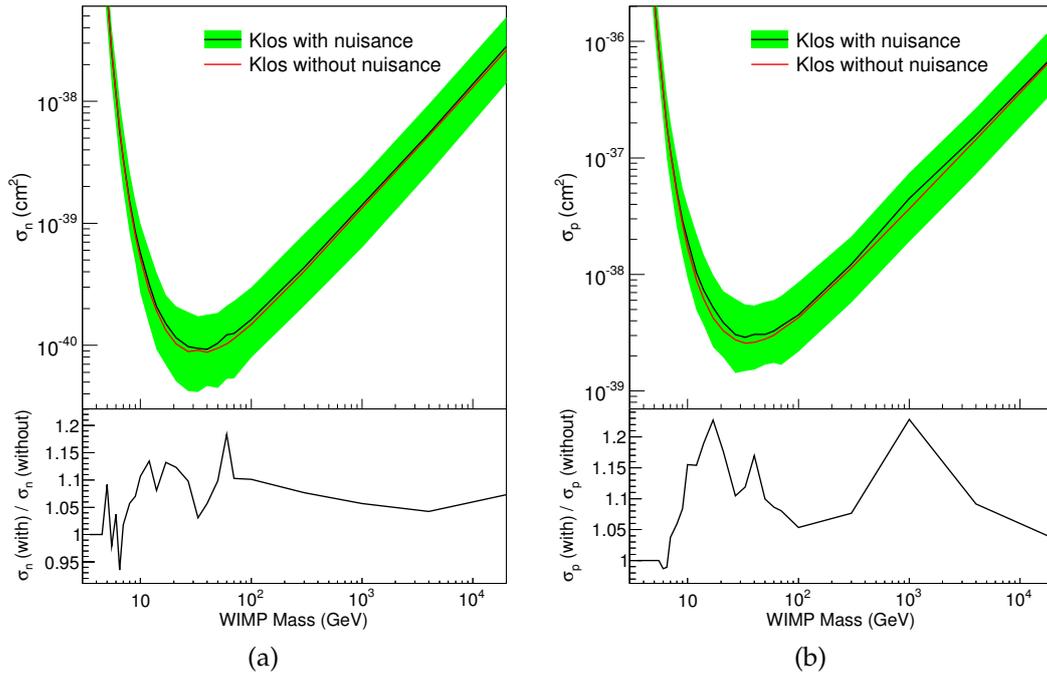


Figure 5.14: Upper limits on the SD WIMP nucleon cross sections at 90% C.L. with and without nuisance parameter variation. The top panels show the absolute cross section limits, and the bottom panels the ratio between the two. Note that when the ratio is exactly 1 at low masses it is because the limit was not calculated with nuisance parameters. The $\pm 1\sigma$ band is shown for the limits with nuisance parameter variation. (a) SD neutron-only cross section. (b) SD proton-only cross section.

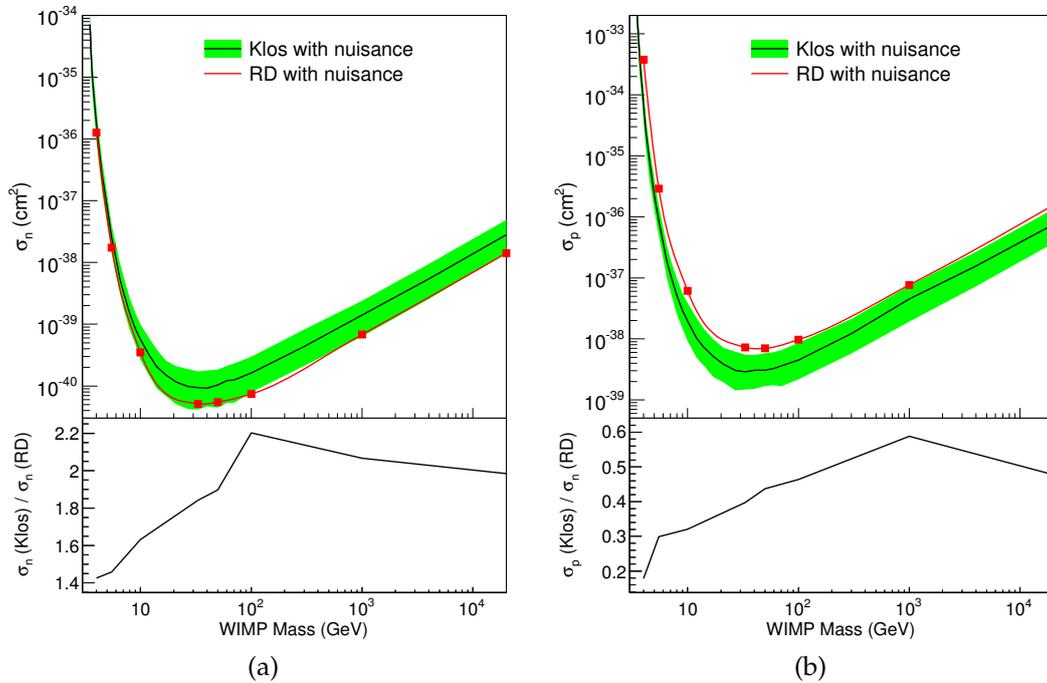


Figure 5.15: Upper limits on the SD WIMP nucleon cross sections at 90% C.L. with structure functions from Klos [236] and Ressel and Dean [238]. The top panels show the absolute cross section limits, and the bottom panel the ratio between the two different structure functions. The $\pm 1\sigma$ band is shown for the limits from the Klos structure functions. (a) SD neutron-only cross section. (b) SD proton-only cross section.

5.7 Conclusions

The results from the reanalysis of the 2013 LUX data set the most stringent limits on the SD WIMP-neutron scattering cross section from a direct search experiment. At low masses there are tighter (but more model-dependent) constraints from collider experiments, but at high masses this result sets the lowest limit on these interactions. For SD proton-only scattering, the most stringent limit is from indirect searches and PICO is the most competitive direct search.

With the inclusion of two-body currents in the spin structure functions for xenon isotopes there is substantial sensitivity even to proton-only coupling. The projected limit for LZ is approximately two orders of magnitude below the current limits from direct searches. The SD limits from LUX will be improved after Run 4, with a total of 300 live days acquired in 2015/16.

Future analyses may also use an effective field theory (EFT) treatment of the WIMP-nucleon interaction [243], as the standard SI and SD responses are not the only possible interactions. There are several possible operators, and constraints can be set on each one instead of SI or SD limits. Two of the operators have a coupling to spin, with the projection of the spin either parallel or perpendicular to the momentum transfer. The standard SD response is a linear combination of these two operators.

Chapter 6

Experimental Study of Electron Emission from Thin Wires in Xenon

The operating voltage of liquid xenon TPCs is often limited by the onset of anomalous electron emission from cathodic wires at certain locations on the electrode grids. This is observed at fields between 20 and 100 kV/cm on the wire surface, but electron emission processes are not expected until much higher fields. There are several possible sources of these emitters, such as defects or impurities on the wire surface.

A grid conditioning campaign was carried out in January 2014 after LUX Run 3 was completed, where the grid voltages were increased past the onset of emission (into a discharge regime) in order to “burn-off” the emitters and raise the operating voltages of the grids. In this work, the conditioning data are used to study the nature of the electron emission into gaseous and liquid xenon before the onset of macroscopic discharge.

This chapter begins with an introduction on the possible emission sources and mechanisms. Section 6.2 describes the analysis method and the results from the investigation of the electron emission. The results are summarised in section 6.1. A search for possible light emission associated with the electron emission is carried out in section 6.4. Discussion on what the analysis says about the emission mechanism and overall conclusions are in sections 6.5 and 6.6 respectively.

6.1 Electron Emission Mechanisms

The most common mechanism for releasing electrons from a cold metal is field emission. When an electric field is applied to the surface the potential barrier acquires a finite width and conduction electrons can tunnel through. This quantum mechanical treatment was first presented by Fowler and Nordheim in Ref. [275]. The effect of applying an electric field is shown in Fig. 6.1. The electric field causes the barrier to decrease away from the wire, and the rounding at the interface is due to the image charge seen by an electron [276].

Fields of the order 10 MV/cm are required to give a significant rate from field emission. This field is orders of magnitude larger than the typical fields which induce

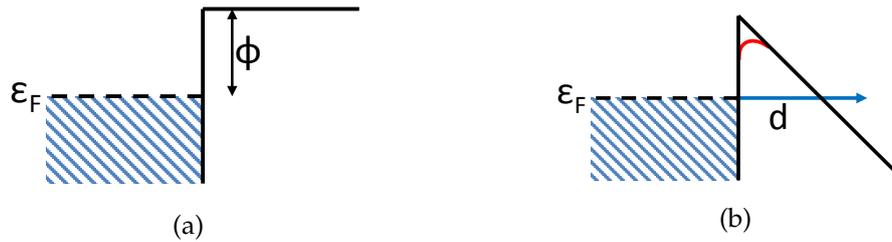


Figure 6.1: Potential barrier at the metal surface in the Fowler-Nordheim theory. (a) Before any electric field is applied (at 0 K); ϵ_F is the Fermi level of the metal with the shaded area representing the partially filled conduction band and ϕ is the work function of the metal. (b) With an electric field, the potential barrier obtains a finite width, d . The red line indicates the effect of the image charge. Electrons near the Fermi level can tunnel through the barrier, as indicated by the blue arrow [277].

emission onset from wires in TPCs – only 20–100 kV/cm on the wire surface. Therefore, field emission alone cannot explain the localised electron emission in TPCs at these fields.

One possible explanation for the anomalous emission observed in some broad-area cathodes is that there are sharp points protruding from the metal surface. Around these points the electric field is enhanced and could be large enough to induce field emission. The field enhancement factor β is used to quantify this, and is given by the ratio between the microscopic field at the tip and the usual field expected on the perfect surface. To reach fields of the order 10 MV/cm on wire surfaces in TPCs requires enhancement factors of 100–200. Measurements of emission from stainless steel in Refs. [278] and [279] gave β factors between 100 and 700. Tips capable of this enhancement have been observed by electron microscopes at emitter sites in Refs. [280] [281] [282] and [283]. However, no protrusions were observed on a new wire before an electric field was applied, only after macroscopic discharge or sparking occurred. Therefore, either the pre-breakdown currents have a different origin, or the protrusions form as the electric field is applied. There are other reasons to disfavor this explanation, which are discussed later.

Other sources of emitter must be considered in order to explain the pre-breakdown emission. One alternative is from contaminants, either embedded into the metal or particles on the surface. Another way of enhancing the field is conducting particles on the cathode surface [284] [285]. These act in a similar way to protrusions from the wire as the field around the particle is enhanced due to the geometry.

There is also evidence of insulating or semiconducting contaminants being the source of electron emission [279] [286] [287]. There are several proposed theories of how insulator materials can cause field emission, and the process depends on the exact conditions and materials involved. One example is that at some critical electric field an electroforming process takes place in the insulator, creating a conducting filament through the impurity [288]. Another involves the high field penetrating through a thin layer of insulator, leading to an altered band structure as shown in Fig. 6.2. The field is enhanced at the metal-insulator junction, allowing tunneling of electrons into the insulator. The electrons are then emitted thermionically over the potential barrier at the

insulator surface [289]. Another possibility is that the adsorbate has an energy level near to the Fermi level of the metal and electrons tunnel coherently into the adsorbate, and then to the outside [290]. This produces a resonance effects, where the emission is enhanced at particular fields.

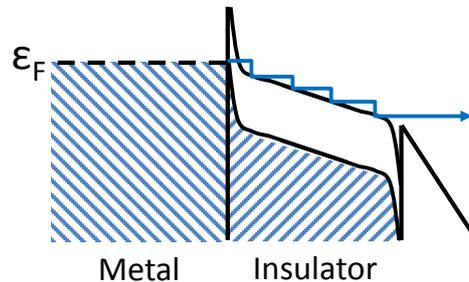


Figure 6.2: Potential for electron emission through an insulating impurity, the blue arrow represents the path of the electrons. As a result of the altered band structure, electrons tunnel into the insulator and are then emitted thermionically [277].

There are also other types of surface defect, such as scratches, holes, or grain boundaries. In Refs. [291] emitters were often found in association with such defects. However, when the field enhancement expected from the geometry was calculated the maximum β factor was only 30. These features often have increased concentrations of contaminants. The emission could be caused by a combination of a field enhancement from the geometry and the effect of the impurity. It is also suggested that a reduction in work function due to the impurity could contribute to the emission [292].

In conclusion, there are several possible causes of field emission on broad-area cathodes. Although there is evidence that whisker-like protrusions from the wire can give field enhancements and cause field emission they have only been observed after previous breakdown or sparking, not on fresh wires. There is evidence that pre-breakdown field emission is caused by impurities or contaminants within or on the cathode surface, often in association with other defects in the metal surface. It is possible that multiple types of electron emitter are present.

6.2 Analysis of LUX Electron Emission Data

The HV conditioning campaign was carried out by gradually increasing the gate, anode, and cathode voltages to look for the onset of emission. After emission was observed, the ramp often continued to somewhat higher fields to condition the grids as is typically done for commissioning room temperature TPCs. When the emission was first observed during a test, and at many points as the ramping continued, the grids were held at constant voltage for short times. The conditioning was split into two sub-campaigns: gate-anode and cathode. Here, we study the former and use the stable periods to investigate the nature of the emission.

There are data available with the liquid level above the gate, and with the liquid a few mm below the gate. When the liquid level is above the gate, as for regular data taking, the emission of electrons from the gate wires is separated in time from the

EL emission in the gas. Some number of electrons are released from the wire. There could be multiple electrons released simultaneously from the surface, or one electron which then causes ionisations in the high field around a defect on the wire. The field around the gate wires is high anyway, just above 50 kV/cm, because of focusing of the field. However, it is not high enough to cause electron multiplication in the liquid. The electrons are extracted and produce electroluminescence in the gas as usual. An example waveform is shown in Fig. 6.3a.

For the data where the liquid is below the gate, the signals are quite different. Any electrons released from the gate wire start to produce EL straight away, since this threshold is only ~ 2 kV/cm. Also, the field in the vicinity of the gate wires is large enough to cause electron multiplication. Therefore, extra charge can be due to multiple electrons being released from the wire at once, multiplication from the field around a sharp point on the wire, or from the high field around the gate (slightly away from the wire). In any case, whatever their origin, all electrons are produced near to the wire and so each will have approximately the same photon yield in the gas. The pulse shape will also be different to usual, with a peak at the start due to the non-uniform electric field, as shown in Fig. 6.3b.

There are some aspects common to both the data with the gate in the gas (hereafter “gas data”) and with the gate in liquid (“liquid data”). Firstly, pulses from each emitter must be identified. The pulses are similar to regular S2s or single electrons and the standard LUX pulse classifier could identify them well. The pulse rate can be much higher than in standard WIMP search data, so the maximum number of stored pulses per event was increased from 10 to 100 in the data processing. Before applying other cuts the positions are plotted to define an x - y cut for each emitter. The position is reconstructed using the Mercury algorithm [205], the same as for regular data. As each emitter can exhibit different behavior, it is important to isolate them. Each emitter will be labelled by a capital letter and its date (to distinguish the same emitter in different stable periods).

Several aspects are studied for each emitter. The rate of electron emission can be monitored over time along with the field on the gate wire, E_{wire} (assumed perfect). This is calculated from the grid voltages, as recorded by the slow control. The method is based on the calculation for the field around a single wire grid from Ref. [293], extended to multiple grids [154]. This is useful for monitoring any changes in the emitter, for example to see if the conditioning is reducing the emission rate, or to see the effects of changing the grid voltage. Using the entirety of the LUX detection capability means it is very sensitive to small emission rates, emitters of only 1 electron/s are identified. This is a current of 10^{-19} A; other experiments studying broad-area emission are typically sensitive to \sim pA. So, this is a different emission regime than other tests.

The subsequent stages study the emission during the periods of stable grid conditions. The only further cut applied is on the pulse length, which removes cases where two pulses overlap. The value of this cut varies for each dataset as there are different fields and operating pressures. It is a $\pm 2\sigma$ cut around the peak in the pulse length histogram.

The previous steps apply to both the gas and the liquid data. The remaining tasks involve studying the number of phe or electrons emitted per event. The method for

this is different for the gas data and the liquid data, which are described separately in the following sections. For the liquid data, the possibility of photon production at the wire is also investigated. For example, in Fig. 6.3a the single phe at the gate wire could be associated with the electron emission.

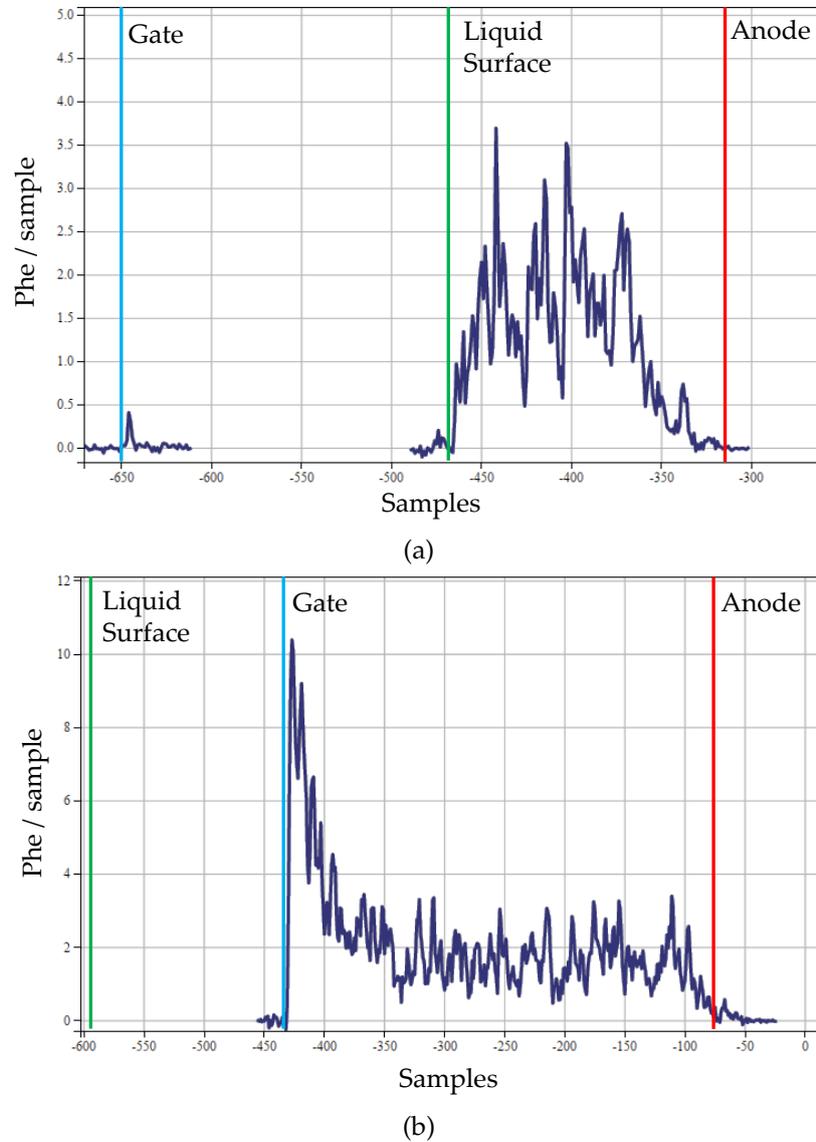


Figure 6.3: Example waveforms from electron emission data. The approximate positions of the gate, liquid surface and anode (in time) are shown. (a) With the gate in the liquid phase. Electroluminescence starts as electrons are extracted from liquid into gas. (b) With the gate in the gas phase. Electroluminescence starts as soon as electrons leave the gate. The field is higher around the gate wire, which produces the peak at the start of the pulse.

6.2.1 Gate in gas

The first gate-anode dataset with the liquid below the gate was taken on Jan. 10th 2014. Previously, tests took place with the gate in the liquid. In the first gas test the gate and anode voltages were gradually increased from ± 2 kV. There were two stable periods

and several different emitters were observed from the gate wires. Figure 6.4 shows the locations of the pulses and Fig. 6.5 shows examples of the pulse rates.

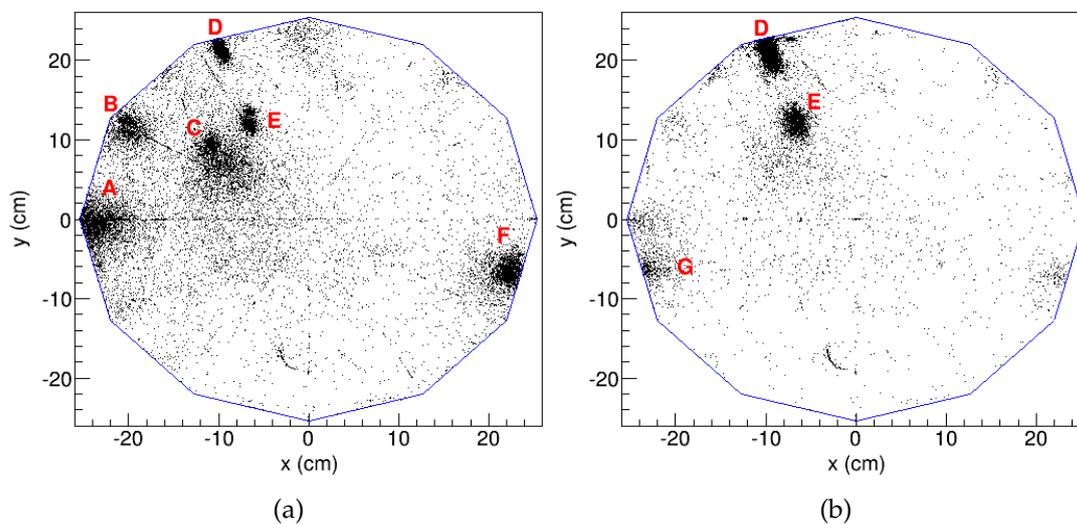


Figure 6.4: Reconstructed positions of electron emitters in gas on Jan. 10th 2014 during the two stable periods. (a) At 7600 s after start. (b) 8600 s after start.

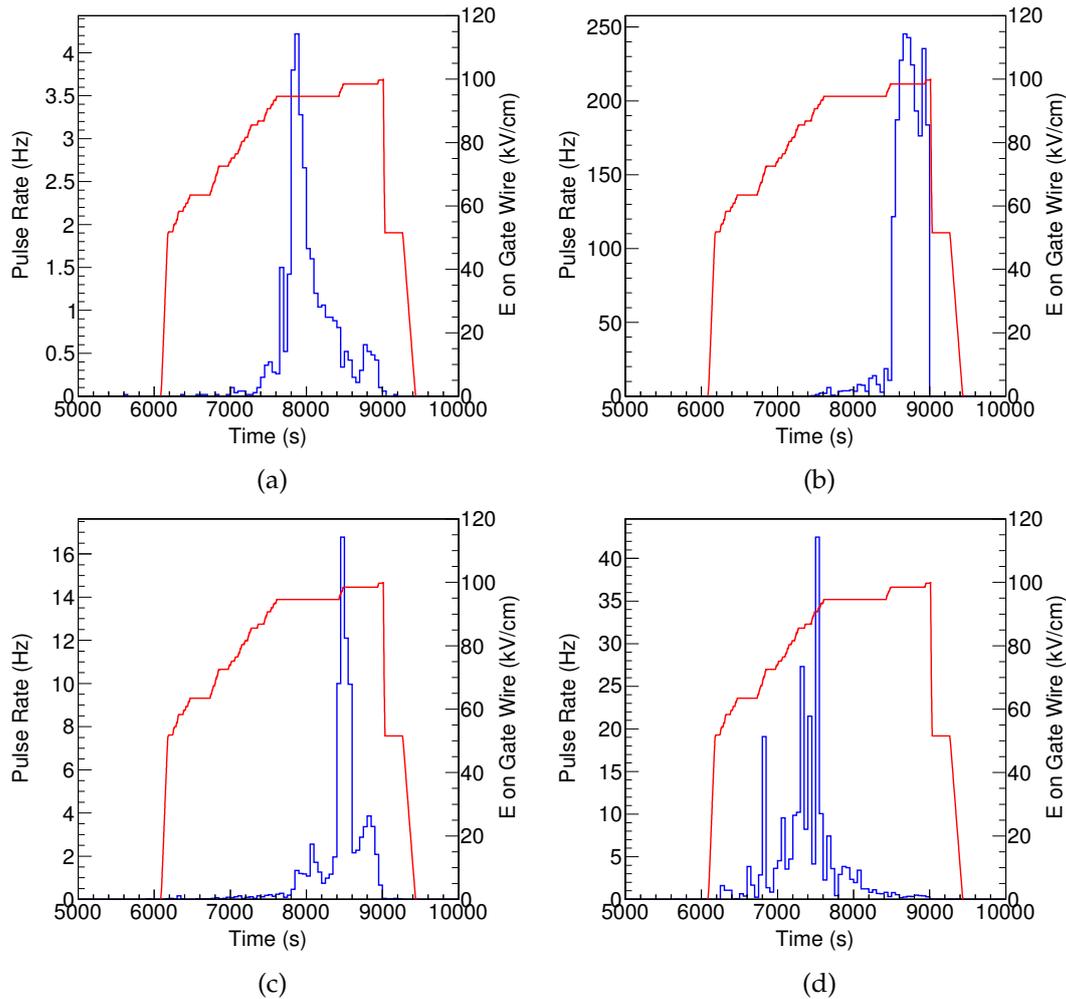


Figure 6.5: Pulse rate (blue) and field on gate wires (red) for emitters in gas on Jan. 10th 2014. For emitters: (a) C (b) D (c) E (d) F.

Most of the emitters, such as C in Fig. 6.5a, show a peak in the emission rate, which then decreases. For A, B, and C the maximum rate (7, 8, and 4 Hz respectively) occurs during the first stable period. For G, the rate is negligible until the second stable period, where it reaches a maximum of 6 Hz. Emitters E and F are similar, but have the maximum rate during ramping. It appears that short peaks in emission rate are common, and those emitters quickly subside. In particular, emitter F completely disappeared by the end of the dataset, only 1000 s after the peak. Different behaviour is observed for emitter D, which has sharp emission onset as the field is increased for the second stable period and the emission rate remains stable. This emitter was observed in the liquid data on Jan. 9th 2014, which is discussed in section 6.2.2.

The emission in the two stable periods is investigated further by looking at the number of detected phe for each pulse. The situation of one electron being emitted from the wire surface can be simulated using Garfield++ [294]. This is a Monte Carlo simulator of electron drifts through gas, which gives the locations of any excitations or ionisations. Further details on the Garfield++ simulation can be found in section 7.3. The scenario simulated here is a plane of parallel wires for the gate, between two planes held at constant voltage for the anode and liquid surface. The simulation pro-

vides a way to distinguish between a single electron emitted from the wire or multiple electrons. If the data agree with the simulation it must be a single electron emitter; one electron at a time released from the wire followed by the charge multiplication one would expect from the “perfect” wire. Otherwise, there are more electrons than expected from at, or near to, the wire surface.

In the first stable period emitter B showed good agreement with the simulation, as shown in Fig. 6.6. The two histograms are only scaled so that the total number of events in data and simulation is the same. In Fig. 6.6a the first two peaks are resolved, which correspond to the events where only one or two electrons were produced. As the simulation only gives the excitations and each is assumed to produce one photon, the photon detection efficiency (PDE) was required. In this case the PDE is found by repeating the simulation with multiplication turned off, to find the number of photons emitted for one electron. This corresponds to the first peak in the data. For the first stable period the simulation gave 730 photons/ e^- when starting electrons from 0° from the horizontal plane and the first peak in the data was 56.4 phe/ e^- , giving a reasonable PDE of 7.7%. With this the simulation agreed with the data well, even for larger numbers of phe where there are many electrons. This means that this particular emitter is only releasing one electron at a time from the wire, and all of the secondary electrons can be explained by the high field around the gate wire.

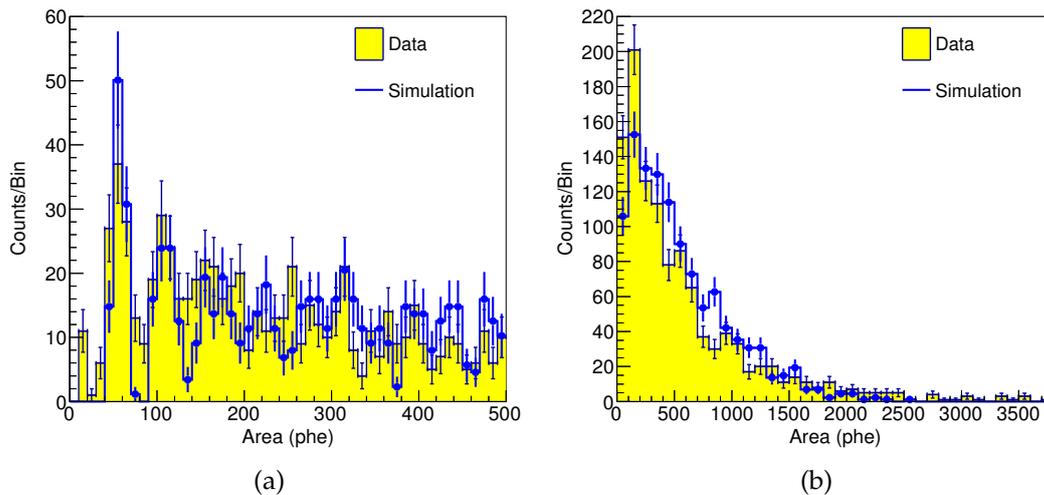


Figure 6.6: Number of phe produced for emitter B on Jan. 10th 2014 during the first stable period with $|E_{wire}| = 95$ kV/cm. Simulation is of individual electrons from the wire surface. (a) up to 500 phe (b) whole range.

Emitter A also showed a peak in the phe histogram as shown in Fig. 6.7. However, the simulation has a larger fraction of events at higher phe, which suggests that electrons are being lost due to the proximity of the wall. The PDE used reproduces the positions of the two observed peaks well. Emitter F showed the same behavior.

As the single electron response in this dataset is known, a histogram of the number of electrons released can also be produced and compared to the simulation. The histogram is fitted with a Polya function which describes the multiplication process in the

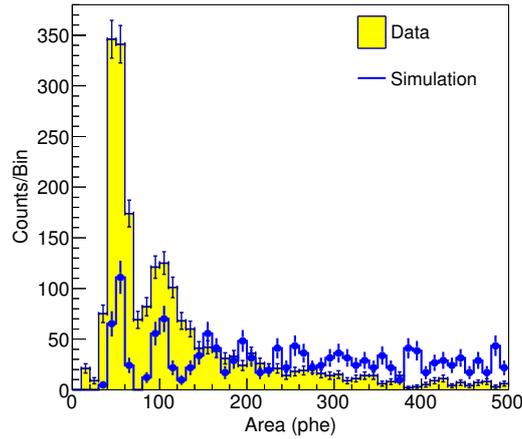


Figure 6.7: Pulse area from emitter A on Jan. 10th 2014 during first stable period ($|E_{wire}| = 95 \text{ kV/cm}$). The simulation is of individual electrons from the wire surface.

non-uniform field around a cylindrical wire [295], and can be written as [296]:

$$P(N) = \frac{1}{\mu} \frac{(\theta + 1)^{\theta+1}}{\Gamma(\theta + 1)} \left(\frac{N}{\mu}\right)^{\theta} e^{-(\theta+1)N/\mu}. \quad (6.1)$$

This describes the probability of observing N electrons for a given mean number of electrons emitted μ , and the shape parameter θ . The Polya is often found to describe electron multiplication processes well, for example avalanches in micromegas detectors [297]. Alternatively, if the process had a fixed low emission rate, it may be expected to follow a Poisson distribution. Therefore, the histogram of the number of emitted electrons is also fitted with a Poisson distribution.

Two examples from the first stable period, emitters B and E, are shown in Fig 6.8. Unlike the single electron emitter B, emitter E in Fig. 6.8b has many more electrons per event than the simulation. Therefore, in E there are extra electrons from at or near to the wire surface. Emitter D was similar to E, with a higher mean of 110 electrons per pulse. From this first stable period alone, we see that there is no apparent relation between the rate and number of electrons emitted per event. For example, emitters D and E are both multiple electron emitters, but have completely different rates. Then, B and E both have low rates ($\sim 10 \text{ Hz}$), but emit different numbers of electrons. Therefore, the nature of the emission depends on the individual site, and not only the macroscopic conditions.

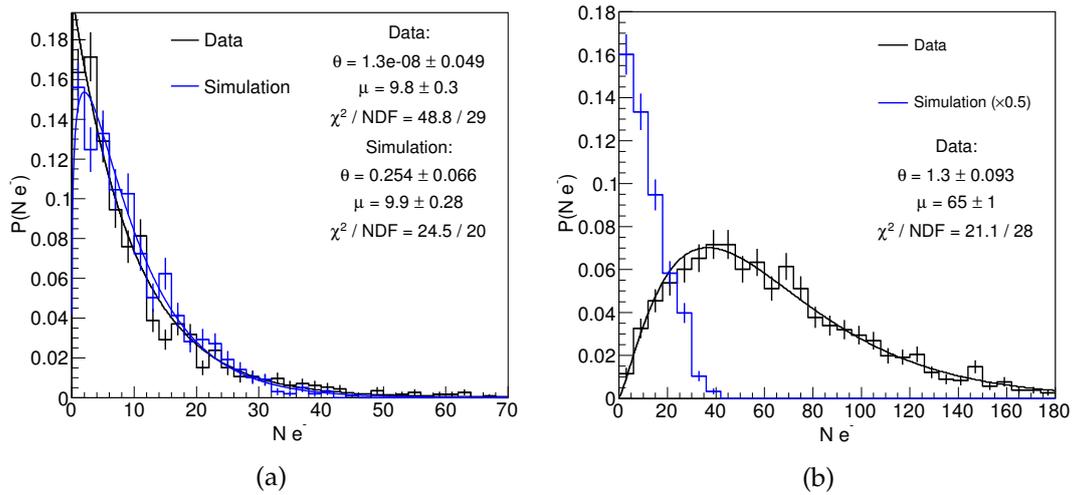


Figure 6.8: Number of electrons per pulse from emitters in gas on Jan. 10th 2014 during the first stable period ($|E_{wire}| = 95$ kV/cm). (a) Emitter B. (b) Emitter E.

The change of voltage for the second stable period affects the emitters in various ways. By this time emitters A, B, C and F have almost completely disappeared. The new emitter, G, is close to the edge and probably lost electrons to the wall (like A in Fig. 6.7). Emitter D underwent a drastic change, with the rate increasing from 10 Hz to 240 Hz, as can be seen in Fig. 6.5b. The mean number of electrons doubles from 110 to 220 per pulse. Conversely, emitter E had a lower pulse rate than the first stable time (although there was a spike during the ramp). Also, the average number of electrons per pulse decreased and the single electron simulation agreed well with the data. This was used as validation in section 7.4 and can be found in Fig. 7.25. It is similar to that shown in Fig. 6.6a. It appears that changes in macroscopic conditions can induce changes in the microscopic emitters, but exactly what the effect is depends on the particular emitter.

After these data the grids were left off overnight and conditioning was performed the following day. The detector was monitored on one channel using an oscilloscope, but full PMT data was not recorded. The conditioning was 30 minutes at 100 kV/cm on the gate wires, and 50 minutes at 110 kV/cm. At the higher field discharge from the gate was observed. The next available data was from a short test on Jan. 12th 2014. Only one emitter was observed at 65 kV/cm on the gate wires, a lower field than the tests on Jan. 10th. The location of the events and the pulse rate are shown in Fig. 6.9.

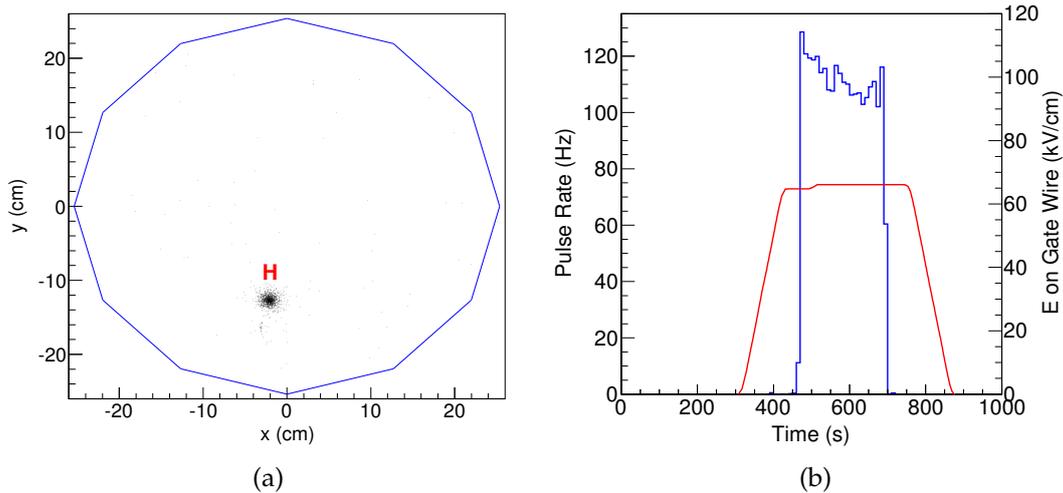


Figure 6.9: Emitter H in gas on Jan. 12th 2014. (a) Reconstructed position of pulses. (b) Pulse rate (blue) and field on gate wires (red).

As there were no peaks observed in the phe histogram, the single electron size was estimated from Garfield++ (416 photons per e^- with a starting angle of 45°) and the 9.1% PDE for the previous emitter away from the detector walls. The number of emitted electrons per pulse is shown for each stable period in Fig. 6.10, the distribution is similar for both. The number of released electrons is higher than the single electron simulation, so there are extra electrons from near the wire surface. The rate is quite stable, only gradually decreasing over time, which carries on even when the field is increased. It appears this emitter was created during conditioning on Jan. 11th.

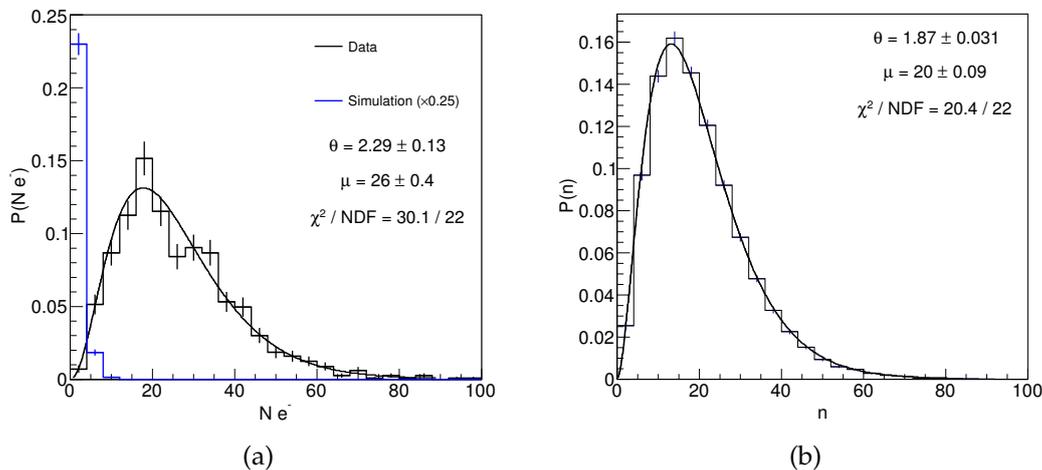


Figure 6.10: Number of electrons from emitter H in gas on Jan. 12th 2014, during the (a) first ($|E_{wire}| = 65$ kV/cm) and (b) second ($|E_{wire}| = 66$ kV/cm) stable period. The simulation is of a single electron released from the wire surface.

Two more tests were conducted on the Jan. 13th 2014 and the same emitter was observed. The emitter location and pulse rate for the morning test is shown in Fig. 6.11, and for the evening test in Fig. 6.12. After of the morning test the grids were left at that voltage for approximately 3 hours for conditioning. In the evening test, the positions

shown in Fig. 6.12 were affected by a PMT being turned off above the emitter, which causes larger pulses to be mis-reconstructed upward. This was done to protect the PMT directly above the emitter from the high emission rate.

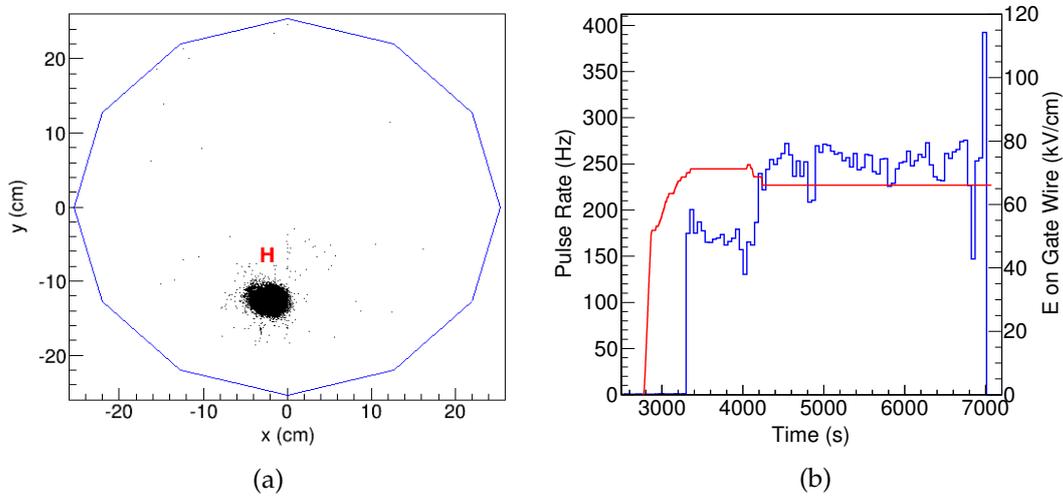


Figure 6.11: Emitter H in gas on the morning of Jan. 13th 2014. (a) Reconstructed position of pulses. (b) Pulse rate (blue) and field on gate wires (red).

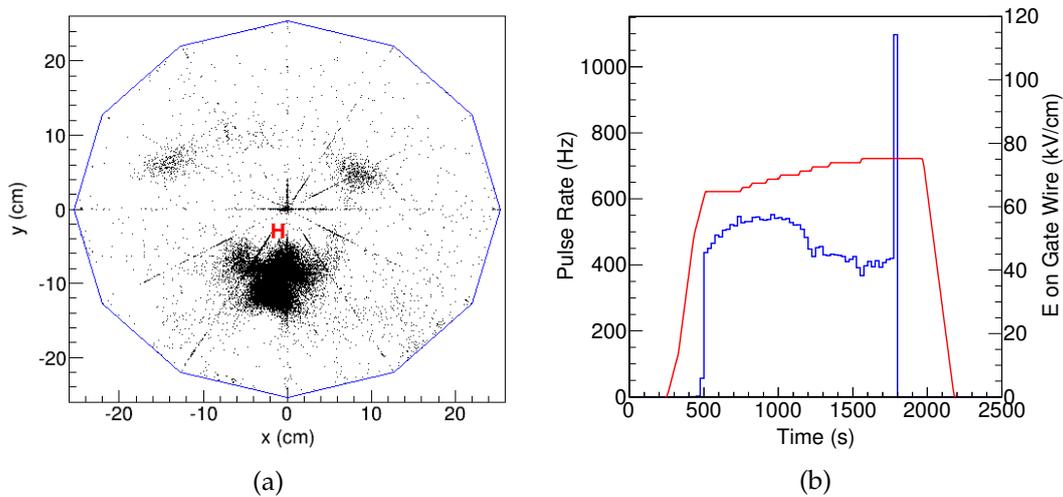


Figure 6.12: Emitter H in gas on the evening of Jan. 13th 2014. (a) Reconstructed position of pulses. A PMT above H was turned off, causing some pulses to be mis-reconstructed upward. (b) Pulse rate (blue) and field on gate wires (red) from emitter H.

In the morning test there were two stable periods, with the latter having a lower field on the gate wires. As shown in Fig. 6.11b, the lower field results in an increase in pulse rate. In the evening test the pulse rate is higher, but does decrease as the test continues. In the first stable period in the morning test the average number of electrons per pulse was 29, but by the last stable period in the evening test it was 16 despite the electric field being similar. Two points can be taken from this emitter so far. Firstly, there is no simple relation between the field on the gate wires and the emission

rate or number of electrons emitted. Secondly, a gradual decrease in rate is observed and so the sustained emission and high field does affect this persistent emitter.

In the next test on Jan. 14th emitter H remains, but a new emitter is also present. The new emitter at the bottom right (Fig. 6.13a) was found to have electron losses on the wall, as the data had less electrons per event than the single electron simulation. The point of emission onset in H is higher than before, at 75 kV, as shown in Fig. 6.13b. This time the pulse area is larger than before, and the mean number of electrons is 43 per pulse, rather than 12 in the last test. All of the pulses are mis-reconstructed upwards. The larger pulse areas (and more electrons per event) could indicate that the emitter site has become more extended.

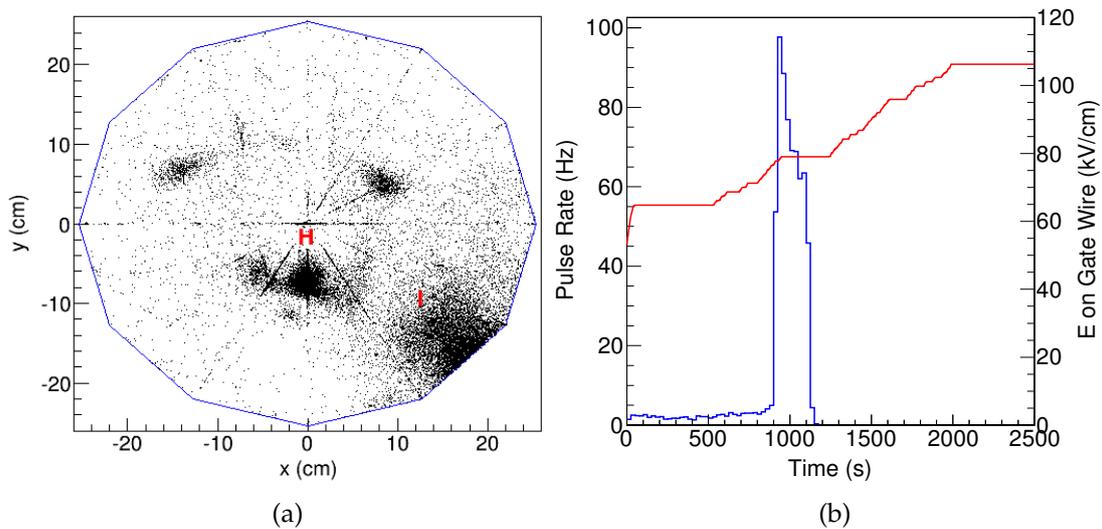


Figure 6.13: Emitters in gas on Jan. 14th 2014. (a) Reconstructed position of pulses. (b) Pulse rate (blue) and field on gate wires (red) for emitter H.

On Jan. 15th 2014 the liquid was raised above the gate for a test, before being lowered again to resume tests in the gas. The only significant emitter in the gas hereafter was emitter H, which reappeared on Jan. 15th 2014 at a field of 85 kV/cm on the gate wire surface. In the subsequent tests on the Jan. 17th and 22nd 2014 the onset threshold varied, as shown in Fig. 6.14.

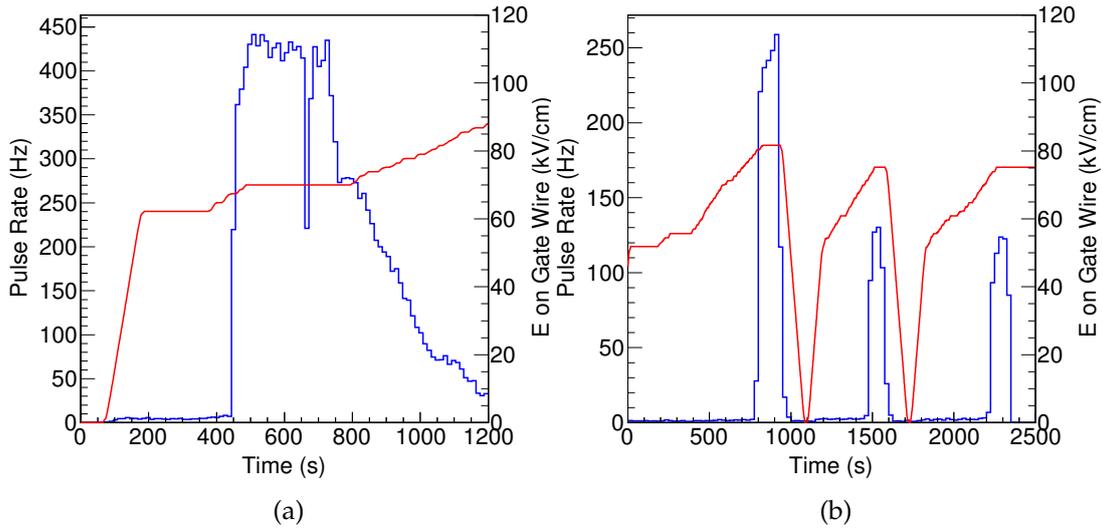


Figure 6.14: Pulse rates for emitter H during tests in the gas. (a) On Jan. 17th 2014. (b) On Jan. 22nd 2014.

Even after multiple ramps, emitter H is not removed. In the last stable period on Jan. 22nd 2014 the rate and number of electrons per pulse (Fig. 6.15) are actually similar to when this emitter first appeared (Figs. 6.9 and 6.10). Although at some points the emission appears to be subsiding, it can increase again in subsequent tests at similar fields on the gate wire. Therefore, it is unclear whether conditioning is effective for this emitter.

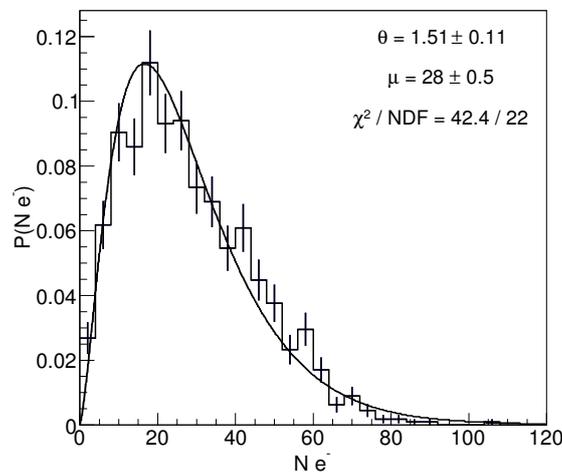


Figure 6.15: Number of electrons from emitter H in gas on Jan. 22th 2014 during the last stable period at 2300 s ($|E_{wire}| = 75$ kV/cm).

The grid conditioning on the anode and gate continued over the following days. There were multiple ramps with the gate in the gas, but to protect the PMTs they were turned off. Data with the gate in the liquid was taken on the Jan. 28th and 29th 2014, before another test in the gas on Jan. 30th 2014. An abrupt onset of emission was observed and pulse rates up to 1000 Hz were measured. However, as soon as this started several PMTs tripped. This makes it very difficult to draw conclusions about these

emitters, as it affects both the position reconstruction and size of the pulses. Although, this does indicate that emitters are still present.

By the end of the campaign, although emitters were still present, emission onset was at a higher field than before. There were many emitters at the start of the campaign which were quickly removed. After conditioning the grid voltages were raised to +7 kV and -1 kV, which is -70 kV/cm on the gate wire surface. During Run 3 the field on the gate wires was -43 kV/cm.

6.2.2 Gate in liquid

The pulse selection, identification of emitters, and stable times are the same as for the gas data. The pulse length cut is also applied. Looking at the number of phe and electrons detected in each event is slightly different to the gas data. As the electric field around the gate wire is not large enough to cause electron multiplication in the liquid, any extra electrons must come from the process at the wire surface. Therefore, there are clearly resolved peaks in the phe histograms for the liquid data.

The distribution of the number of emitted electrons is found by fitting Gaussians to the phe histogram. The expected mean and width for each peak can be written as a function of the second peak (the second peak is used rather than the first, because it is possible that some small single electron signals could be missed). The function for the peak with N electrons and amplitude A_N , as a function of the mean μ_2 and width σ_2 of the second peak is given by:

$$f(x) = A_N \exp\left(-\frac{(x - (N/2)\mu_2)^2}{2(\sqrt{N/2}\sigma_2)^2}\right). \quad (6.2)$$

The above formula is used to set the initial values for the fit. The means and widths are allowed to vary within 10% of the initial values. If the first peak is narrower than expected it would indicate that not all electrons are being detected. The number of events that contain N electrons is found by integrating each Gaussian; then the probability of emitting N electrons is plotted, and fit with the same Polya function as before.

The first data of the conditioning campaign was with the gate in the liquid, taken on Jan. 7th 2014. The gate-anode test started with the Run 3 voltages (anode +3.5 kV, gate -1.5 kV), before gradually ramping as shown in Fig. 6.16. Emission started at 54 kV/cm on the gate wires, and only one emitter was visible. The pulse area was studied for the short stable time with 60 kV/cm on the gate wire.

For the one observed emitter, the pulse rate increases sharply once going above 54 kV/cm. The first three peaks in the histogram are clearly visible and the tail can be reproduced. The Polya distribution fits well, which suggests it is consistent with an electron multiplication process. However, the Poisson fit could not reproduce the probability of one or two electrons.

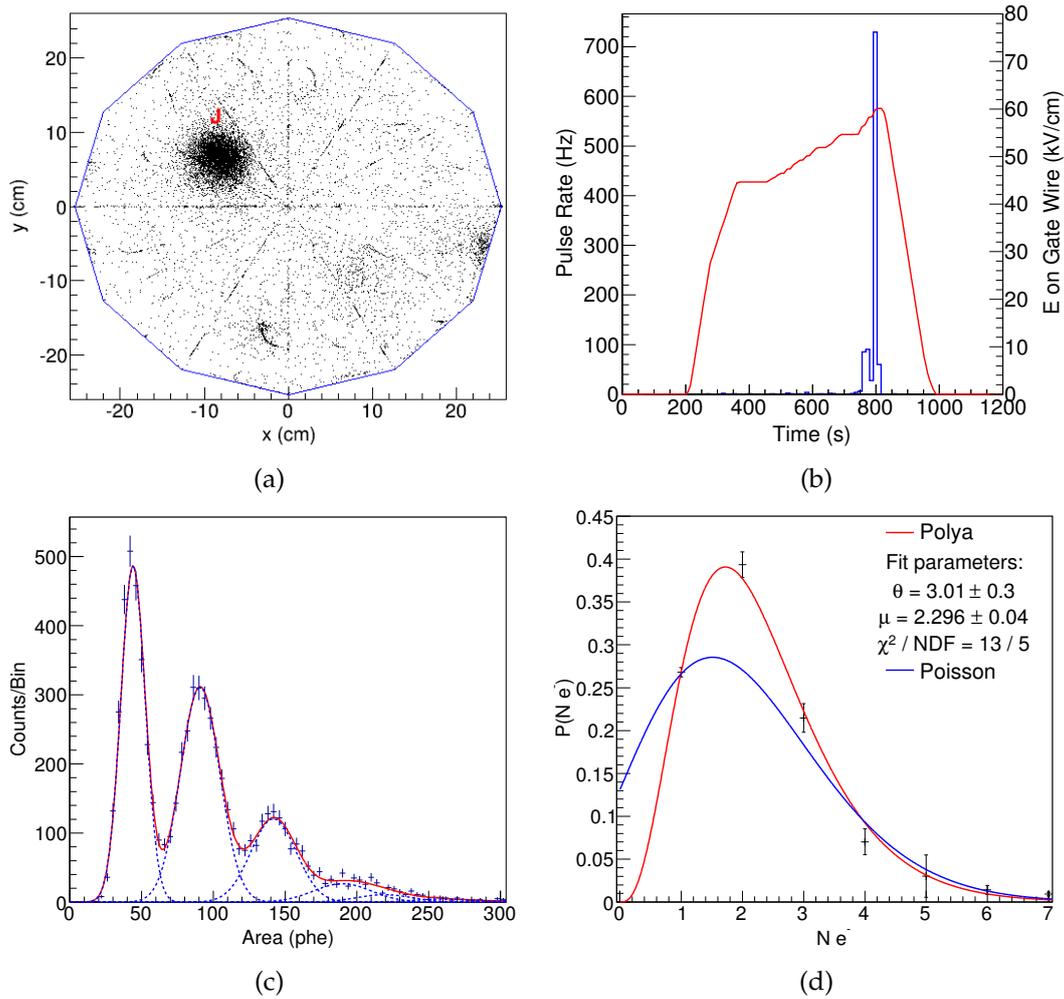


Figure 6.16: Liquid data from Jan. 7th 2014. (a) Reconstructed positions. (b) Pulse rate for the emitter J (blue), and field on the gate wire (red). (c) Pulse areas of J, with the multiple Gaussian fit (red). The blue dashed lines are the individual Gaussians. (d) Number of emitted electrons from J with a Poly and Poisson fit.

The following gate-anode test on Jan. 9th 2014 went to higher fields. The first onset of emission occurs for the faintest of the three observed emitters at a field of 46 kV/cm on the gate wires. The other two emitters start at a field between 60 and 64 kV/cm. Emitter J from the previous test was not present. The reconstructed positions and pulse rates are shown in Fig. 6.17.

Emitters K and D show a peak in the emission rate, which then drops as the field is increased further. L has more sustained emission, but the rate does drop as the field is increased. In the stable time at 1600 s emitter K only releases one electron at a time. A small peak for two electrons can be seen in the next stable time at 1900 s. This is the same as emitter L, another single electron emitter but with a much higher emission rate. The other emitter, D, has multiple peaks in the pulse area histogram. This emitter was later seen in the gas data on Jan. 10th 2014, as discussed in section 6.2.1. In Fig. 6.18 the pulse area histogram and the number of electrons are shown from the stable period at 1600 s. The overall distribution in the number of electrons changes very little even as the field on the gate is changed.

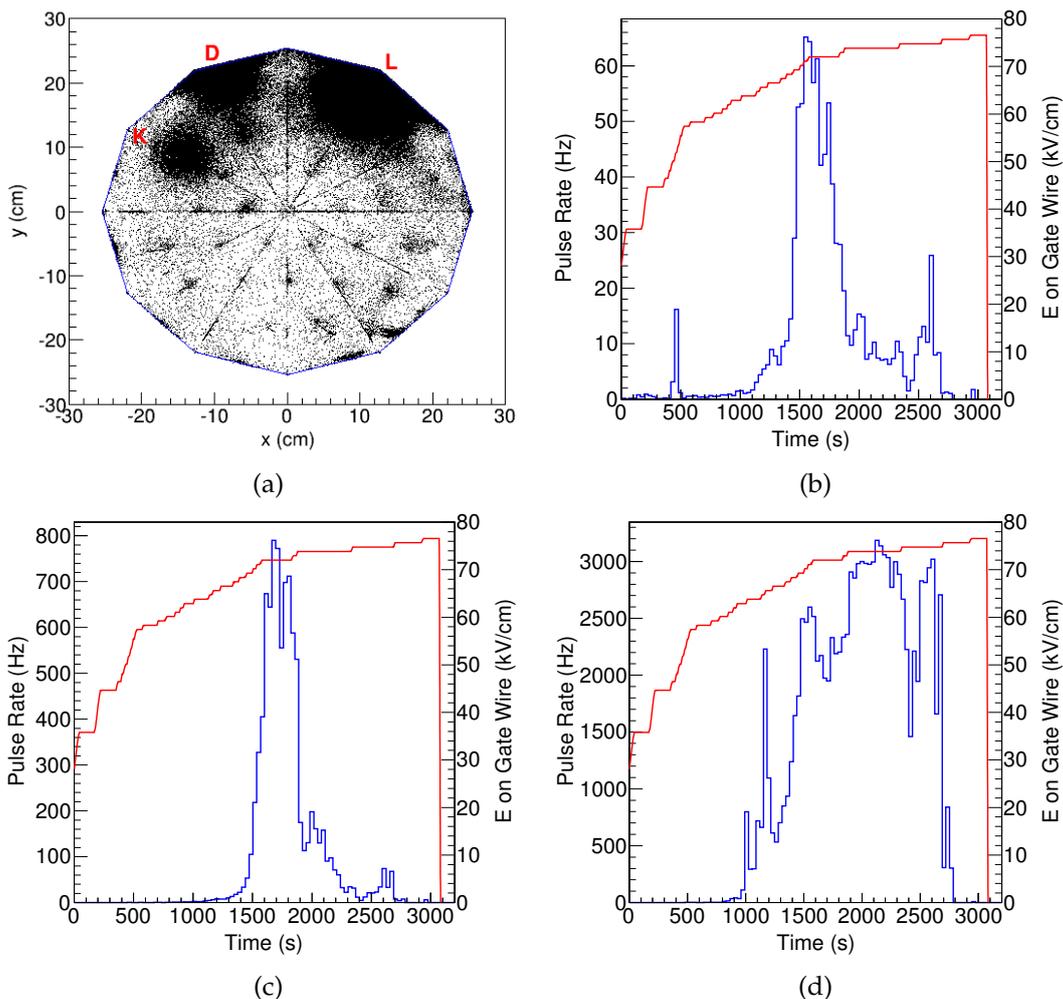


Figure 6.17: Emitters in liquid on Jan. 9th 2014. (a) Reconstructed positions. Pulse rates (blue) and field on gate wires (red) for emitters: (b) K, (c) D, (d) L.

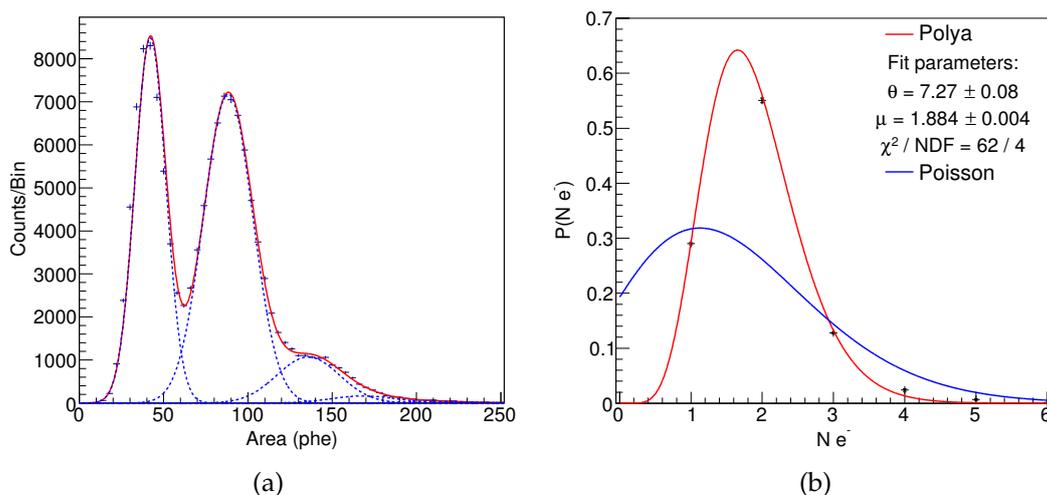


Figure 6.18: Emitter D in liquid on Jan. 9th 2014 during stable grid voltages at 1600 s ($|E_{wire}| = 72 \text{ kV/cm}$). (a) Pulse area. (b) Number of emitted electrons.

These first tests in the liquid are somewhat similar to those in the gas: there are peaks in the pulse rate which then quickly drop. Some emitters almost exclusively emit one electron at a time, whereas another has multiple electrons per pulse. Two of the emitters, K and L, were not present in the subsequent gas data on Jan. 10th 2014.

The next data with the liquid above the gate were taken on Jan. 15th 2014, with tests in the gas taking place since the last liquid test (see section 6.2.1). The locations of the emitters are different to those on the Jan. 9th 2014, as shown in Fig 6.19a. In this test the field on the gate was not raised as much as before, but there were multiple ramps. Emitter A almost exclusively emitted one electron at a time, as shown in Fig. 6.19d, and only had a significant rate during the first stable period. On the other hand, emitter M had a significant rate during each ramp, and had multiple electrons per pulse. The detected phe and number of electrons per pulse for the first and last stable periods are shown in Figs. 6.20 and 6.21 respectively.

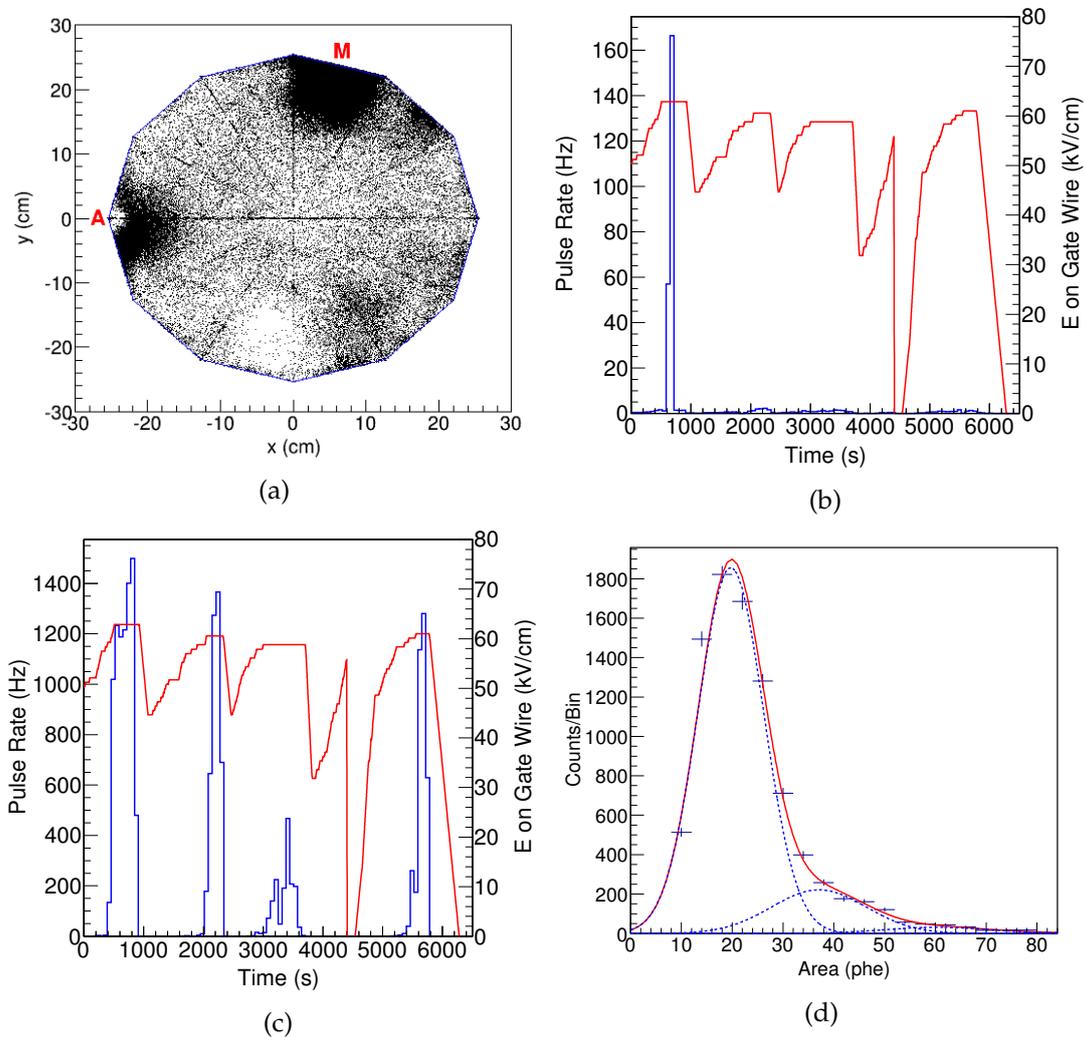


Figure 6.19: Emitters on Jan. 15th 2014 liquid test. (a) Reconstructed positions. Pulse rates (blue) and field on gate wire (red) for emitters (b) A and (c) M. (d) Detected phe for emitter A during the first stable period at 600 s ($|E_{wire}| = 63$ kV/cm).

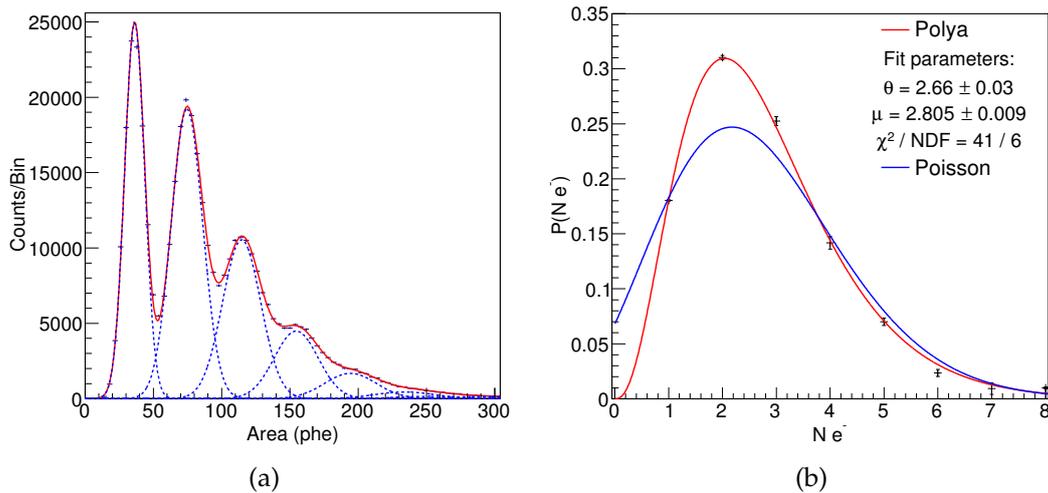


Figure 6.20: Emitter M during first stable time at 600 s ($|E_{wire}| = 63$ kV/cm) in the Jan 15th 2014 liquid test. (a) Pulse area. (b) Number of electrons per pulse.

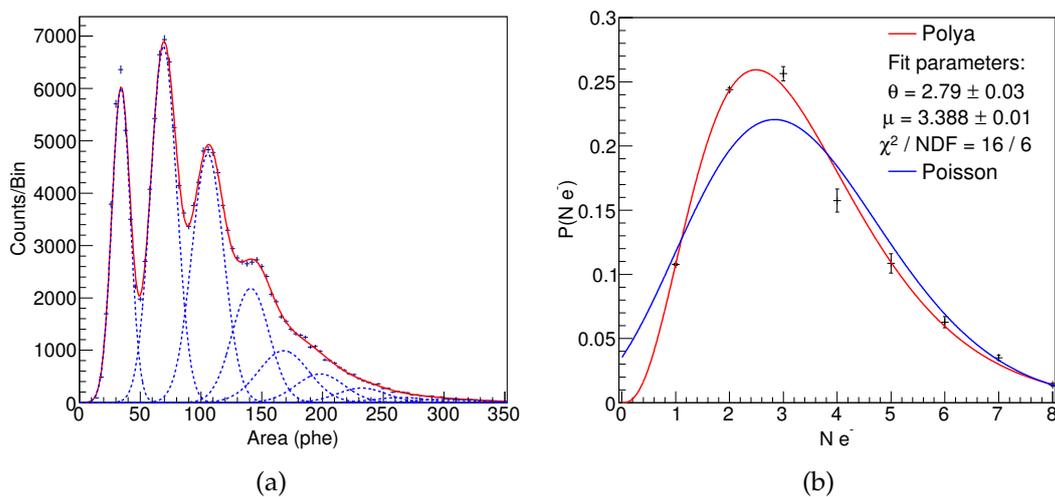


Figure 6.21: Emitter M during last stable time at 5500 s ($|E_{wire}| = 61$ kV/cm) in the Jan 15th 2014 liquid test. (a) Pulse area. (b) Number of electrons per pulse.

For emitter M, although each ramp went to a similar field, the mean number of electrons per pulse changed. It started with an average of 2.8 (Fig. 6.20b), but dropped to 2.4 and 2.5 for the second and third stable periods respectively. Then, in the last stable period there the mean increases to 3.4 electrons per pulse, as shown in Fig. 6.21. This is similar behaviour to emitter H in the gas.

More conditioning in the gas took place on the subsequent days, as described in section 6.2.1. The last liquid data were taken on Jan. 28th and 29th 2014. The electric field on the gate was almost 70 kV/cm during the stable periods and different behaviour was observed at these higher fields. The pulse length was not near to the expected 2 μ s for most pulses. An example waveform from the 28th is shown in Fig. 6.22. As well as the electron pulses there is a high rate of single phes, which could be the onset

of glow – continuous emission of single photoelectrons, attributed to thermoluminescence of the PTFE induced by the intense VUV radiation. With glow present the pulse finder can no longer split the electron pulses correctly. The light from the single phes is spread out evenly over the arrays as shown in Fig 6.23. Therefore, it cannot be due to localised electron emission.

The locations of the emitters are still visible, as shown in Fig. 6.24 for the Jan. 28th 2014 test. The absolute value of the pulse rate is not useful because many counts are missed due to overlapping pulses. Two small patches of emission are visible, one of which appears to be the remnants of emitter M. Some small spikes in the emission rate occurred, but with a maximum measured rate of <5 Hz. This is a clear improvement from the liquid tests on Jan. 15th 2014. Emitter O is very similar. The main emitter here is N, which has a significant emission rate as shown in Fig. 6.24b. This was similar on the single ramp during the Jan. 29th 2014 test, where the maximum emission rate was 95 Hz during a stable period with a field of 71 kV/cm on the gate wires. Although these data cannot be used to investigate the number of emitted electrons per pulse, it does indicate that the pulse rate increases significantly just above 70 kV/cm on the gate wires.

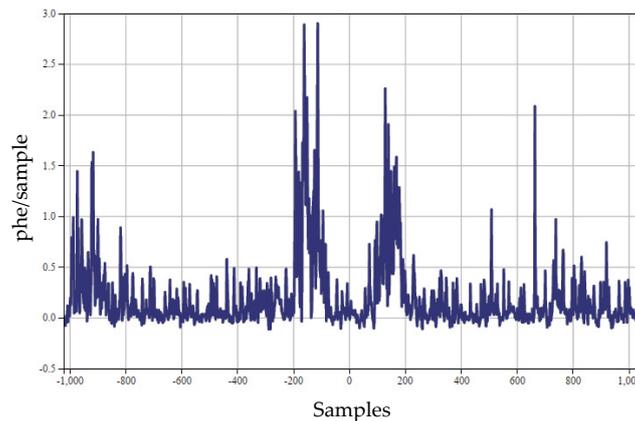


Figure 6.22: Example waveform from Jan. 28th 2014 data. The continued emission of single phes from glow stopped the pulse finder from properly separating the electron pulses.

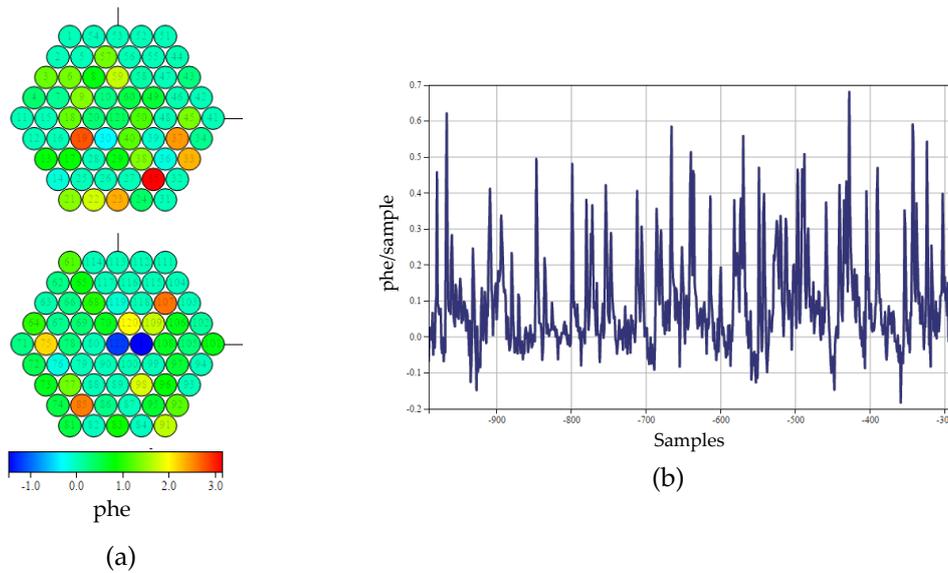


Figure 6.23: Example of glow from the Jan. 28th 2014 data. (a) PMT hitmap. The light is not localised to any particular area. (b) Typical waveform, multiple single phes.

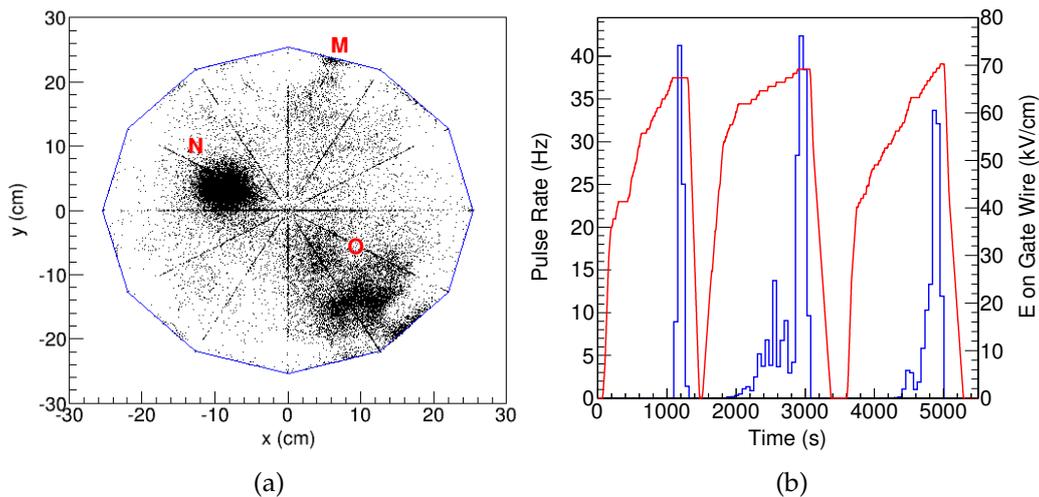


Figure 6.24: Emitters in liquid test on Jan. 28th 2014. (a) Reconstructed positions. (b) Pulse rate (blue) and field on gate wires (red) for emitter N.

6.3 Summary of Electron Emission Data

Although the behaviour of the emitters varied greatly there are some conclusions that can be made, it is now a good point to summarise what was found so far. Table 6.1 shows the location of each emitter, and at which stable periods it was observed in.

Emitters that release one electron at a time were observed in both the gas and the liquid data. In the gas data there are emitters with up to 700 electrons per pulse, but it is not possible to know how many are from the emitter at the surface or due to the high

field around the gate wire. In the liquid data the largest number of electrons released at once was 8. For the gas data, during stable grid conditions, the mean emission rate was 168 Hz and the average field on the gate wires was 81 ± 13 kV/cm. For the liquid data these were 800 Hz and 69 ± 6 kV/cm. The range of observed emission rates was large: from a few Hz up to 550 Hz in the gas data, and from a few Hz up to 3200 Hz in the liquid data. The mean emission rate during stable periods was 430 ± 340 Hz in the gas data and 1060 ± 950 Hz in the liquid data. The average field on the gate at emission onset was 77 ± 12 kV/cm and 58 ± 4 kV/cm for the gas data and the liquid data respectively.

The nature of the emission in the gas and liquid data was quite similar. Many emitters have short spikes in their emission rate, which then quickly subsides. There are several examples where emitters are removed, particularly in the first liquid data on Jan. 7th 2014 and the first gas data on Jan. 10th 2014. However, some emitters persist throughout many ramps. There are also examples where increasing the field reduces the emission rate, in both liquid and gas data. The emitters undergo several changes during the tests, as can be seen in Table 6.1 by looking at the pulse rate during different stable periods.

Differences between emitters were even seen when the macroscopic conditions were the same. One example from the liquid data is from the Jan. 9th 2014 test, which had two single electron emitters, but another that released up to 4. Similarly, the gas data on Jan. 10th 2014 had both single and multiple electron emitters. This shows that the microscopic conditions at each emitter location are different.

The distribution of the number of emitted electrons per event consistently fits a Polya distribution for both the gas and the liquid data. For the gas data there is always multiplication in the high field around the wire which is difficult to distinguish from what is happening on or near to the wire surface. In the liquid, however, the expected field is not high enough to cause electron multiplication, so the extra electrons have to come from the wire surface. It is also important to note that a Poisson distribution was not a good description for the number of emitted electrons. Therefore, the extra electrons are consistent with a multiplication process just at the wire surface.

Table 6.1: Summary of electron emitters during stable grid conditions. For liquid data, $N e^-$ from wire shows how many electrons come from at or near to the wire surface and the mean is shown in the next column. In the gas data there can be extra electrons from the emitter, or from the high field away from the surface. The electron drift simulations could determine whether it was a single electron at the wire surface or more, which is shown in $N e^-$ from wire. Mean $N e^-$ for the gas data includes all extra electrons. When the number of electrons from the wire is unknown it is due to electron losses on the walls.

Emitter	x (cm)	y (cm)	Date	Stable time (s)	Gas/Liquid	E on gate wire (kV/cm)	Mean pulse rate (Hz)	$N e^-$ from wire	Mean $N e^-$
J	-8.1	6.6	Jan. 7th	800	Liquid	60	335	1-6	1.94
K	-13.9	8.9	Jan. 9th	1600	Liquid	72	48	1	1.0
				1900		74	9.6	1-2	1.0
				2400		75	9.5	1-2	1.0
L	11.8	20.2	Jan. 9th	1600	Liquid	72	2156	1-2	1.0
				1900		74	2984	1-2	1.0
				2400		75	2387	1-2	1.0
D	-9.1	22.4	Jan. 9th	1600	Liquid	75	673	1-4	1.8
				1900		74	103	1-4	1.8
				2400		75	29	1-4	1.9
				7600	Gas	95	6.3	Multiple	110
				8600		98	210	Multiple	220
				7600	Gas	95	1.6	1	9.2
C	-10.7	9.2	Jan. 10th	7600	Gas	95	2.0	1	9.8
				7600	Gas	95	2.5	Multiple	65
E	-6.6	11.9	Jan. 10th	7600	Gas	98	5.3	Multiple	21
				8600		98	2.3	?	3.2
F	22.3	-6.8	Jan. 10th	7600	Gas	95	4.1	?	5.2
A	-23.5	-0.4	Jan. 10th	7600	Gas	95	35	1-3	1.07
			Jan. 15th	600	Liquid	63			
H	-2.1	-12.7	Jan. 12th	400	Gas	65	389	Multiple	26
				550		66	524	Multiple	20.0
			Jan. 13th (M)	3400	Gas	71	161	Multiple	28.9
				4300		66	227	Multiple	16.3
			Jan. 13th (E)	500	Gas	65	493	Multiple	22.1
				1600		75	401	Multiple	11.8
			Jan. 14th	900	Gas	79	52	Multiple	42.8
			Jan. 17th	500	Gas	70	396	Multiple	27.6
			Jan. 22nd	800	Gas	82	245	Multiple	71.3
				1500		75	126	Multiple	30.0
			2300		75	96	Multiple	28.1	
100	Gas	65	9.9	?	1.2				
M	6.2	23.5	Jan. 14th	600	Liquid	63	1215	1-7	2.38
			Jan. 15th	2000		61	1004	1-6	2.07
				3000		59	165	1-6	2.14
			5500		61	845	1-8	2.89	

6.4 Search for Light Emission from Wires in Liquid Data

Now we describe the search for any photon emission from the wires associated with the electron emission. With the liquid above the gate, any light production from when the electrons are released from the wire can be separated from the EL production in the gas. The drift time between the gate and liquid surface is 1-2 μs , depending on the liquid level. Therefore, the search is for small photon signals just before the electron pulses.

Electron pulses are selected in the same way as for the previous analysis. To search for photon signals an isolation cut is applied, which requires there to be no pulses identified as a single electron or S2 within 8 μs before the start of the electron pulse. The quiet time is significantly larger than the expected search window for the photons to ensure that any tails of pulses are not included, and to allow for estimation of the background as explained later on.

The pulse finding efficiency for signals of only a few phe is low in these data. There are many pulses in each event and large pulses are prioritised over smaller ones. Therefore, to search for the photon signal the analysis is done on the evt files, which contain the whole waveform. To find whether there is a photon signal the waveforms from each emitter are aligned by the start time of the each electron pulse. Then, the average is taken in each time bin to get an averaged waveform. Examples are shown in Fig. 6.25. The peak at the start of the electron pulse is an artifact of the alignment as there has to be phe to define the start of the pulse, but afterwards there can be small breaks in the emission.

For some emitters there was a clear peak just before the average electron pulse. The time between the photon and electron was either 0.8 μs or 2.0 μs . The expected drift time between the gate and the anode for the nominal liquid level is approximately 1.8 μs (with some dependence on the electric field). The 2.0 μs drift time is consistent with this, and is likely slightly longer due to the non-uniform field and electrons not taking a straight path to the surface (for example if the emitter was on the side of the wire). The shorter drift time implies that the liquid level was lower at that time. This is consistent with the level sensor readings on the slow control, which were lower than their nominal values. Therefore, the photon emission is associated with the electron emission from the gate. This is the first time it has been observed in a liquid noble element detector. However, the photon peak was not seen for all of the emitters.

A further selection cut is applied based on the number of electrons in each event. This was a 1σ selection around the mean of each peak in the electron pulse area histogram for each dataset. With this the number of photons and electrons emitted in each event can be investigated. For finding the pulse area in the photon window, a signal region and three background regions are selected, as shown in Fig. 6.25. The signal region is based on the location of the peak in the average waveform. Having a long isolation cut allows three well-separated background regions to be found, which are taken to be the same width as the signal region. The pulse area is found by integrating the waveform within each region. For emitters where the photon peak was not visible in the average waveform the signal region is taken to be the same as that of other

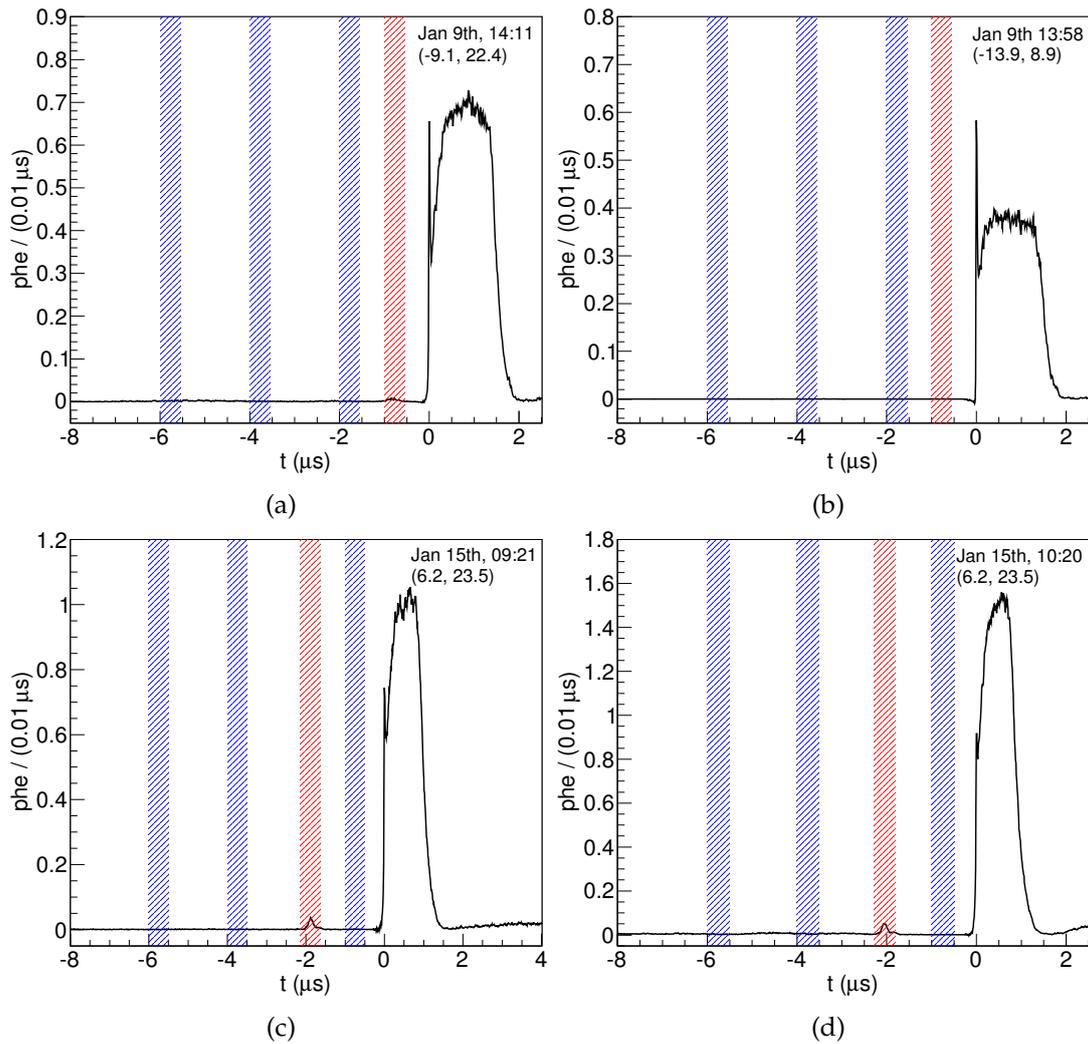


Figure 6.25: Averaged waveforms from liquid data with the signal (red) and background (blue) regions. (a) Emitter D on Jan. 9th 2014, 2400 s after start. (b) Emitter K on Jan. 9th 2014 at 1600 s. (c) Emitter M on Jan. 15th 2014 at 2000 s. (d) Emitter M on Jan. 15th 2014 at 5500 s.

emitters on that day. For each number of electrons the phe area is plotted with the background subtracted.

To find which photon distribution can reproduce the observed number of phe a toy Monte Carlo simulation is carried out. The probability of emitting N photons is taken as an input. For each trial a number of photons is sampled from this and the photon detection efficiency is applied. The pulse size of a detected phe is determined from the background regions in data, which gives a Gaussian with a mean of 1.1 phe and width of 0.25 phe. The detection efficiency from the wire surface was determined from Geant4 simulations in LUXSim. Photons are generated from a toroid around the gate wire, at the x - y location of the particular emitter. The detection efficiency ranged from 6% for emitters close to the walls to 8% for emitters towards the centre. The detection efficiency did not vary significantly with the angle around the wire. It is also important to account for the probability of one photon creating two phe in the PMT, which occurs 20% of the time [255]. The dominant effects are the photon detection efficiency and two

phe emission probability. Even if several photons are emitted, it is very unlikely that more than one or two are detected, and so there is no sensitivity to the underlying photon distribution.

On the other hand, there is some sensitivity to the mean number of photons emitted. For the following simulations the mean number of photons was taken to be constant. The mean number of photons that gives the best agreement with the data is found by a simple χ^2 test between the two phe histograms. During the tests on Jan. 9th 2014 there were two emitters which only released one electron at a time. The data and simulation are shown in Fig. 6.26. There are a few events with phe in the signal region (even after background subtraction), but the lowest χ^2 was with zero emitted photons. Simulations with the mean number of photons slightly increased are also shown for comparison, and can be taken as an upper limit.

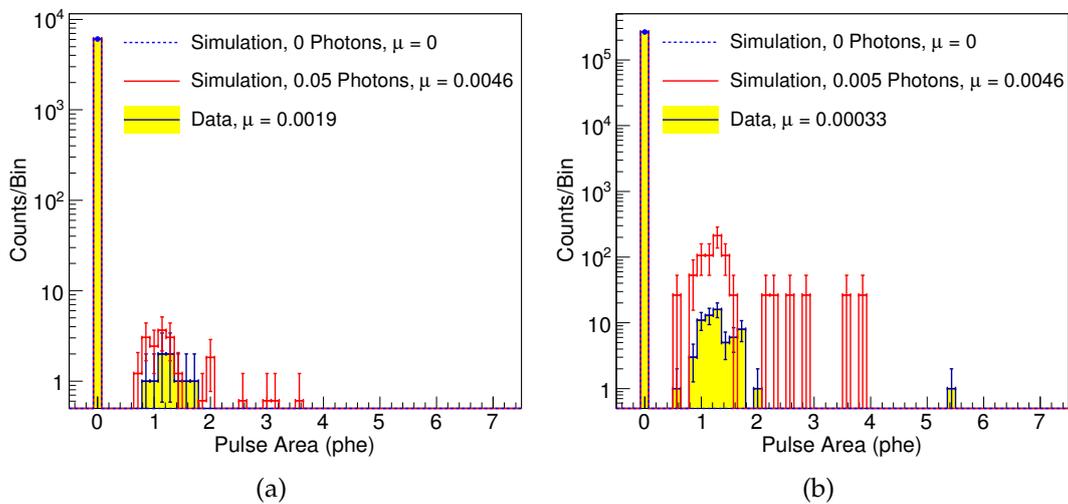


Figure 6.26: Number of detected phe from the gate for Jan. 9th 2014 emitters during stable time at 1600 s, which almost exclusively emit one electron at a time. The mean number of detected phe, μ , is shown for both the data and the toy Monte Carlo simulations. Emitter (a) K and (b) L.

Example phe histograms from multiple electron emitters on Jan. 9th and Jan. 15th 2014 are shown in Figs. 6.27 to 6.29. The simulations reproduce the shape of the observed phe spectrum reasonably well, which confirms that the photon detection efficiency and two-photon emission are the dominant effects. For the two multiple electron emitters the best fit number of photons is found for each number of electrons emitted for each of the stable times. The results are shown in Fig. 6.30.

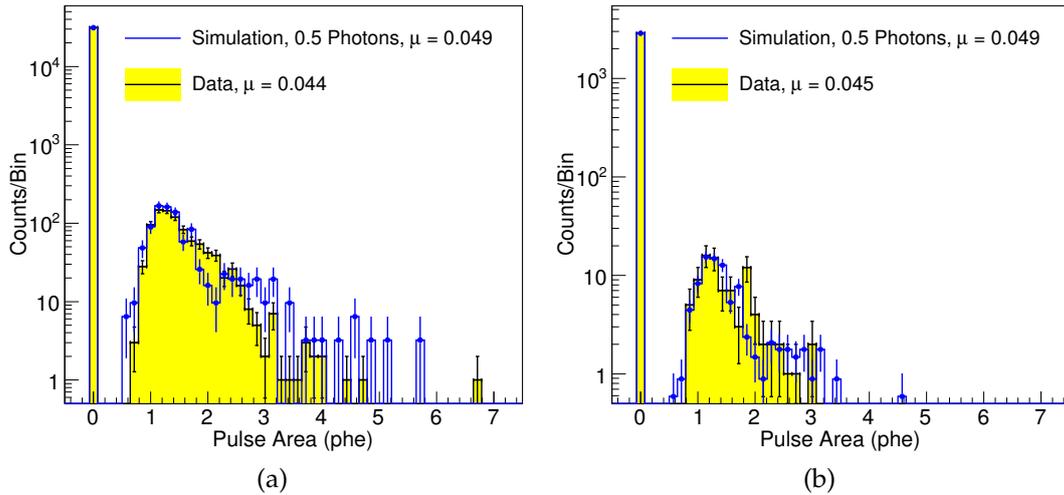


Figure 6.27: Number of detected phe from the gate for emitter D on Jan. 9th 2014, during stable time at 1600 s. This emitter releases between 1 and 4 electrons at once. The mean number of detected phe, μ , is shown for both the data and the toy Monte Carlo simulations. (a) Phe histogram when only 1 electron was emitted. (b) Phe histogram when 4 electrons were emitted.

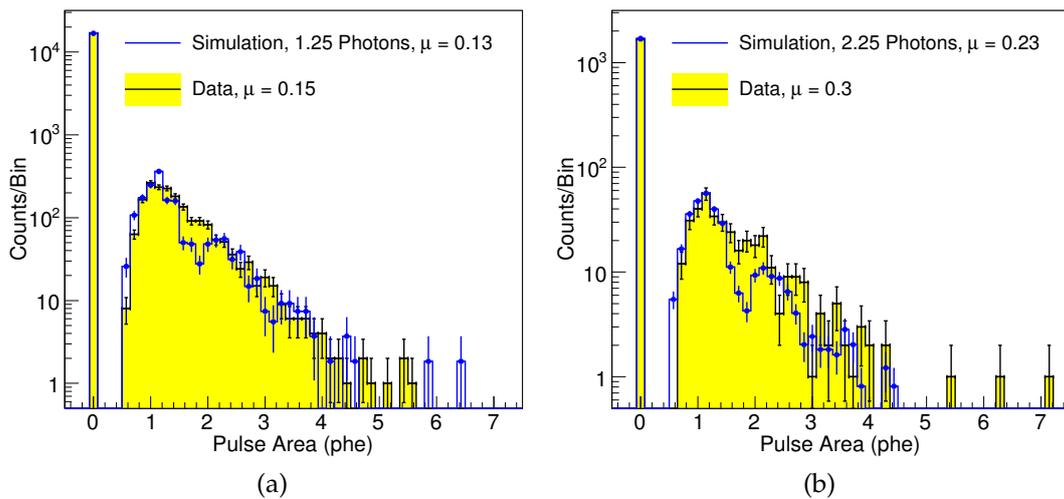


Figure 6.28: Number of detected phe for emitter M on Jan. 15th 2014, during stable time at 3000 s. This emitter releases up to 8 electrons at once. The mean number of detected phe, μ , is shown for both the data and the toy Monte Carlo simulations. (a) Phe histogram when only 1 electron was emitted. (b) Phe histogram when 5 electrons were emitted.

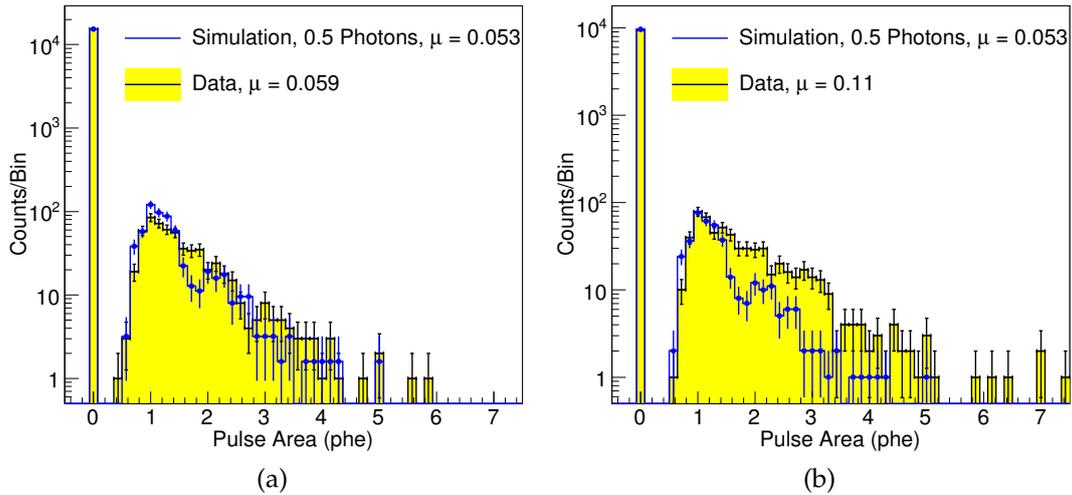


Figure 6.29: Number of detected phe for emitter M on Jan. 15th 2014, during stable time at 5500 s. This emitter releases up to 8 electrons at once. The mean number of detected phe, μ , is shown for both the data and the toy Monte Carlo simulations. (a) Phe histogram when only 1 electron was emitted. (b) Phe histogram when 5 electrons were emitted.

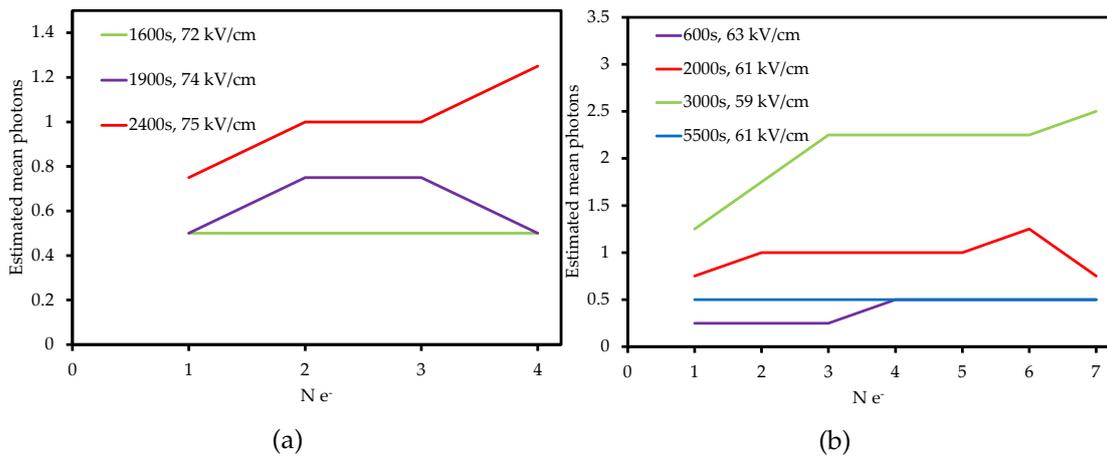


Figure 6.30: Estimated mean emitted photons against the number of emitted electrons during times with stable grid voltages. The mean is estimated to the nearest 0.25 photons. (a) Emitter D on on Jan. 9th 2014. (b) Emitter N on Jan. 15th 2014.

From Fig. 6.30 we see there is no clear correlation between the number of electrons and photons emitted. During most of the stable times the average number of photons emitted remains constant even as more electrons are emitted. There are two where there is a slight increase in the mean number of photons emitted. In general, the mean number of emitted photons is similar to the mean number of electrons and this is an important conclusion in itself.

For the single electron emitters there is very little, if any, photon emission. If they had emitted a similar number of photons per electron as the multiple electron emitters, which was at least 0.1 photons per electron, this would clearly be visible in the phe

histograms in Fig. 6.26. There are a few detected phe in the signal region even after background subtraction, and so there is a small probability to emit a photon.

6.5 Discussion: Which Emission Mechanism?

Firstly, we consider the possibility that the emission is caused by a locally enhanced field. This could cause scintillation and multiplication in the xenon just around the emitter, which would be consistent with the observation of a few photons and electrons released at a time. The Polya distribution, which is used to describe electron multiplication processes, fits the number of emitted electrons well even for the liquid data.

If an enhanced field is caused by an extrusion from the wire then the Fowler-Nordheim theory should be applicable. The current of electrons emitted from the wire can be calculated for a given field. An approximate form following Ref. [277] is:

$$R_e(E) = \frac{A}{e} \frac{1.54 \times 10^{-6} E^2}{\phi} \exp\left(\frac{-6.83 \times 10^9 \phi^{3/2}}{E}\right), \quad (6.3)$$

which gives the number of emitted electrons per second as a function of electric field at the emitter E in V/m; ϕ is the work function of the metal in eV, e the electron charge in Coulombs and A is the area of the emitter in m^2 . The Fowler-Nordheim theory has been successful in describing emission in the ideal case of a metal tip cathode, for example in Refs. [298] [299]. The Fowler-Nordheim theory describes a straight line when plotting $\ln(R_e/E^2)$ against $1/E$.

Equation 6.3 is now used to estimate the field required at the emitter to give the observed emission rate. The work function is that for the metal-xenon interface. For liquid xenon, the vacuum work function must be modified by the energy of the electronic conduction level V_0 of the liquid, so the effective liquid work function is $\phi_l = \phi + V_0$. The work function for stainless steel is 4.4 eV, and $V_0 = -0.61$ eV for liquid xenon at 165 K [300]. Then, the enhancement factor $\beta = E/E_{wire}$, where E_{wire} is the calculated field on the surface assuming a perfectly smooth wire. For stainless steel, emitters have been reported with areas between 1 nm^2 and $100 \mu\text{m}^2$ in Ref. [278] (which estimates the area of the emitter from measurements at varying electric field using a Fowler-Nordheim plot). In Ref. [279] two emission sites were observed and investigated with a scanning electron microscope, which found areas of 0.5 nm^2 and 13 nm^2 . For the observed emitters in the LUX data, β was between 100–400 depending slightly on the unknown emitter size A . This implies that fields between 100 MV/cm and 250 MV/cm are present at the emitter.

For a protrusion from the wire the enhancement factor β depends on the geometry. It has been calculated for a hemispheroidal [301] and cylindrical [302] projections from the surface. These are shown in Fig. 6.31. The enhancement factor depends on the ratio of the width and height.

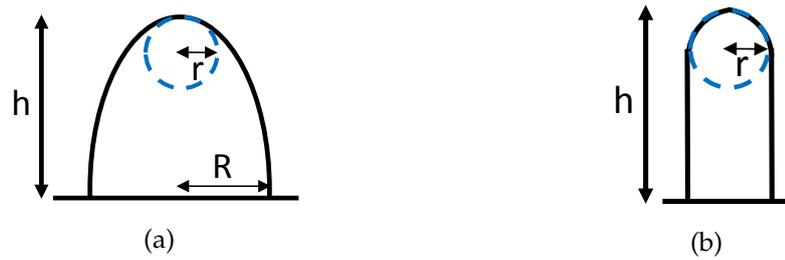


Figure 6.31: Parameters for (a) spheroid and (b) cylindrical protrusions from wire surfaces.

The β factors calculated using equation 6.3 for the LUX emitters have been found in other studies. Experiments are able to inspect the metal surface where an emitter was observed with an electron microscope. Studies of field emission from stainless steel in vacuum in Refs. [278] and [279] observed protrusions with β s between 100 and 700 at emitter sites. In Refs. [283], [303] and [304] cylindrical and spheroidal protrusions from stainless steel were found with diameters of $\sim 1 \mu\text{m}$ and height $\sim 20 \mu\text{m}$, and β between 30–300. Similarly, Refs. [280] [281] and [282] found 0.1-2 μm tall cylindrical protrusions at emitter sites on various metal surfaces, including stainless steel. However, the protrusions are never observed before sparking or discharge took place. In the LUX data there was no sparking before the first test, but emitters were still present. Therefore, this disfavours projections from the wire causing an enhanced field.

It is possible that the field enhancement could be from conducting particles on the surface, which have been observed with $\beta \sim 100$ [284] [285]. This would still follow the Fowler-Nordheim behaviour. However, there are reasons to disfavour field emission from any locally enhanced field. Equation 6.3 can be used to calculate the ratio between the emission rate in the liquid and in the gas. This removes the uncertainty due to the emitter area A . The average field during stable times was 81 kV/cm in the gas and 69 kV/cm in the liquid, which is multiplied by β . For the liquid the work function is lowered by V_0 . The ratio $R_{\text{gas}}/R_{\text{liquid}}$ is calculated for the minimum and maximum β observed, given the average fields. For the minimum β of 100, the rate in the gas is predicted to be 100 times larger than that in the liquid. When going to a β of 400 the ratio decreases to 4. However, the average emission rate in the liquid is 800 Hz, compared to 168 Hz in the gas, and so the observed ratio is less than 1. Therefore, this is not consistent with the expected behaviour for field emission.

The strongest evidence against field emission is the change in emission rate as the field is increased. Most of the emitters in the LUX data have a peak in emission rate at a certain field, then the rate decreases again as the field is increased further. For field emission the emission rate should increase sharply as the field is raised. A possible explanation is that the emitters are being “burnt off”, by loss of material or heating. However, the emission rates are relatively small, particularly for the liquid data. For a rate of 100 electrons per second, the power dissipated is $\sim 10^{-13}$ W. This is a tiny amount, which would not do anything to an emitter with dimensions on the μm scale. Even for the gas data where rates of 1000 Hz were observed, $\sim 10^{-12}$ W is still small. These energies cannot cause any heating or bubble formation in the liquid, so cannot

explain why the rate drops within only a few minutes. This suggests that the emission cannot solely be field emission.

As the emission appears to be enhanced at particular fields, it could suggest a resonant process. Tunnelling through an absorbate with an energy level near to the Fermi level of the metal is possible [290]. However, it is not clear whether this process would give electron multiplication or scintillation.

A possible emission process is the temporary or permanent electroforming of conducting filaments in insulating materials, which takes place at a particular voltage across the impurity [305] [306]. For some insulating impurities there are observations of electroluminescence associated with the electron emission process from copper (with the emitter in a vacuum) [307] [288]. The wavelength of the emitted light was between 500 and 900 nm with one or more peaks. The filament behaves like a metallic whisker in a dielectric, which gives a field enhancement at the end of the tip [289]. The scintillation with emission into vacuum was attributed to electron scattering in the insulating material, but with an emitter in xenon this could also take place in the high field region around at the tip.

Another mechanism is that the applied electric field can penetrate through a dielectric impurity to the metal, and distort the band structure as shown in Fig. 6.2. The penetrating field is able to accelerate electrons in the impurity, which can give electron multiplication or scintillation [289].

A likely source of insulator is an oxide layer on the metal. The nature and quality of this layer is possibly an important factor for field emission, as there are often imperfections in the surface quality or grain boundaries. These features could create surface states which are populated or depopulated at certain fields. The electron multiplication and photon emission could come from a process such as those described above. Alternatively, there could be a small field enhancement from a surface defect in addition to the oxide [291]. The β factor alone may not be enough to induce a significant field emission current, but even a β of ~ 10 – 20 could give a field above the electron multiplication threshold in liquid xenon (725 kV/cm [308]).

The electron emission process was also investigated using a small dual-phase xenon test chamber at Imperial College. This consisted of a single PMT, with an anode grid, gate grid, and single cathode wire. The voltage on the cathode is ramped up to observe the field emission. The cathode wire sample can be changed, allowing different materials and wire treatments to be tested. Similar behavior to the LUX tests was found; emission onset is observed, but the emitters disappear as the field increases further. It suggests some kind of resonance effect. In these tests there was certainly no sparking or discharge, which further disfavours projections from the wire. In addition, chemical passivation of stainless steel wires was found to significantly reduce electron emission. This treatment involves using an acid solution to remove the surface iron, exposing the other alloy components to air so that they form a protective oxide layer. This creates a very smooth oxide layer that would remove emitter sources from imperfections or grain boundaries.

6.6 Conclusion

The emission process was found to emit a few electrons and a few photons at once, producing a similar number of each. Resonant effects in the emission rate are observed. This strongly disfavours field emission from a locally enhanced field, as the rate should increase sharply with increasing field. Although it is not possible to tell exactly what mechanism causes electron emission from this study, it is possible that impurities on the wire, or an oxide layer are responsible. The electroluminescence could come from the emission process in the impurity itself. The impurities are often associated with other wire defects, and so the emission could be a combination of a resonant effect in the impurity along with a smaller field enhancement, which would give electron multiplication in the xenon.

Chapter 7

Design of the LZ Electroluminescence Region

The performance of the electroluminescence (EL) region was investigated in order to obtain a suitable design for LZ. The impact of the main design parameters on the EL response is assessed by simulating electron drifts in the various scenarios. This work contributed to the LZ Conceptual Design Report [197], where acceptable designs were presented, and later to the Technical Design Report [130], where changes were made to the gate-anode configuration due to other constraints besides the EL response.

The EL region is the gas gap between the liquid surface and the anode. A high electric field (~ 10 kV/cm) extracts electrons from the liquid and drifts them towards the anode; as they travel they collide with and excite xenon atoms in the gas, with the subsequent photon emission creating the S2 signal. The yield for saturated xenon vapour is of the order of 1000 photons per cm at typical operating pressures [167]. Electric fields are usually applied by wire grids held at large potential differences. Various electrode configurations are possible, some of which are shown in Fig. 7.1. In LUX [155] a very fine woven mesh was used for the anode and coarser parallel grids for the others. In XENON100 all electrodes were hexagonal meshes [309]. ZEPLIN-III had a different design, where the anode was a solid plate and there was no gate grid [310]; this was instead of a top array of PMTs, so photons were reflected here rather than being detected.

The EL response should have a low variance in the photon yield, otherwise both discrimination and energy resolution can be degraded. The mean yield should also be large, which helps to detect signals with few extracted electrons and improves the relative variance. The overall performance of the TPC can be affected by limitations in the EL region design. For example, with the wire grids in Fig. 7.1 electrons can have different path lengths to the anode wires and experience different fields on the way, which means that the photon yield depends on the vertex position in the liquid. Also, high fields near the wires can cause ionisation and increase the variance in photon yield.

For LZ, the discrimination between electronic and nuclear recoils should be high in order to reject ERs from solar-pp neutrinos and from backgrounds in the LXe (e.g. radon progeny). The conservative baseline estimate for LZ is an average of 99.5%. In previ-

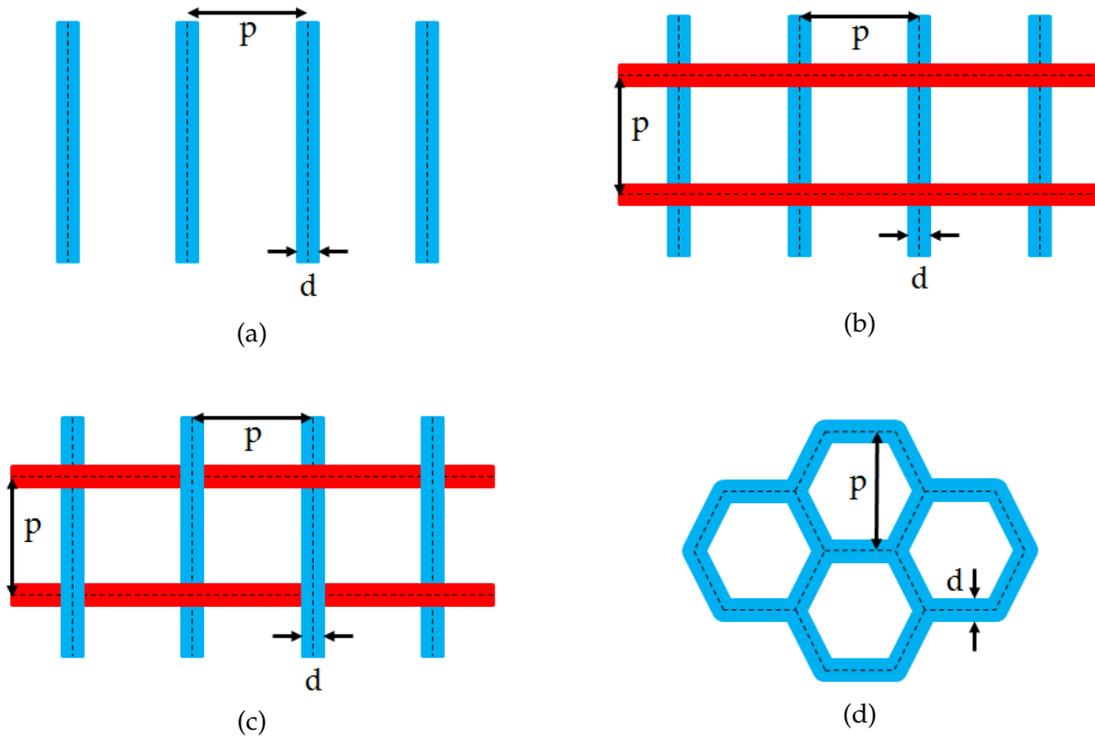


Figure 7.1: Common wire grid geometries described by a pitch p (wire centre to wire centre) and wire diameter d . (a) is a parallel wire grid, (b) a crossed wire grid (one set of wires on top of the other), (c) a woven wire mesh and (d) a hexagonal mesh.

ous experiments, ZEPLIN-III achieved a discrimination of 99.99% [177], higher than the 99.8% from LUX [178] and 99.75% in XENON100 [151]. This may partly be attributed to the high electric field in the drift region of ZEPLIN-III of nearly 4 kV/cm, compared to 0.53 kV/cm in XENON100 and 0.17 kV/cm in LUX. However, there are other differences between these detectors that could affect the discrimination. The planar geometry in ZEPLIN-III gives a very uniform field which must have intrinsically lower variance in S2 photon production than a wire anode. There is no light obscuration from the gate or anode or grid deflection to consider. Other differences are that ZEPLIN-III did not purify the xenon by recirculating during operation (hence providing a potentially quieter liquid surface), and the shorter active region gives less diffusion of electrons in the liquid. In the WIMP search region where the signals are relatively small, the variance in the yield is usually dominated by stochastic fluctuations in the number of detected S1 photons, which cannot be reduced for a fixed light collection efficiency. Therefore, although the EL design should have little impact on the resolution when compared to statistical fluctuations, it is important to ensure that this is the case.

Energy resolution is also important at high energies. In neutrinoless double beta decay (NDBD) searches, the signal is expected to be a peak at 2.458 MeV, the Q-value of the ^{136}Xe decay. This is near to background peaks from ^{214}Bi at 2.448 MeV and ^{208}Tl at 2.614 MeV, so the resolution affects the number of background events that can be rejected. Therefore, the half-life sensitivity is very dependent on the resolution that can be achieved. Also, improved resolution of γ -ray photo-peaks at MeV energies enables

a better understanding of backgrounds for WIMP searches. For large signals, fluctuations in the number of quanta are relatively small, so the intrinsic detector effects can have a significant impact on the resolution.

The EL response is investigated by simulating single electron drifts through the gas. We utilise software which simulates the microscopic transport, giving the location and timing of any excitations or ionisations in the gas phase. The electron drifts can be used for studying the single electron signal, and larger signals can be created by combining many such tracks.

Firstly, the various considerations to be accounted for are described in section 7.1. Many of these motivate conflicting requirements, so a compromise must be made. In section 7.2 the effect of the energy resolution at high energy is discussed, in relation to the NDBD searches. Sections 7.3 and 7.4 outline the simulation methodology and their validation, respectively. Next, simulations of existing experiments (LUX, ZEPLIN-III and XENON100) are presented in section 7.5, which include further validation against experimental data. Simulating each of the designs allows the effect of the wire geometry on the resolution to be investigated. Once a particular type of grid geometry is adopted for LZ, different possible combinations of pitch and wire diameter are simulated as described in section 7.6. After an acceptable solution is found for small signals, the variance of higher energy events is investigated in section 7.7.

As the studies progressed, at various stages there were changes to the design of the gate-anode system. Initially, the gate-anode distance was 1 cm with a 0.5 cm gas gap and the potentials were -3.5 kV on the gate and $+4$ kV on the anode. The initial studies were conducted with these pre-CDR parameters. This was soon changed to -4 kV on the gate, which gives a small change to the electric fields (the new voltages deliver approximately 10 kV/cm in the gas) and was adopted for the CDR model. Later, for a pre-TDR scenario, the gas gap was increased to 1 cm (with a 1.5 cm gate-anode distance) in order to give a larger EL yield. The potentials on the grids were increased to keep the field in the gas at approximately 10 kV/cm. The subsequent studies include the CDR and pre-TDR configurations. The baseline eventually adopted in the TDR an intermediate scenario with the same gas field and an 0.8 cm gas gap (1.3 cm gate-anode distance). More details on each configuration is presented in section 7.6.

7.1 Design Considerations

7.1.1 Photon yields

One of the main parameters for the S2 signal is the overall gain: how many photons are emitted per electron extracted from the liquid. This should be sufficiently large so that the signal-to-noise ratio is high as even the smallest signals from single electrons are important. Another benefit of high gain is to help with position reconstruction in the horizontal plane. This is particularly important for events from decays originating on the walls of the detector. These events undergo charge loss on the walls, giving a reduced S2 signal [311] [254]. In addition, the generated S2 light is more likely to undergo reflection on the walls in the gas phase. Therefore, the position of these

events is sometimes mis-reconstructed towards the centre of the TPC. Having higher photon statistics in the top PMT array helps to reduce this spatial bias. Given these constraints, a minimum photon yield of 50 photoelectrons per electron was set as a requirement (compared to ~ 30 photoelectrons in LUX). With the initial estimate of the photon detection efficiency of 9% this corresponds to a minimum of 550 photons per electron, which was the yield for the first simulations of the CDR scenario studied here. The estimate of the photon detection efficiency was later revised to 7.5% [197]. In the pre-TDR scenario a larger yield of 1190 photons/electron was considered [130] due to the other benefits discussed here, so a scenario which gives over 50 photoelectrons per electron with the lower photon detection efficiency is also simulated.

Although a larger photon yield is useful, there is an upper limit due to saturation of the PMT readout. Apart from WIMP searches, there are larger signals that must be reconstructed accurately. Calibration sources are used with energies of hundreds of keV. It is important that these do not saturate the PMT systems too significantly, as this would affect the calibrations in position and energy reconstruction. For LZ, the requirement is that non-linearities for the ^{83m}Kr calibration source should be less than 1%, for both S1 and S2 from interactions anywhere within the TPC. For higher energy signals, such as background γ -rays at MeV energies and NDBD, many PMTs in the top array will saturate even with a gain of 550 photons per electron.

There are other challenges associated with achieving a large photon yield that may be more restrictive than the PMT or DAQ saturation. The yield depends on the electric field strength, gas pressure and length of the gas gap. Data show that the yield is a linear function of the field [167] [168]. In initial scenarios with a gas gap of 0.5 cm and a gas pressure of 1.6 bar (as described in [197]), to obtain 550 photons per extracted electron requires 10.4 kV/cm in the gas. Increasing the electric field introduces difficulties related with high fields near both the anode and gate wires and at the edge of the TPC (see section 7.1.4).

Alternatively, increasing the drift length can help, and eases some mechanical constraints from grid deflection. However, it can be problematic for S2-only searches. These are searches below the S1 threshold, which means that the drift time cannot be used to determine the vertical position of interactions in the detector. Instead, some depth information can be obtained from the S2 pulse width, because it contains information on the amount of longitudinal diffusion in the liquid. A longer drift time in the gas introduces more gas diffusion, masking the liquid diffusion.

7.1.2 Extraction efficiency

The extraction efficiency (also known as the emission probability) is the probability of an electron that reaches the liquid surface to be emitted into the gaseous phase. This parameter depends on the the electric field in the liquid and the height of the potential barrier at the surface. Ideally, this should approach 100%, which gives larger S2 signals while removing a source of variance and the possibility of charge accumulation at the liquid surface. Typically, measurements from [312] are used. The extraction efficiency asymptotes above approximately 6 kV/cm in the liquid, and it is assumed that this represents unity emission probability, although this is currently unclear. In LUX, the

extraction field was 2.9 kV/cm during Run 3, which corresponds to an extraction efficiency of 50%. For LZ the aim is to achieve an extraction efficiency of over 95%, from a field of 5.6 kV/cm. Again, the same challenges of achieving a high electric field affect this quantity and are discussed in section 7.1.4.

7.1.3 Resolution of S2 signals

Many factors contribute to the overall variance of S2 signals. Ideally, the variance should be limited by statistical fluctuations, and the impact of the EL region design should be negligible. Table 7.1 describes the different stages in producing S1 and S2 light from an interaction. We consider that electronics effects downstream of the PMTs are sub-dominant.

Table 7.1: Processes contributing to the variance of S1 and S2 signals. Stages 3, 4 and 5 only affect S2 signals, the others affect both S1 and S2.

	Stage in Event	Source of Fluctuations
1	Production of primary quanta	Statistical fluctuations in N_{ex} and N_i
2	Recombination	Probability r for e^- to recombine with ions
3	Liquid transport	Xenon purity (electron attachment)
4	Extraction into gas	Extraction efficiency
5	S2 photon production	Statistical and E field in gas (shape and magnitude)
6	Photon transport	Reflection/absorption on surfaces
7	Photon detection	PMT quantum efficiency
8	Photoelectron gain	Multiplication gain fluctuations in PMT

The first stage is the interaction of the primary recoil, where ionisation and excitation of xenon atoms are produced (step 1 in Table 7.1). There are fluctuations in the number of excitations N_{ex} , ionisations N_i and the division between them. Subsequently, there is an additional probability that an ionised electron can recombine with a neighbouring ion replacing the electron with a new photon (2), as described in section 3.1. These processes are fundamental to the interaction, and their variance can only be reduced by increasing the electric field. After recombination, the number of ionisation electrons N_e and prompt photons released from the interaction site is determined. Next, electrons are transported through the liquid (3) and some can be lost by attachment to electronegative impurities in the xenon. If the purity is high, variations in electron losses can be reduced, and can be corrected for in data. Another source of variance is the extraction efficiency (4), which was mentioned previously. After extraction comes the production of the EL photons (5), the main topic of this study, which is discussed below in more detail. The next two stages are photon transport (6) and detection (7). Photons can be lost due to absorption on surfaces such as the PTFE walls (reflectivity $\approx 97\%$ [130]) and grid wires (reflectivity $< 50\%$ for stainless steel [170]) or due to absorption in the liquid (we expect the absorption length to be $\gg 10$ m in LZ). Afterwards, one must consider the quantum efficiency of the PMTs. These photon transport effects can be collectively accounted for by a binomial probability to detect a photon, estimated to be 7.5% for LZ [130]. Grid design can affect this because the wires can block the light, and this is discussed further in section 7.1.5. The last step is the

conversion of photons into an electronic signal in the PMTs (8). The single photoelectron response fluctuates, as measured by the width of the photoelectron distribution, and this is the last source of S2 signal variance which we consider.

In LZ, the requirement on the combined S1+S2 energy resolution is a maximum of 2% at 2.5 MeV. This is set in order to enable a NDBD search and to perform γ -ray spectroscopy to inform background models. In this chapter the S2-only resolution is investigated, which is a component of the combined energy resolution.

For the S2 photon production, the variance is heavily influenced by the gate-anode region design. In an ideal case, for a perfectly uniform electric field, the only variance would be from statistical fluctuations in the number of produced photons, intrinsic to the random walk problem, which is small. However, this ideal case is technically difficult to achieve. The anode and gate are usually wire grids or meshes, to allow light to pass through to the top PMT array. At the anode the field is focused onto the wires, meaning that it is not uniform and can be large enough to cause electron multiplication near to individual wires. The first Townsend coefficient of ionisation is 1 electron/mm at ≈ 35 kV/cm in gaseous xenon at 1.6 bar [169], [313]. Any electron multiplication will inevitably increase the variance in photon production, as secondary electrons also produce EL (although only for short track lengths). Another related effect is that the path lengths electrons take can vary, depending on where they come from below the anode, so the photon yield is position dependent. This effect can be significant even for large signals of many electrons. Additionally, the path length and electric field may vary across the horizontal plane due to deflection of the grids due (mostly) to electrostatic attraction.

The overall effect on the resolution can be quantified. The mean S2 size is given by:

$$\mu_{S2} = \bar{N}_i(1-r)(1-a)\epsilon \bar{N}_{EL} p \bar{A}, \quad (7.1)$$

where N_i is the number of initial ionisations:

$$\bar{N}_i = \frac{E_R}{W}(1 + N_{ex}/N_i). \quad (7.2)$$

The excitation-to-ion ratio is taken to be a constant equal to $N_{ex}/N_i = 0.2$ for electron recoils; W is the average energy required to create a quantum, r is the recombination probability, a the average probability of electron attachment in the liquid, ϵ the electron extraction efficiency, N_{EL} the number of electroluminescence photons per extracted electron, p the photon detection efficiency (in the gas phase) and \bar{A} the average PMT gain. The relative width of the S2 distribution can be found by adding the (statistically independent) individual contributions in quadrature:

$$\begin{aligned} \left(\frac{\sigma_{S2}}{\mu_{S2}}\right)^2 &= \left(\frac{\sigma_i}{\bar{N}_i}\right)^2 + \left(\frac{\sigma_e}{\bar{N}_e}\right)^2 + \left(\frac{\sigma_t}{\bar{N}_t}\right)^2 + \left(\frac{\sigma_x}{\bar{N}_x}\right)^2 + \\ &\quad \left(\frac{\sigma_\gamma}{\bar{N}_\gamma}\right)^2 + \left(\frac{\sigma_{phe}}{\bar{N}_{phe}}\right)^2 + \left(\frac{\sigma_G}{\bar{G}}\right)^2. \end{aligned} \quad (7.3)$$

A similar approach was followed in [314] [315]. Here, \bar{N} refers to a mean, and σ the standard deviation. The subscript i refers to the initial number of ionisations in the interaction, e to electrons that escape recombination, t is for electrons that have been transported to the liquid surface (not lost to impurities), x is for extracted electrons, γ for EL photons and phe for detected photoelectrons. \bar{G} is the sum of the individual PMT gains.

The first term arises from microscopic fluctuations in the number of quanta created and is described by the Fano Factor, $F = \sigma_i^2 / \bar{N}_i = 0.059$ for liquid xenon [316]. Fluctuations in the number of electrons is entirely due to the recombination process. This does not have the expected standard deviation for a binomial process [178] [207] but instead was found to scale with N_i :

$$\sigma_e = C\bar{N}_i, \quad (7.4)$$

where C is an empirical constant, determined to be $C = 0.067$ from LUX data. Although recombination does depend on electric field, there is not expected to be a large effect until the field is much higher than that of LUX (170 V/cm) or LZ (300–700 V/cm). Therefore, the difference in recombination between LZ and LUX should be small.

The second term in equation 7.3 becomes:

$$\left(\frac{\sigma_e}{\bar{N}_e}\right)^2 = \left(\frac{C\bar{N}_i}{(1-r)\bar{N}_i}\right)^2 = \frac{C^2}{(1-r)^2}. \quad (7.5)$$

The recombination probability r is taken from the value obtained from FastNEST simulations, where it is tuned to give the best agreement with experimental data.

The next stage is transport of electrons through the liquid. The probability of attachment at a drift time t is given by $P(t) = 1 - e^{-t/\tau}$, where τ is the free electron lifetime. The average probability is given by integrating over all drift times, with $\Delta t = t_{max} - t_{min}$:

$$\begin{aligned} a &= \frac{\int_{t_{min}}^{t_{max}} (1 - e^{-t/\tau}) dt}{\int_{t_{min}}^{t_{max}} dt} \\ &= \frac{\tau (e^{-t_{min}/\tau} - e^{-t_{max}/\tau}) + \Delta t}{\Delta t}. \end{aligned} \quad (7.6)$$

Then this is taken as a binomial probability for an electron not to reach the surface. With the probability of reaching the surface $s = 1 - a$, the mean and variance on the are given by:

$$\bar{N}_t = s\bar{N}_e, \quad (7.7)$$

$$\sigma_t^2 = s\bar{N}_e(1-s) = \bar{N}_t(1-s), \quad (7.8)$$

The electron attachment term can then be written as:

$$\left(\frac{\sigma_t}{\bar{N}_t}\right)^2 = \frac{a}{\bar{N}_t}. \quad (7.9)$$

Similarly, electron extraction and photon detection are binomial processes. Following the same treatment gives:

$$\left(\frac{\sigma_x}{\bar{N}_x}\right)^2 = \frac{1-\epsilon}{\bar{N}_x}, \quad (7.10)$$

$$\left(\frac{\sigma_{phe}}{\bar{N}_{phe}}\right)^2 = \frac{1-p}{\bar{N}_{phe}}. \quad (7.11)$$

For the EL light production, the total number of photons \bar{N}_γ comes from many different electrons. Then, $\bar{N}_\gamma = \bar{N}_x \bar{N}_{EL}$ where \bar{N}_{EL} is the mean number of EL photons produced for each extracted electron and the standard deviation $\sigma_\gamma^2 = \bar{N}_x \sigma_{EL}^2$. The EL term becomes:

$$\left(\frac{\sigma_\gamma}{\bar{N}_\gamma}\right)^2 = \frac{1}{\bar{N}_x} \left(\frac{\sigma_{EL}}{\bar{N}_{EL}}\right)^2. \quad (7.12)$$

Similarly, we adopt a constant mean PMT gain \bar{A} and describe the width of the single photoelectron response by σ_A , which is considered equal for all PMTs. Using $\bar{G} = \bar{N}_{phe} \bar{A}$ and $\sigma_G^2 = \bar{N}_{phe} \sigma_A^2$ gives:

$$\left(\frac{\sigma_G}{\bar{G}}\right)^2 = \frac{1}{\bar{N}_{phe}} \left(\frac{\sigma_A}{\bar{A}}\right)^2. \quad (7.13)$$

Taking all of equations 7.10 to 7.13 allows equation 7.3 to be rewritten as:

$$\begin{aligned} \left(\frac{\sigma_{S2}}{\mu_{S2}}\right)^2 = & \frac{F}{\bar{N}_i} + \left(\frac{C}{(1-r)}\right)^2 + \frac{a}{\bar{N}_t} + \frac{1-\epsilon}{\bar{N}_x} + \\ & \frac{1}{\bar{N}_x} \left(\frac{\sigma_{EL}}{\bar{N}_{EL}}\right)^2 + \frac{1-p}{\bar{N}_{phe}} + \frac{1}{\bar{N}_{phe}} \left(\frac{\sigma_A}{\bar{A}}\right)^2. \end{aligned} \quad (7.14)$$

For LZ the expected parameters are: $a = 38\%$ (with a minimum $\tau = 843 \mu\text{s}$ corresponding to one drift length, $T_{min} = 54.6 \mu\text{s}$ and $T_{max} = 818 \mu\text{s}$), $\epsilon = 95\%$ and $p = 7.5\%$. For the PMT gain fluctuations, σ_A/A is approximately 38% for the PMT model in question.

The dispersion in EL light production is the quantity measured by the simulation studies in this chapter. In equation 7.14 one assumes the same Gaussian mean and width for all extracted electrons. Therefore, it does not account for the x - y position dependence, or a non-Gaussian EL distribution (eg. due to electron multiplication).

The effect of position dependence is discussed further in relation to the large signals for NDBD in section 7.2.

The resolution for LZ from equation 7.14 is shown in Fig. 7.2 for different σ_{EL} . The mean EL yield is taken to be 1190 photons per extracted electron from the TDR scenario. The electron lifetime is 5 ms, which was the maximum achieved in the EXO detector [201]. This shows that unless σ_{EL} is very large, the S2 resolution should not be affected. The dominant contribution to the S2 width is from the recombination fluctuations. However, due to the anti-correlation between S1 and S2 in recombination, when using the combined S1 + S2 energy scale recombination cannot contribute to the resolution. This is because if an electron recombines with an ion, it instead produces a photon, conserving the total number of response quanta.

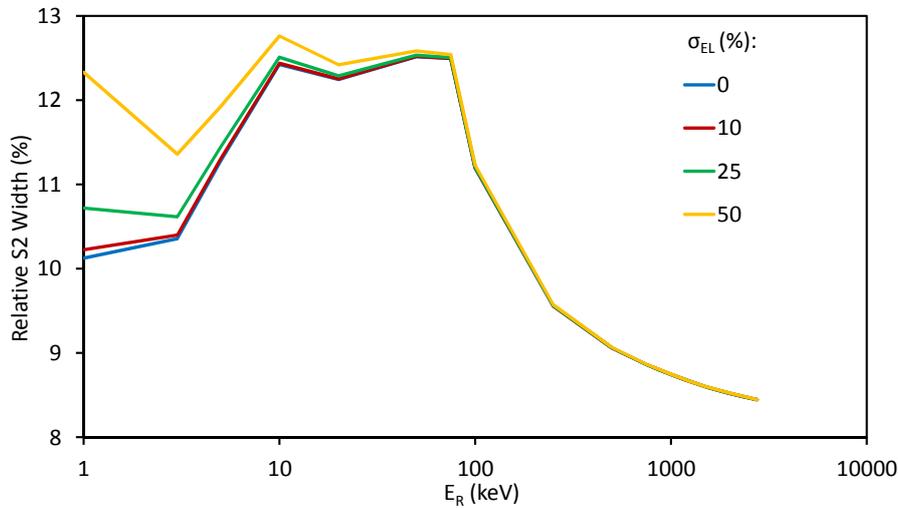


Figure 7.2: Effect of EL light production variance on overall S2 resolution for electron recoils. The relative S2 width (σ_{S2}/μ_{S2}) is calculated from equation 7.14 for various values of σ_{EL} with $\bar{N}_{EL} = 1190$ photons per electron and assuming the same EL mean and width for all electrons. The piece-wise behaviour at low energies are from changes in the recombination fraction (measured values [178]) which is the dominant contribution to the S2 width. Note that the resolution in combined S1 + S2 is much smaller than S2 only due to complete anti-correlation between S1 and S2 in the recombination process.

The effect of σ_{EL} on the discrimination for small signals in a WIMP search is quantified using FastNEST simulations. This is a Monte Carlo simulation that generates the S1 and S2 light for a nuclear or electronic recoil of a given energy. The typical (S2/S1) vs. S1 scatter plot can be created to measure the discrimination. The variance from the EL light production can be increased to quantify its the impact on discrimination.

Two datasets are generated in FastNEST, one for ER and one for NR, both with flat energy spectra (1 to 100 keV). The detector parameters are the same as used above. The discrimination is measured using the leakage fraction, defined as the number of ER events lying below the NR median divided by the total number of ER events in each S1 bin. This parameter is therefore proportional to the ER background of the experiment assuming a simple analysis. An example plot of the leakage fraction is

shown in Fig. 7.3. Next, Fig. 7.4 shows the average leakage fraction for $\sigma_{EL} > 0$ divided by that with $\sigma_{EL} = 0$.

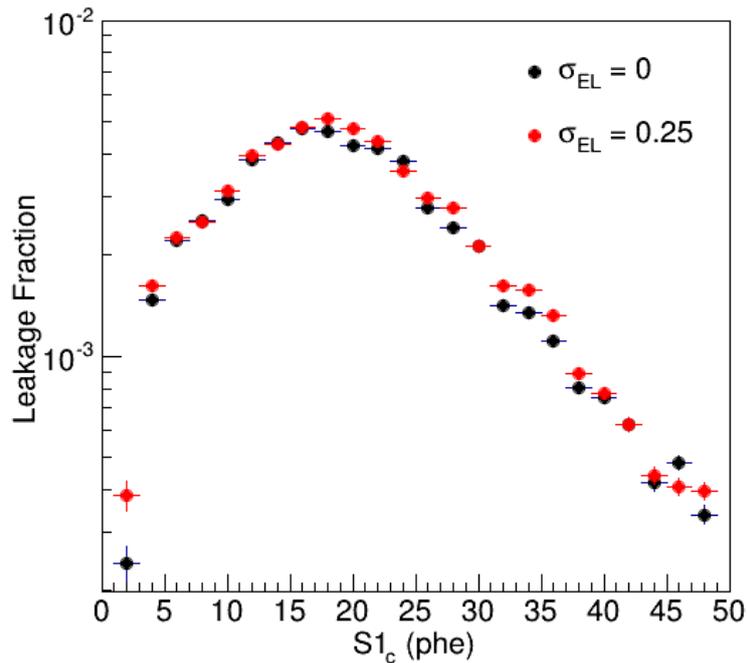


Figure 7.3: Leakage fraction from NEST for typical LZ parameters. EL light production is Gaussian with a mean of 1191 photons per extracted electron and a variable width. The plot shows the leakage fraction with zero width on the EL light production and with an increased relative width of 25%.

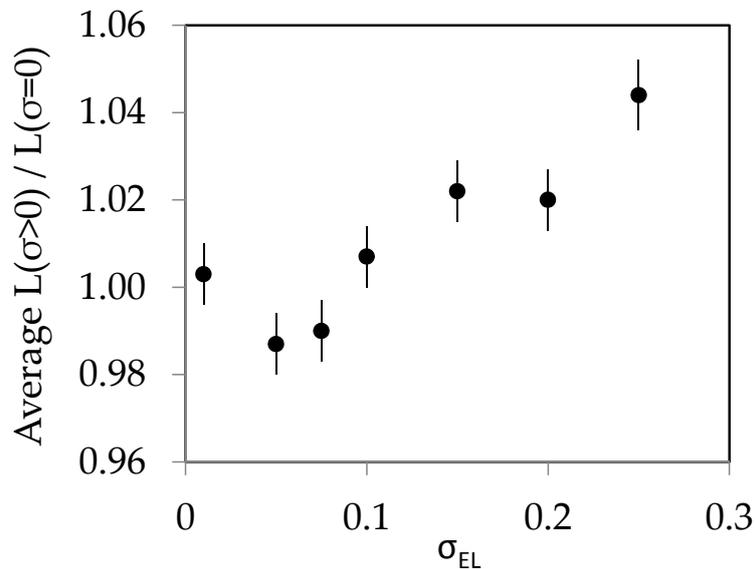


Figure 7.4: Mean value of the relative leakage fraction for various σ_{EL} .

The effect on the discrimination can be seen at $\sigma_{EL} = 15\%$ and above. At 10% and below no significant increase can be seen. Therefore, a safe upper limit on the EL light production variance from the point of view of discrimination is approximately 10%.

Even above this the effect is small, only approaching a 5% increase in leakage for a 25% relative width in the EL photon yield.

In conclusion, the dispersion in EL photon production should add but a small contribution to the total S2 resolution at low energy. However, note that we neglect the actual spatial distribution of the ionisation which is especially important for higher energy interactions, as will be discussed later. There is a small effect on the discrimination for $\sigma_{EL} > 10\%$. The treatment so far assumed a Gaussian distribution for the EL light production and did not account for position dependence. Therefore, it is important to check the effect with the EL yield distribution from the full simulation. The EL results can be incorporated into FastNEST code to simulate whole events accounting for the spatial distribution of energy deposits, including the generation of ionisation along the particle tracks.

7.1.4 Maximum electric fields on wire surfaces

A high electric field in the gas is required to generate an appropriate photon yield and high electron extraction efficiency. However, there are challenges associated with this. In previous noble liquid detectors the maximum field on the surface of cathodic wires was between 50-100 kV/cm and in LUX the maximum sustainable fields were only ≈ 19 kV/cm on the cathode and 46 kV/cm on the gate. Above these values, light and charge emission from the wires is observed. This effect is the main topic of Chapter 6, so here only the limitations on the LZ design are discussed. For the anode, the problem is instead the effect of the field on electrons drifting in the gas towards these wires. The high fields near the wire surface can cause ionisation, and hence more electroluminescence, which will degrade the S2 resolution.

The maximum field can be reduced by either decreasing the pitch of the grid or increasing the wire diameter. Both changes affect the optical transparency of the grid, which is a dominant driver of the LZ optical model.

For parallel wire grids, the electric field on the wire surface can be calculated once the field above and below the grid are known. The treatment follows [293], which takes the field as two components, a constant field plus one that varies around the wire. This gives:

$$E_W = \frac{a(E_a - E_b)}{2\pi r} + (E_a + E_b) \sin \theta, \quad (7.15)$$

where E_a and E_b are the electric fields above and below the wire grid (in the far field), a is the pitch, r the wire radius and θ the angle around the wire (measured from the horizontal plane of the grid). This equation also assumes that the wire radius is much smaller than the pitch, which usually holds true.

The upper limit on the maximum electric field on cathodic wires is taken to be 50 kV/cm based on limits observed in previous experiments and on the R&D ongoing at Imperial College to understand and mitigate this issue. This limit was adopted as a requirement [197].

7.1.5 Optical transparency

In order to maximise the photon detection efficiency, there should be minimal obstruction of light by the grids themselves. This affects the S2, but is particularly important for S1 signals. To increase the optical transparency requires a larger pitch or smaller wires. However, this is contradictory to the restrictions from the maximum electric fields and variance. This is where a compromise must be made, between having an acceptable optical transparency and a high uniform field. The transmittance at normal incidence (T_N) is used to quantify the optical transparency of the grid. For a single wire grid of pitch p and wire diameter d , T_N is given by:

$$T_N = \frac{p - d}{p}. \quad (7.16)$$

Other possible shapes are crossed grids, woven meshes or hexagonal grids. T_N for the crossed grids or woven meshes is given by:

$$T_N = \left(\frac{p - d}{p} \right)^2. \quad (7.17)$$

The hexagonal grid is also the same as the crossed/woven mesh, as the area of a hexagon with pitch p is $\sqrt{3}p^2/2$ and the constants cancel out.

In LUX, the transparency of the woven mesh anode was 80%, and for the single wire grid gate 98%. For LZ the requirement is that the combined transparency of all grids in the liquid is greater than 80% [130].

7.1.6 Mechanical and electrostatic considerations

LZ has a larger TPC diameter than LUX, 1.5 m compared to 0.5 m, which presents additional mechanical challenges that affect all of the other considerations. The wire tension must be large enough to combat electrostatic attraction and the weight of the grid. For the EL region, this creates non-uniformities across the horizontal plane of the detector as the gas gap will decrease with radius. In turn, this changes the electric field in the gas gap and therefore both the S2 pulse size and the pulse width. Although most non-uniformity can be corrected for from data by looking at the variation in S2 pulse width across the detector, it should not be too large. A high tension opposes these effects. For the gate-anode a maximum deflection of 2 mm at the centre of the detector can be tolerated [130].

The mechanical stability of the grid depends on the choice of grid geometry. Parallel wire grids are particularly bad due to displacement of individual wires and all of the tension being in one direction. This is not a good configuration for the manufacture of large grids as it would require extremely strong ring holders in the TPC. Also, if a wire were to fail and snap it would likely cause irreversible damage. A crossed wire grid, with two planes of orthogonal wires, reduces the ring deformation. However, the wire planes still repel each other. A woven wire mesh requires less vertical space

than the crossed grid, and less tension as the woven wires support each other through friction against deformation and electrostatic repulsion. It is possible that wires could be damaged during the weaving process, but overall this is a good configuration.

The rest of the properties are better than the other grid geometries. A chemically-etched grid (such as that used in XENON100) is robust and easier to attach to the ring holders, but difficult to manufacture on the scale of LZ. Woven or crossed grids also help against electron emission, as the electric field on the wire surfaces is reduced.

Although in LUX single wire grids were used for all grids except the anode, in LZ woven meshes are preferred for all grids. However, the very fine LUX anode mesh (254 μm pitch) cannot be used in LZ. The woven mesh was commercially available for the size of LUX, but a larger one for LZ is not. Therefore, it has to be manufactured by hand by the project. The mechanical constraints and limitations in manufacturing the grid mean the pitch has to be larger than in LUX. This could increase the variance in photon production, but will have a larger optical transparency. The pitch for LZ will be at least ≈ 2 mm. The wire diameter also needs to be larger so that the wires can withstand the higher tension required; the minimum considered for LZ is 75 μm .

Grid deflection affects various aspects of the detector performance as discussed above. The effect on S2 photon yield, electric field in the gas, electric field on the gate surface and S2 width are investigated after the simulation studies in section 7.6.4.

7.1.7 Other spurious effects

Using the tracking and timing information of electron transport in the gate-anode region the pulse shape can be reconstructed to check that the S2 response is acceptable. Non-uniformities in the field will have an effect on the pulse shape. For example, the high field near to anode wires could yield more photons per unit time at the end of S2 pulses.

Another possibility is that electrons could pass between anode wires, into the region above the grid. In LUX, as the pitch was very small, this did not occur and has not been observed in other xenon dark matter TPCs. However, for LZ the anode pitch is likely to be larger for the mechanical reasons discussed in section 7.1.6. Typically, in dual-phase xenon TPCs the field above the anode is below the EL threshold. If electrons go through the anode into this region they stop emitting photons for some time, before they come back down onto the anode wire and emit more photons. This would result in the pulse being split, with some light arriving after the main EL emission. It is important to check that there are no spurious light signals after the S2.

7.1.8 Summary of design considerations and limitations

Given the mechanical constraints, an EL region design that gives acceptable performance must be found, and here we summarise the most salient points from the above discussion. These apply especially to the small signals relevant for the WIMP search.

The manufacturing considerations limit the pitch to be greater than ~ 1 mm, with a wire diameter of over $75 \mu\text{m}$. The EL photon yield needs to be sufficiently large, of the order 1000 photons per electron, since a minimum of 50 phe per extracted electron is required. This is to improve the position reconstruction and performance in detecting small signals of only a few electrons. Upper limits on the photon yield are from the linearity of the PMT response, which is required to be linear to within 1% at the energy of the $^{83\text{m}}\text{Kr}$ calibration source.

The variance in photon production is likely to be a small contribution to the overall S2 variance. For a relative width above 10% there is a small effect on the average discrimination, but a lower resolution should be achievable. Previous calculations assumed a Gaussian distribution for the EL photon yield, which may not be the case if there is significant amounts of electron multiplication. The x - y position dependence of the yield was also not included. The grid pitch for LZ is larger than in previous liquid xenon TPCs, so it is necessary to check for these effects.

7.2 High Energy Signals: Neutrinoless Double Beta Decay

LZ may also search for neutrinoless double beta decay (NDBD). For xenon experiments, the liquid is both the source (8.86% ^{136}Xe) and detector medium. NDBD is expected to produce two electrons with a total energy equal to the Q-value of the decay, 2.458 MeV for ^{136}Xe [317]. On average this interaction creates 1.18×10^5 electrons in the liquid and most interactions create single scatter events. There are γ -ray backgrounds from ^{208}Tl (^{232}Th decay chain) at 2.614 MeV and ^{214}Bi (^{238}U decay chain) at 2.447 MeV. Therefore, a better energy resolution increases the separation between the signal and background peaks and improves the sensitivity to the NDBD signal. The predicted half-life sensitivity is between 1.8×10^{26} and 4×10^{26} years for a 2.0% or 1.0% energy resolution at the NDBD Q-value respectively [130]. This also depends on the achieved spatial resolution (better separation of vertices helps to distinguish multiple scatters) and background levels.

For large signals such as NDBD, the energy resolution should be good due to large electron statistics. However, detector effects can increase the resolution even at high energy due to the position dependence of the interactions which was discussed in section 7.1.3. Whereas in the previous section we considered point-like vertices, here the extended nature of the higher-energy interactions becomes a prime consideration. How this affects the resolution even for large signals is illustrated in Fig. 7.5. NDBD events have back-to-back electrons with tracks folded over approximately 3 mm when they are created. This is comparable to the ~ 1 mm pitch considered for the gate. Therefore, the means and variance in EL production that are being sampled by the electrons depend on the event position relative to the gate and anode wires, which was not accounted for in section 7.1.3. In the limit of multiple interactions spread over many gate-anode unit cells the S2 yield will sample from all positions.

For LZ the pitch of the gate grid will be comparable to the size of the vertex in NDBD. Therefore, this effect could be significant. One other consideration is that diffu-

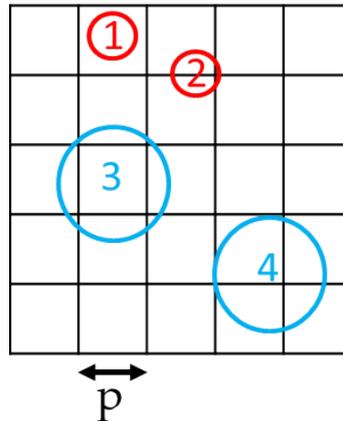


Figure 7.5: Effect of position dependence of interactions relative to the anode wires. The black lines represent the anode grid (of pitch p). The circles represent ionisation cloud of electrons in a single event. If the spread of electrons is smaller than the pitch, as in circles 1 or 2, then the paths the electrons take when drifting to the anode can vary significantly between interactions. Those in 2 are more likely to go straight to the anode wires. For circles 3 and 4, where the electrons are spread out slightly more than the anode grid, different parts of the unit cell are being sampled. Therefore, there could still be position dependence here, although less than for the case of circles 1 and 2.

sion in the liquid will cause a larger electron cloud, so this must be taken into account. The overall resolution is quantified by simulating whole NDBD events using FastNEST combined with the Garfield++ simulations for the electroluminescence yield. Diffusion in the liquid can be accounted for, so that the size of the electron cloud at the liquid surface is known, which depends on the depth of the interaction. As well as doing this for NDBD, it is also useful to look at lower energy interactions. The effect of position dependence could still be important for signals within the WIMP search range and the results can be compared to the FastNEST simulations with a Gaussian EL production variance as described in section 7.1.3.

7.3 Towards an LZ Design: Simulation Method

A combination of software is used to create the wire geometry and drift electrons through the gas phase. The electron transport is modeled using a Monte Carlo simulation, Garfield++ [294]. Some examples of electron drift lines are shown in Fig. 7.6. The electrons are tracked microscopically as they travel and locations of any excitations or ionisations are recorded. The mechanism of de-excitation in gaseous xenon (as described in section 3.1.3) means that an excitation can be assumed to produce one photon, and ionisations produce extra electrons which are also tracked. Garfield++ can calculate electric field maps for simple configurations where an analytical solution exists. This limits it to 2D geometries consisting of planes and parallel wires. However Garfield++ can read electric field maps externally generated by Finite Element software. To calculate the electron transport properties of the gas, Garfield++ is interfaced to the MAGBOLTZ software [318] which includes elastic and inelastic cross-sections

for many gases and mixtures, to calculate the trajectories, diffusion and gain of the electrons.

Apart from the parallel wire grid, the geometries considered here do not have 2D symmetry. For these other software is required to calculate the electric field. ELMER is a finite element method solver [319], which can calculate the field for almost arbitrary 3D structures. This requires a geometry and mesh to be created, which is done using GMSH [320]. A 3D unit cell is implemented and the boundary conditions set. GMSH then splits the region into sections, triangles in 2D and tetrahedrons in 3D, as shown for example in Fig. 7.7b. ELMER reads the mesh from GMSH, and calculates the electric field at the nodes of the mesh elements. An example of electric field contours from ELMER is shown in Fig. 7.7c. This ELMER field map is then read by Garfield++, which interpolates between the nodes to obtain the electric field at any location. The unit cell is also repeated in Garfield++. More details on creating each different geometry can be found in the following sections.

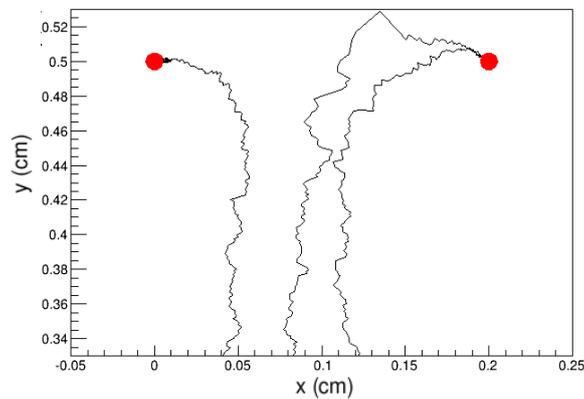


Figure 7.6: Example of three electron drifts from a liquid surface to anode wires (red circles). The line connects all of the points where excitations occurred.

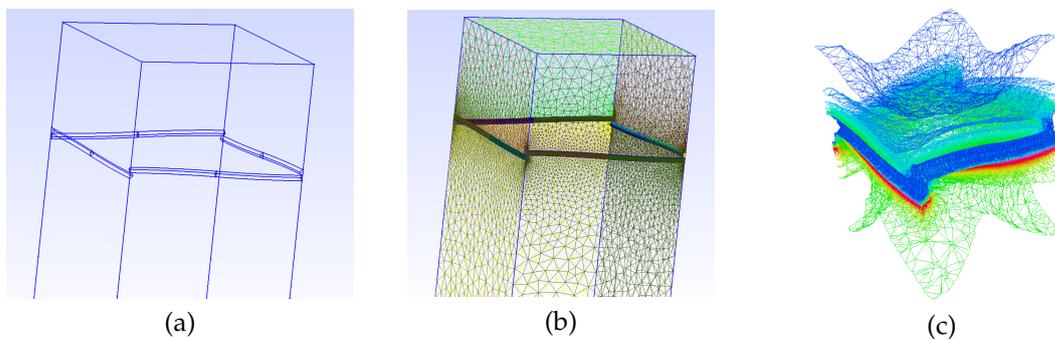


Figure 7.7: Example of steps required to create electric field map. (a) wire geometry in GMSH; (b) 2D meshing of the woven wires; (c) ELMER electric field contours.

7.4 Validation

The Garfield++ package has been extensively used for other applications. This includes energy resolution studies in NEXT, a gaseous xenon electroluminescence TPC for neutrinoless double beta decay [314], designing a proportional counter for the ALICE detector [321], and optimising Gas Electron Multipliers (GEMs) [322] [323]. The latter application involves electron multiplication as well as electroluminescence. Although this is reassuring, further validation against experimental data is carried out, since our application involves saturated vapour rather than room temperature gases.

The first test was to compare the number of electroluminescence photons produced per unit distance in the simulation to experimental data. There are several datasets for the electroluminescence yield in xenon gas, but the most reliable are those from our Coimbra collaborators who measured this yield in xenon vapour at low temperature, in equilibrium with the liquid phase, as a function of electric field [167]. Another widely used dataset is from Ref. [168], and that in Ref. [324] is another study of the electroluminescence yield obtained from Garfield++. There are differences between measured data from different authors, with the results usually being below the theoretically predicted values [169]. The measurement is technically challenging due to the difficulties in calibrating absolutely the light collection and detection efficiency at the xenon wavelength.

The experimentally measured photon yield from Ref. [167] is described by equation 3.3, which is a parameterisation from Ref. [170]. Fig. 7.8 shows the yield from equation 3.3, experimental data and simulation for two different pressures; 1.6 bar the operating pressure for LUX, ZEPLIN-III and possibly for LZ, and 2.2 bar the gas pressure used in XENON100 [309]. The pressure dependence is factored out by plotting the reduced electric field (in Td, where $1 \text{ Td} = 10^{-17} \text{ V}\cdot\text{cm}^2$) against the photon yield per cm divided by the number density of the gas.

This shows that the data and simulation agree to within $\sim 10\%$ over a wide range of fields. The simulation slightly overestimates the photon yield between ~ 8 to 15 Td. Given the potential systematics in the data, this agreement is acceptable. The onset of electroluminescence is also reproduced well. Deviations from the linear trend are seen in the Garfield++ data at 1.6 bar above approximately 17 kV/cm, which is due to the onset of electron multiplication. The Garfield++ data at 2.2 bar also agrees with the data well.

The electron drift velocity is also validated, as it is important for reproducing the pulse shape. Various measurements of the drift velocity in gaseous xenon are available and the agreement between measurements is reasonable, as shown in Fig. 7.9. The Garfield++ (bulk) drift velocity in the gas is measured by simulating multiple drifts at each constant field, and measuring the average time taken to complete the drift. This is compared to two of the more recent datasets from [325] and [171] (points reproduced from plot shown in [326]).

The agreement to the data is good over most electric fields. There is some disagreement at very low field, but this is below 100 V/cm at 1.6 bar. The minimum electric field in the simulations is over 100 V/cm, so the drift times will be reproduced well.

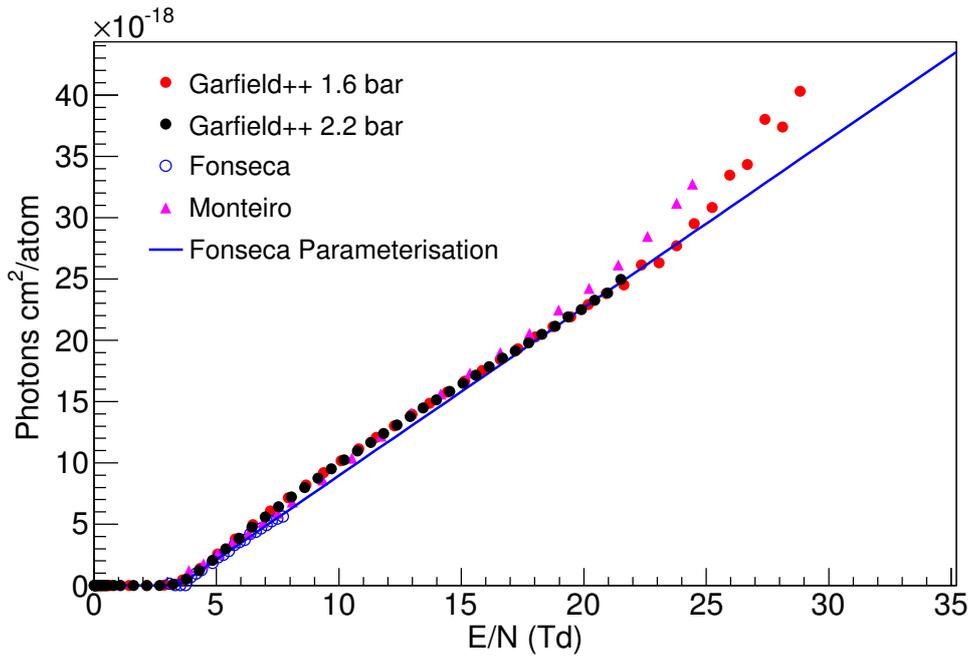


Figure 7.8: Photon yield as a function of reduced electric field for gaseous xenon vapour. Points from Garfield++ are plotted for 1.6 and 2.2 bar. It is compared to experimental data from Monteiro *et. al.* [168] and Fonseca *et. al.* [167]. The line is a parameterisation from [170] of the data in Ref. [167].

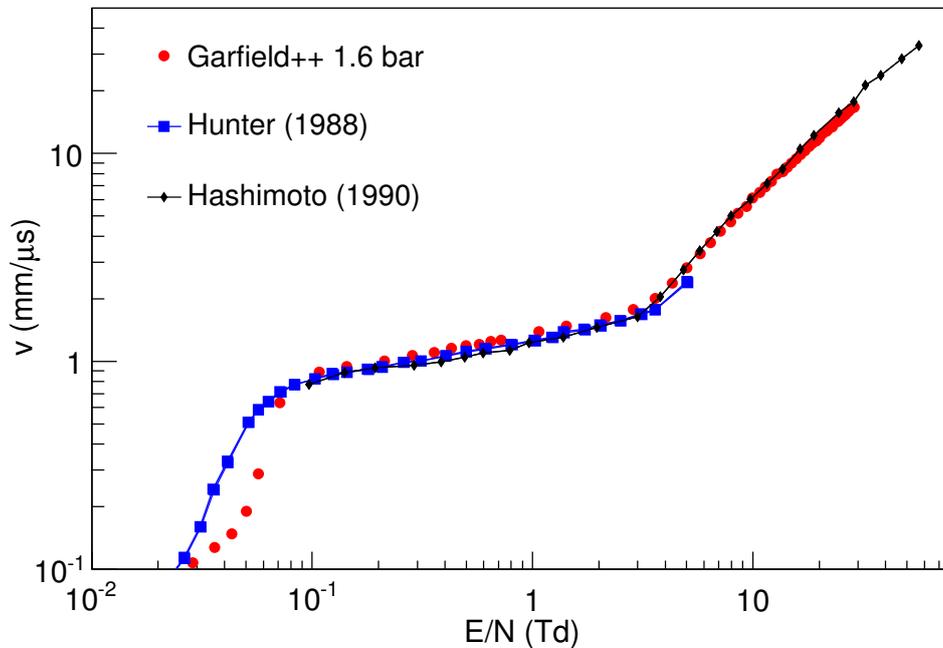


Figure 7.9: Drift velocity from Garfield++ against reduced electric field. Also shown are data points from Hunter *et. al.* [325] and Hashimoto *et. al.* [171].

7.5 Simulations of Existing Experiments

The first simulations are of the previous experiments (LUX, ZEPLIN-III and XENON100) to give guidance to the type of design LZ should adopt. The simulation results are also compared to experimental data as extra validation. To assess the impact of the field shape due to the anode geometry, a planar version of LUX and XENON100 were created. The field and other parameters (gas height, pressure) are set to be the same, to isolate the effects on the resolution from the anode design. The layout and electrostatic parameters of the experiments are summarised in Fig. 7.10.

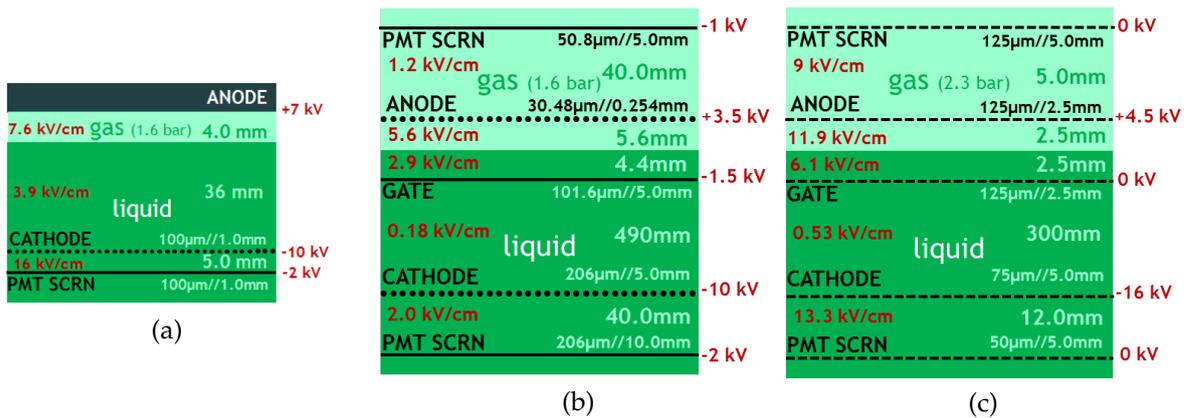


Figure 7.10: Geometry and parameters for the various dual-phase xenon experiments. The distances between each set of wires are shown, along with the wire diameter//pitch of each one. (a) ZEPLIN-III; the cathode and PMT screen are parallel wire grids, the anode is a solid plate [310]. (b) LUX; the anode is a woven mesh and the others are parallel wire grids [155]. (c) XENON100; all grids are hexagonal etched meshes [309].

7.5.1 Creating field maps for previous experiments

In this section the details of creating the electric field map and running the electron drifts are described for each experiment.

ZEPLIN-III

For ZEPLIN-III the entire simulation can be done in Garfield++. Due to the planar geometry, the field is assumed to be uniform (there is no gate grid in this 2-electrode design and the anode is a solid plate). Edge effects are not accounted for here. The simulated region is simply a box of gas at constant pressure and temperature. Electrons are drifted from the bottom to the top of the region in Garfield++.

LUX Run 3

The anode woven mesh in LUX is the example geometry shown in Fig. 7.7a. In order to build the wires in GMSH two semi-circular surfaces are created at two opposite corners. In Fig. 7.11a the corner at positive x and y is shown. Next, these are extruded along the arc of a circle with radius R :

$$R = \frac{p^2 + r_w^2}{2r_w}, \quad (7.18)$$

through an angle α :

$$\alpha = \sin^{-1}(p/R). \quad (7.19)$$

These quantities are determined from the pitch p and wire radius r_w . For the upper semi-circle, the point which anchors the rotation is at $(p, p, -R + r_w)$ and for the lower one $(p, p, R - r_w)$ (taking $z=0$ to be at the anode centre). This gives the two wire sections shown in Fig. 7.11b. Next, the ends of these wires are extruded again, but the points they are rotated around are different. For the top wire, the rotation is around $(-p, p, R - r_w)$ and for the bottom wire $(p, -p, -R + r_w)$. This is shown in Fig. 7.11c. The other corner is done in the same way. After the wires are made, the remaining part is a box for the gas. A plane at the top and bottom are made and joined to the wires. This leaves the final geometry shown in Fig. 7.11d. With this method, the pitch and wire radius can be changed to create woven meshes of any dimension.

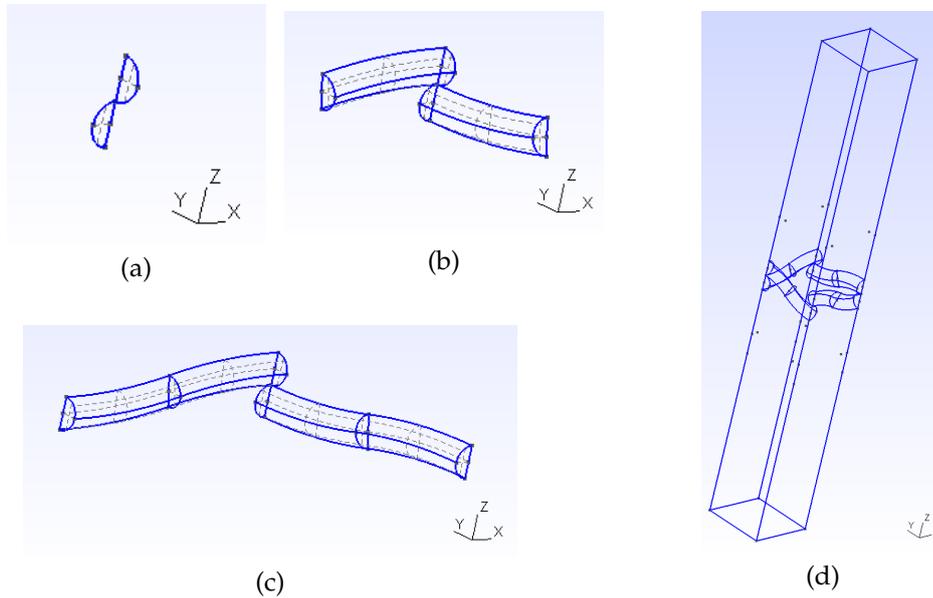


Figure 7.11: Method for making a woven wire mesh geometry in GMSH. (a) initial semicircles for creating the wires. (b) first extrusion for the wires. (c) second extrusion of wires. (d) adding outer box to complete the unit cell.

With the geometry implemented, the next step is to set the boundary conditions. The bulk electric fields in LUX are found using the calculation for the field around a single wire grid [293], extended to multiple grids [154]. This accounts for leakage of the electric field between different regions. This effect occurs because not all of the field

lines terminate on the grids, some pass into the next region. In the case of LUX, the anode has a very small pitch so there is minimal leakage of the field through this grid. The values of the fields away from the wires are shown in Fig. 7.10b. Once these are known, the potential at the liquid surface and above the anode can be calculated using $E = V/d$, which works here because the leakage through the anode is negligible. At the liquid surface the potential is 0.37 kV. Instead of using the entire distance between the anode and top PMT shield several cm away, the geometry is shortened. The field above the anode is uniform away from the wires, so a plane of constant potential is created 0.3 mm above the anode with the expected potential at that location (3.1 kV in this instance). The potential at the wire surface is fixed to be 3.5 kV. The symmetry for this unit cell (Fig. 7.11d) is a mirror symmetry, in both x and y .

The height of the unit cell is shortened to reduce the computational requirements on the simulation. The size of the mesh created from GMSH depends on the number of nodes in the geometry. A larger mesh takes longer to create, and increases the time taken for the Garfield++ drifts to run. One way to improve the performance is to have a varying mesh density. In GMSH, the density at a particular point in the geometry can be specified. As the field away from the wires quickly becomes uniform, a high mesh density is not required here. At the corners of the box, the density is set to be 0.2 mm but on the wires it is smaller. In between the density is gradually changed (automatically by GMSH) so that it reaches the values set at these points, as can be seen in Fig. 7.7b.

To find the optimal trade-off between computational efficiency and accuracy, the LUX field map was created using different mesh densities on the wire surfaces. Then the field along the wire surface is plotted and compared. Four different lines were chosen: one is at the corner of the unit cell (where the wires cross) and is a circle around the bottom wire. The other three are along the wire in the y direction, at the bottom, side, and top of the wire. Five different mesh element sizes on the wire surface were chosen between 50 and 0.5 μm . The results from each of these are shown in Fig. 7.12.

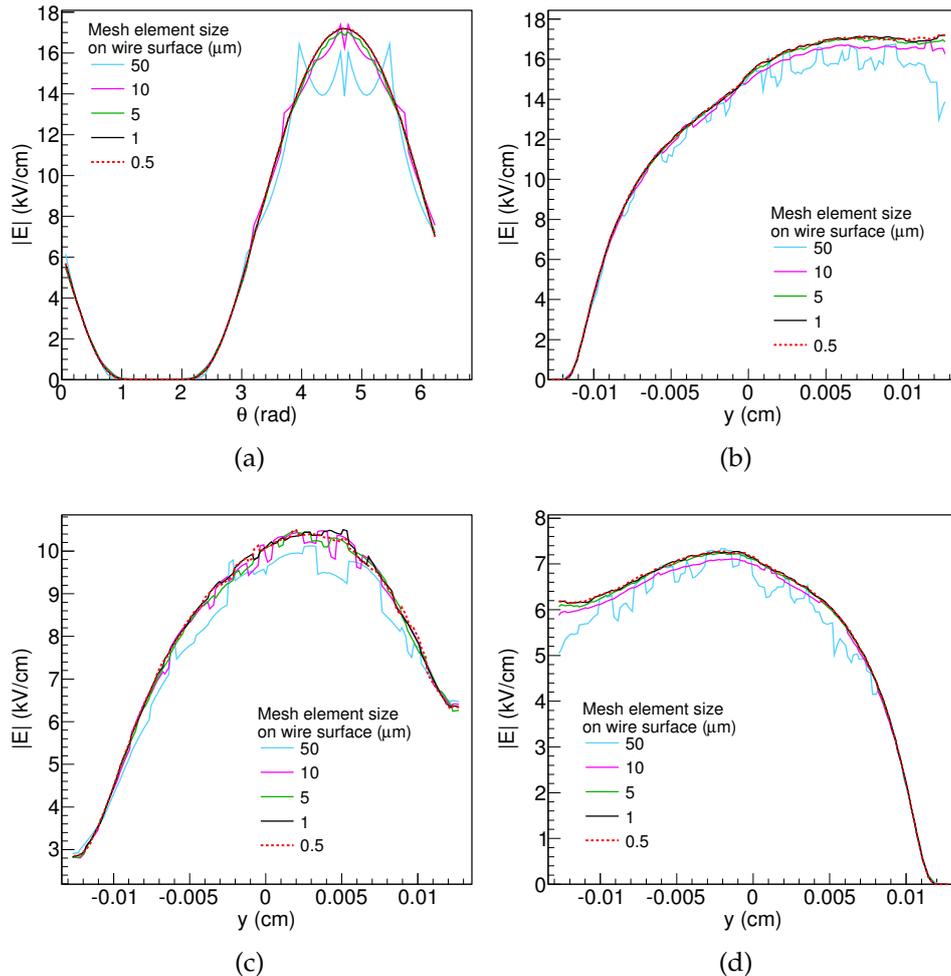


Figure 7.12: Electric field strength on wire surfaces for LUX with different mesh element sizes on the wire surfaces. (a) Around the bottom wire at the corner of the unit cell, θ specifies the angle from the horizontal. (b) Along the bottom of a wire segment (one side of the unit cell). (c) Along the side of a wire segment. (d) Along the top of a wire segment.

In Fig. 7.12 the three highest mesh densities agree well, but the lowest two show deviation from the other lines. Therefore, mesh element sizes of 5 μm and below are likely to be acceptable. Electron drifts will be simulated with each mesh element size to compare the effect on the photon yield and variance. Other examples of the field are shown in Fig. 7.13.

For all of the following electron drifts in the LUX simulations, the starting position is sampled at random from the bottom of the geometry (liquid surface). The effect of focusing from the gate wires is ignored here, as the gate pitch is much larger than the anode pitch, so many anode unit cells are being sampled. Therefore, the sampling of the position is uniform.

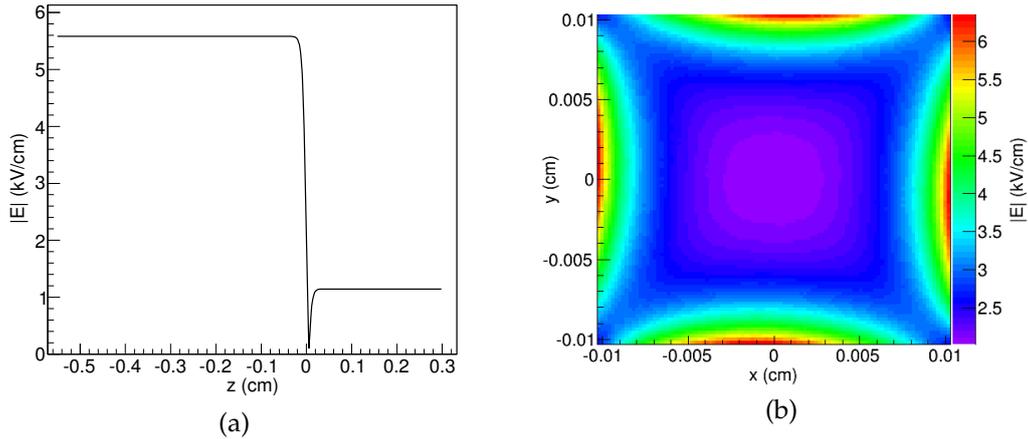


Figure 7.13: Electric fields in the LUX simulation. (a) Along the centre of the unit cell ($x=y=0$). The liquid surface is at $z=-0.56$ cm and anode is at $z=0$ cm. (b) In the x - y plane at the anode ($z=0$).

XENON100

For XENON100 the method for creating the geometry is different to that used for LUX. Here the gate and anode grids are the same geometry and pitch, so focusing in the liquid is important. Relative to the gate, the anode is offset by half the pitch. It is designed so that the electrons are focused onto the anode wires by the gate [309]. The arrangement of the gate, anode and top grids can be seen in Fig. 7.14. A triangular unit cell is used here, which includes all three grids. To make this, the geometry was imported to GMSH from a CAD drawing. The ability to import a triangular unit cell was not included in Garfield++ by default, but this was added by the Garfield++ authors at our request. We give special thanks to Joshua Renner from LBNL for implementing this extra feature.

Initially, the geometry was created with completely sharp-edged wires. As with the LUX scenario, the effects of the mesh density was investigated. Five different mesh element sizes on the wire were chosen for the square wires and the field plotted along the top, side and bottom of one of the anode wires. The details of the field on the anode wires is more important than the gate wires here, as electrons are focused away from the gate. The field on the wire surfaces is shown in Fig. 7.15.

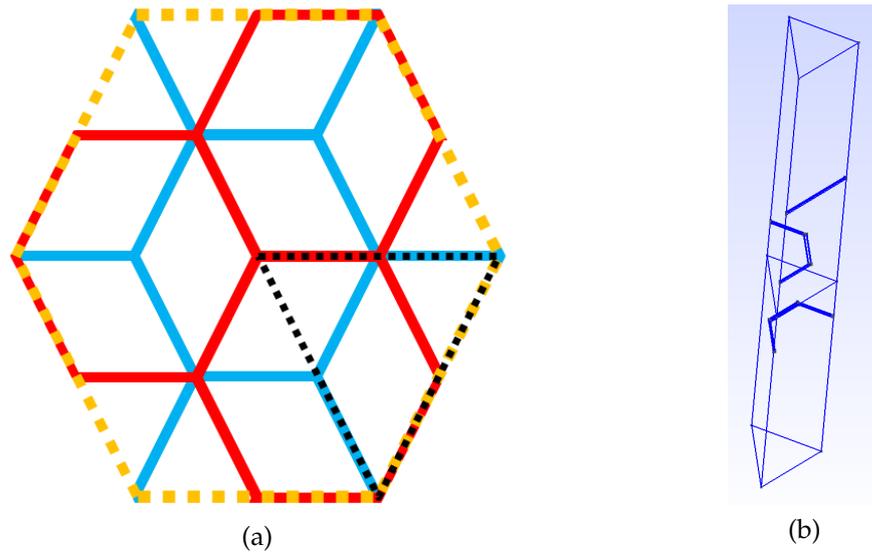


Figure 7.14: Layout of the hexagonal grids in XENON100. The anode and gate are offset from each other by half the pitch. The top grid has twice the pitch of the anode and gate. (a) grids viewed from above, yellow dotted line is the top grid, the red line is the anode and blue for the gate. The black dotted line shows the triangular unit cell. (b) Triangular unit cell. This is reflected at each edge to create the whole geometry.

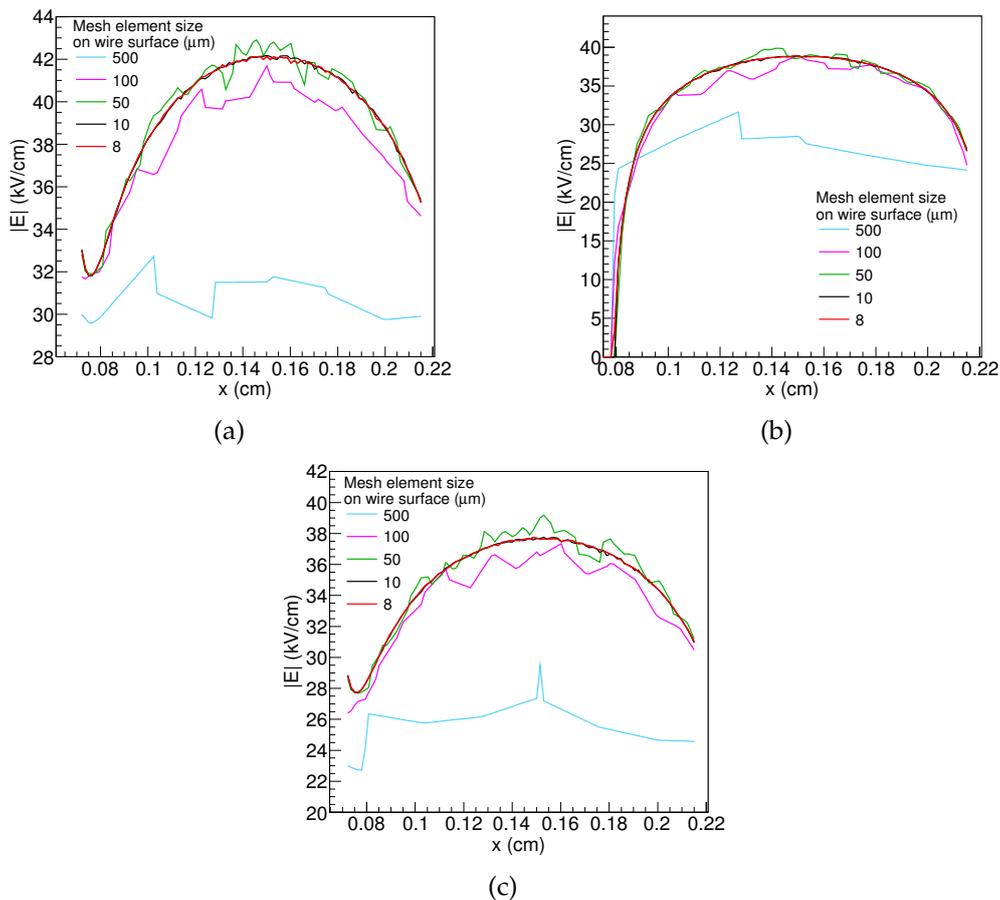


Figure 7.15: Electric field strength on the anode wires for XENON100 with different mesh element sizes on the wire surfaces. (a) Along the bottom of a wire segment (one side of the hexagon). (b) Along the side of a wire segment. (c) Along the top of a wire segment.

The XENON100 geometry is a larger volume than the LUX woven mesh and three grids are included. This means that there are more severe limitations on the mesh size. The lowest mesh density here did not reproduce the field on the wire well. The next mesh element size, $10\ \mu\text{m}$, was only slightly off. The other three reproduced the mean value well, with the smallest two being smooth. However, the smallest mesh element size greatly increases the simulation time in Garfield++, from a few hours to over 12 hours per electron drift. It is important to note that the field away from the wires is similar in all of the cases, with only a slight disagreement in the lowest mesh density as can be seen in Figs. 7.16c and 7.16d. Therefore, to check that the mesh density does not introduce any errors, the simulation with sharp wires will be done for each mesh. It is important to check the electron multiplication too, which may occur in the high field near the wires and this could be affected by the mesh density. The effect should be small at least for the three highest mesh densities, as the fluctuations in the field are only of the order $\sim 2\%$ and it only affects the very last part of the drift.

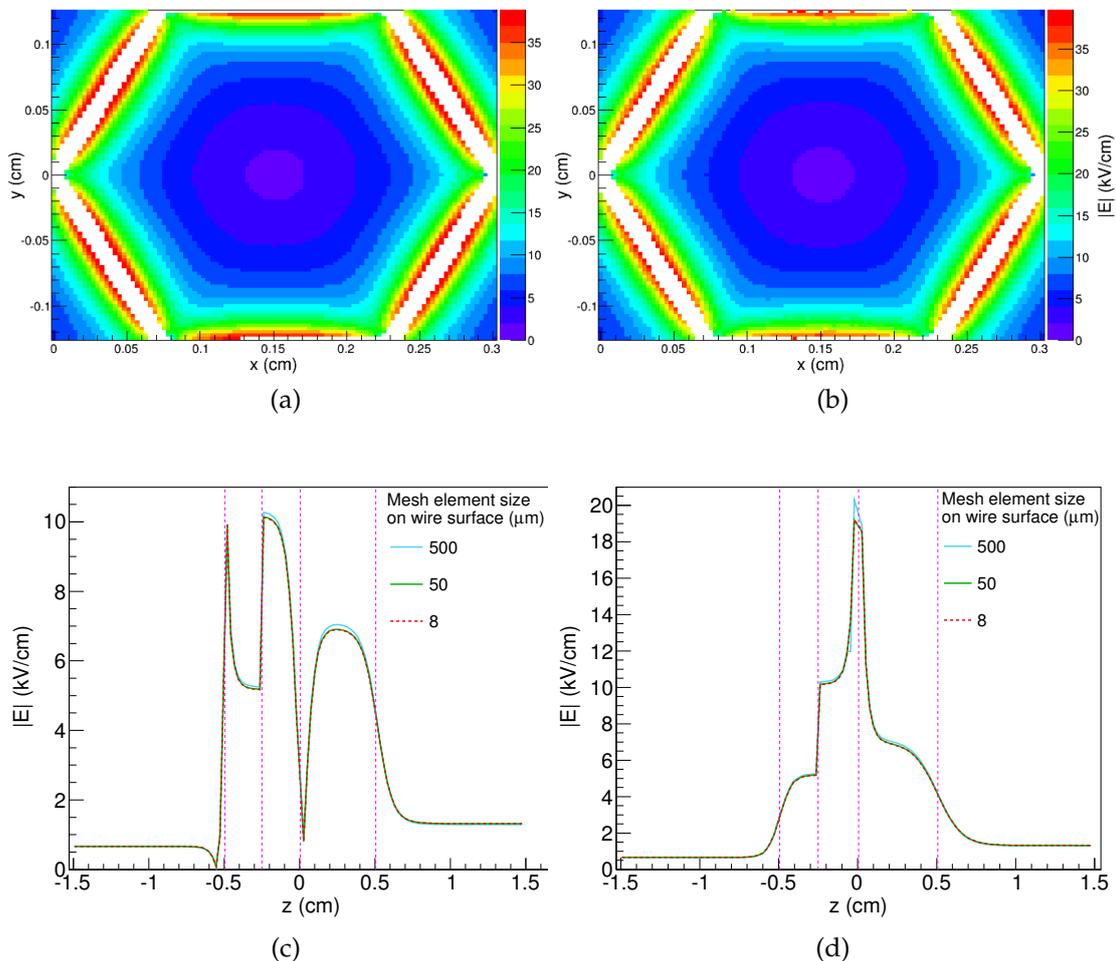


Figure 7.16: Electric fields in the XENON100 simulation. (a) E field at the anode in the x - y plane, for a mesh element size of $8\ \mu\text{m}$ on the wire surface. (b) Same as (a), but with a mesh element size of $50\ \mu\text{m}$ on the wire surface. (c) E field along z , through the centre of the anode. (corner of the gate). The vertical magenta lines show the positions of the gate, liquid surface, anode and top PMT shield from left to right. (d) E field along z , passing through a corner of an anode hexagon (centre of gate), vertical lines are the same as (c).

The cross-sectional shape of the wires is not known, and the sharp edges are not likely to be the case in reality as the XENON100 grids are made by chemical etching from a sheet. Therefore, the wires will have some rounding. Along with the sharp wires, two rounded rectangles were chosen with the rounded corners making up either 20% or 80% of the circumference. The different wires are shown in Fig. 7.17. After an appropriate mesh density was found, electron drifts were simulated for each wire rounding.

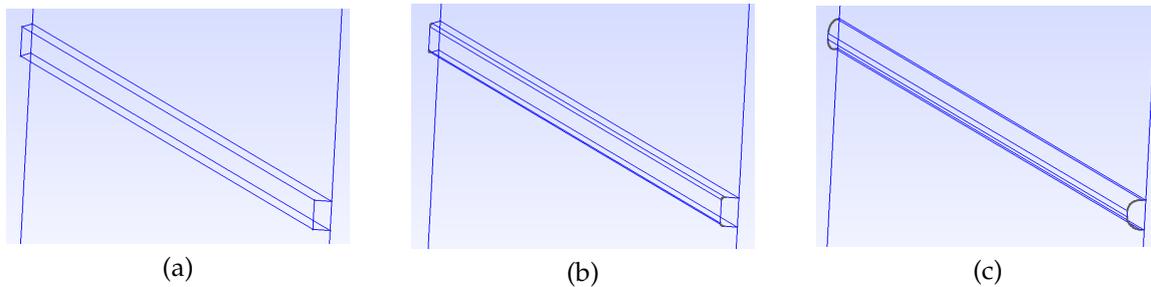


Figure 7.17: Different roundings of XENON100 wires. A certain percentage of the wire surface is part of the curved corners. (a) 0% rounding. (b) 20% rounding. (c) 80% rounding.

The boundary conditions here are set in the same way as the LUX field map. The potentials at the top and bottom of the unit cell are chosen so that the fields in Fig. 7.10c are reproduced with the appropriate voltages set on the wires, using the same values as Ref. [327]. Again, the boundary conditions at the side of the cell are that the electric flux through them is zero.

Another difference from the LUX scenario is that the field in the gas drift region is not as uniform. This is due to the combination of high electric field and a shorter gas gap. However, the design of the gate-anode focusses electrons, which reduces the variance in the drift paths. This is accounted for by finding electron drift lines through the liquid to the surface, as Garfield++ can only do drifts in the gas phase. An example set of drift lines from the XENON100 simulation is shown in Fig. 7.18. Then, electrons start positions are sampled at random from within where the drift lines cross the liquid surface. In the parallel plate version, the constant field is set to the average field in the gas gap, so that the mean yield is the same as the actual geometry.

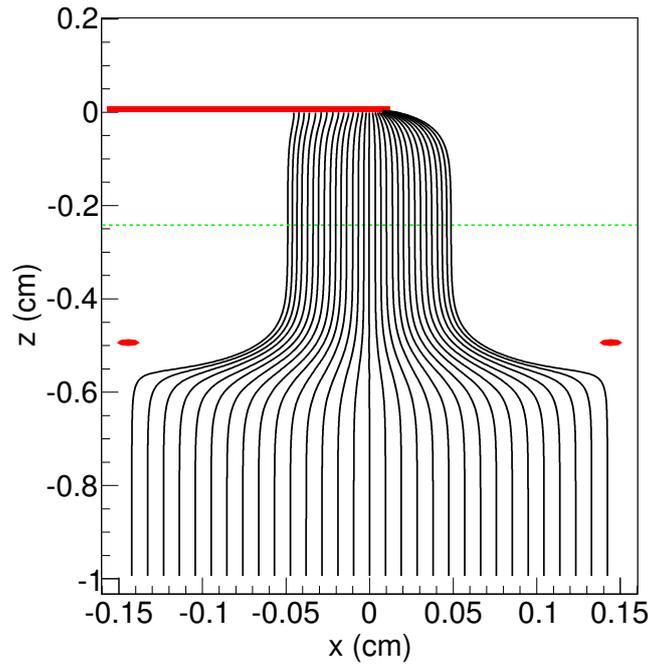


Figure 7.18: Drift lines in the XENON100 simulation. The wires locations in this plane ($y=0$) are shown in red, and the liquid surface in green.

7.5.2 Mesh density effects on photon and ionisation yields

The first results are from the LUX and XENON100 geometries (with sharp wires) with different mesh densities on the wire surface. Although the fields on the wire surfaces were found to be very similar for all but the lowest mesh density, it is important to check that the mesh density does not affect the mean photon yield or variance. The results from the LUX drifts are shown in Fig. 7.19 and summarised in Table 7.2, followed by the same for XENON100 in Fig. 7.20a and Table 7.3. In addition, Fig. 7.20b shows the number of secondary electrons produced in each XENON100 simulation.

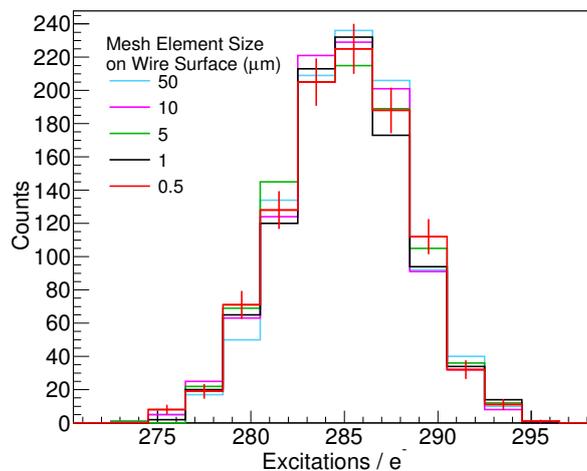


Figure 7.19: Photons emitted per electron in the LUX simulation with different mesh densities on the wire surface. For clarity, error bars are only shown for the highest mesh density.

Table 7.2: Mean and relative widths for the LUX simulation, with the field map created using different mesh element sizes.

Mesh element size at wire surface (μm)	γ/e^-	σ/μ (%)
50	285.2 ± 0.1	1.13 ± 0.02
10	284.9 ± 0.1	1.15 ± 0.02
5	285.1 ± 0.1	1.18 ± 0.02
1	285.0 ± 0.1	1.16 ± 0.02
0.5	285.0 ± 0.1	1.18 ± 0.02

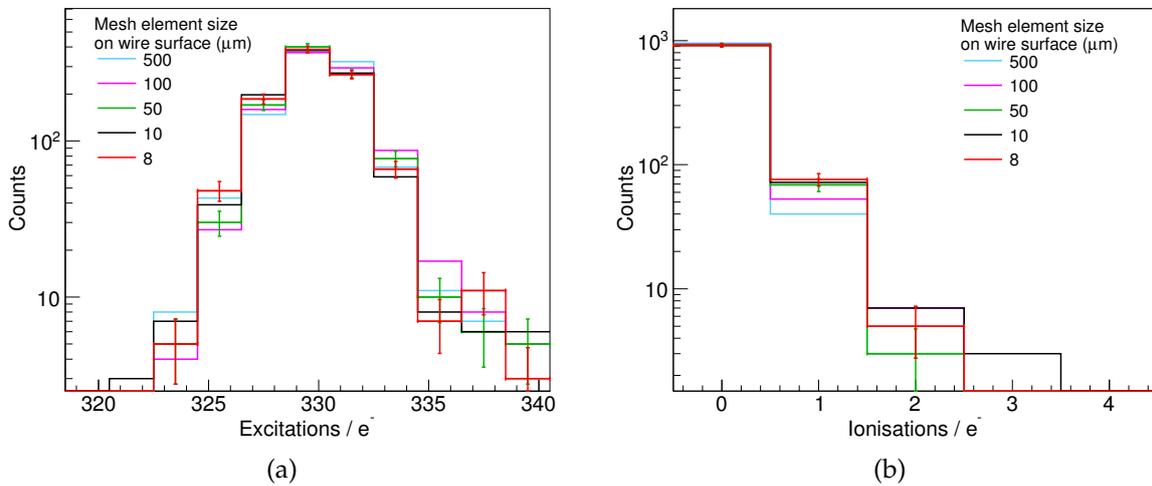


Figure 7.20: XENON100 simulation results for varying mesh densities on the anode surface. (a) Number of excitations (photons) per electron. Only the main peak is shown here for clarity, there are also events with more excitations when ionisation also occurs. (b) Number of ionisations (secondary electrons produced) per electron drift.

Table 7.3: Mean and relative widths for the XENON simulation, with the field map created using different mesh element sizes. The error shown only accounts for the statistical error on the mean, and not any systematic error.

Mesh element size at wire surface (μm)	γ/e^-	σ/μ (%)	ions/ e^-
500	330.4 ± 0.1	1.00 ± 0.02	0.05 ± 0.01
100	330.9 ± 0.2	1.45 ± 0.03	0.07 ± 0.019
50	330.5 ± 0.1	1.15 ± 0.02	0.08 ± 0.01
10	330.3 ± 0.1	1.31 ± 0.03	0.10 ± 0.01
8	330.3 ± 0.1	1.13 ± 0.02	0.09 ± 0.01

From these it is clear when the mesh density starts to have an effect. In the LUX simulation, even the largest mesh element size of $50 \mu\text{m}$ did not have significantly different results. The two larger sizes tried for XENON100 did have a small effect on the number of ions, which can be seen in Fig. 7.20b and the relative widths in the photon histograms. Therefore, $50 \mu\text{m}$ is taken to be the maximum mesh element size on the wire surface for all of the following simulations. It is possible that a lower mesh density could be used for the woven wire meshes, but the computation time for this chosen density is acceptable.

7.5.3 Results from simulations of experimental geometries

Firstly, electron drifts were simulated for each experimental geometry and for the corresponding planar-anode versions to find the mean and variance in the photons produced per electron. Figure 7.21 shows the ZEPLIN-III simulation and Fig. 7.22 for LUX and its planar version. Figure 7.23 shows the photon variance for each of the XENON100 simulations with different wire roundings and the planar version. Here, the number of secondary electrons produced is compared in Fig. 7.24. This occurs near to the anode wire so could be affected by the wire rounding. The results from the different geometries are summarised in Table 7.8.

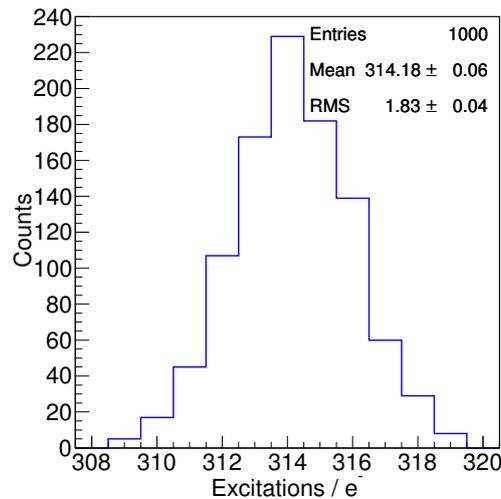


Figure 7.21: Excitations per electron in the ZEPLIN-III planar geometry simulation, $P = 1.6 \text{ bar}$, $E_g = 7.6 \text{ kV/cm}$, $h_g = 4.0 \text{ mm}$.

The planar ZEPLIN-III scenario does have very good resolution in S2 light production, as expected. The relative width is 0.58%, which should not increase the variance even for small signals. However, ZEPLIN-III only had one PMT array which gave a lower photon detection efficiency that would have degraded the S2 resolution.

The LUX simulation with a wire mesh anode has poorer resolution (1.16%) than ZEPLIN-III. The overall photon yield for both is similar, and the gas pressure is the same. Therefore, the difference has to be due to the electric field, either the strength or shape. The gas field in ZEPLIN-III was 7.6 kV compared to 5.6 kV in LUX. The comparison between the LUX simulation with the anode mesh (Fig. 7.22a) to the planar version (Fig. 7.22b) did not show any large difference in the relative width. The planar

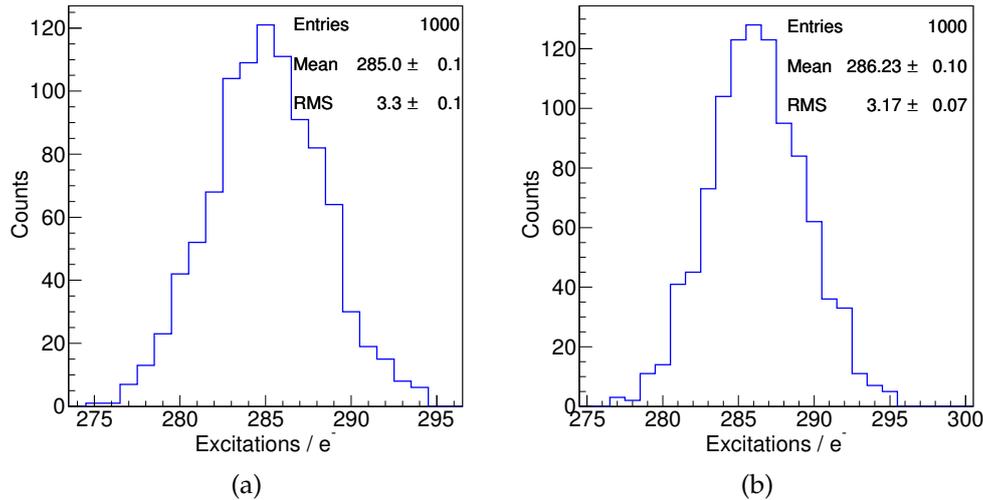


Figure 7.22: Excitations per electron in LUX simulations with parameters $P = 1.6$ bar, $E_g = 5.6$ kV/cm, $h_g = 5.6$ mm. (a) Woven wire mesh anode. (b) LUX-like planar geometry.

version is slightly smaller at 1.11%, but it is clear that the field shape near the anode does not degrade the resolution in LUX. Therefore, the smaller variance in S2 light production in ZEPLIN-III must be due to the higher electric field.

The first three XENON100 simulations of hexagonal meshes with different wire roundings all gave similar results. The means are within 1% of each other and the RMS within 35%. It appears that the 80% rounding case has the largest width, due to the tail to the right of the main peak. This tail is caused by secondary electron production near to the anode wires. Therefore, it cannot be seen in the planar version in Fig. 7.23d. Even with the multiplication, it is not increasing the variance to a prohibitively large amount: for the sharp wires simulation the variance is only 1.15%. However, the planar version is only 0.44%. The geometry does have a significant effect here, but the higher field relative to LUX is compensating for the extra variance from the multiplication.

The last aspect to check for the XENON100 scenario is if there are any differences in the number of secondary electrons produced per drift. Fig. 7.24 shows the number of ionisations with each of the three different roundings. For clarity, error bars are only shown for the 0% rounding. It can be seen that the agreement is reasonable, even between the 0% and 80% case. The 80% rounding does have more ionisations overall in this test, but not by a statistically significant amount.

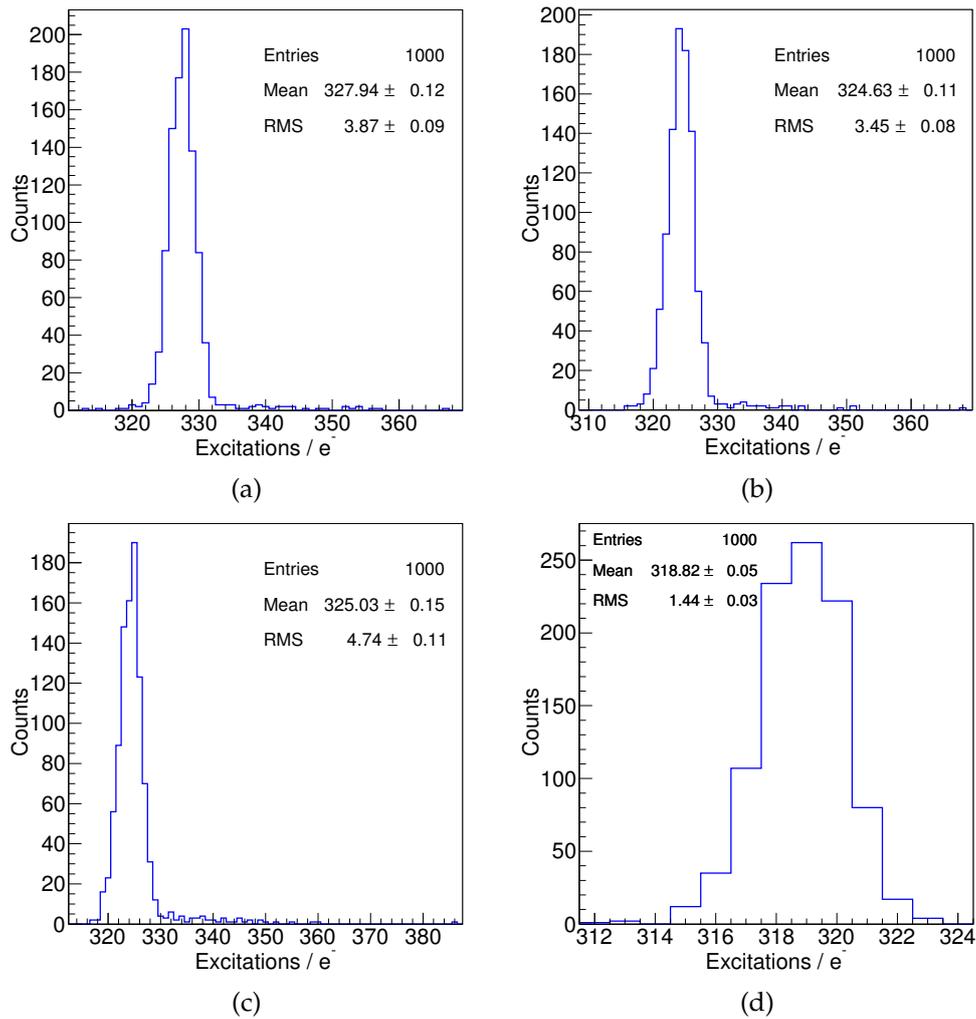


Figure 7.23: Excitations per electron in the XENON100 simulations, $P = 2.3$ bar, $E_g = 11.9$ kV/cm, $h_g = 2.5$ mm. (a) Hexagonal mesh, no rounding. (b) Hexagonal mesh, 20% rounding. (c) Hexagonal mesh, 80% rounding. (d) XENON100-like planar geometry.

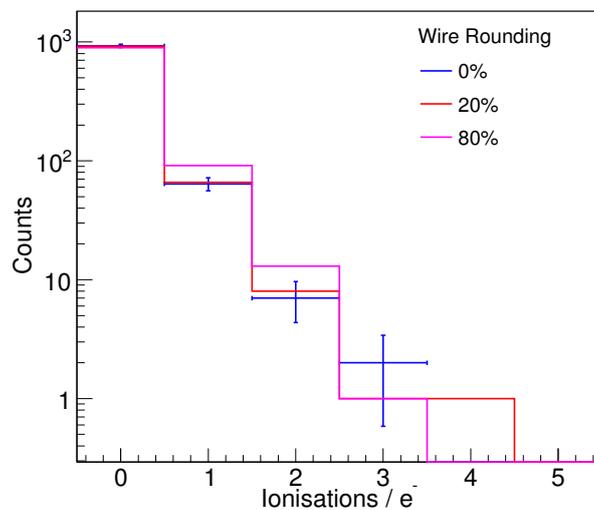


Figure 7.24: Ionisations per electron in the XENON100 simulations with different wire roundings. The planar version at 11.9 kV/cm does not show any secondary electron production.

On the other hand, the high field in XENON100 does give the largest electron emission probability of the three experiments (although this does not affect our results, which are only for emitted electrons), and there is less obscuration of light by the grids than in LUX. Both of these would increase the size of S2 signals, therefore giving better photon statistics.

The maximum electric fields on the wires were also found for each case. If the grid is a woven wire mesh or hexagonal mesh and is included in the simulation, then the maximum field is found by sampling points close to the wires. For the woven wire mesh, the maximum field occurs where the wires cross over, at the bottom of the wire. In the hexagonal XENON100 mesh the maximum was found to be on the bottom corners of the wires, half way between two corners of the hexagon. This was found to decrease quickly away from this point. Other values of the field along the wires can be seen in Fig. 7.15. For parallel wire grids that are not included in the electric field map, the maximum field is calculated using equation 7.15.

Table 7.4: Results from simulations of the three experimental geometries. h_g is the gas gap height, P_g the pressure in the gas, E_g the average field in the gas, γ/e^- the number of photons produced per electron and σ/μ the width over mean. The optical transmission is the probability for light to pass through the grid at normal incidence. For the max field and optical transmission, A, G and C refer to the anode, gate and cathode respectively. The maximum electric fields are found from the field map by scanning along the wire surfaces. The XENON100 result shown here is for the wires without rounding. Therefore, the maximum field on the wire for this case should be taken as an upper limit. The electron emission probability from the liquid is calculated using a fit to [312]. The LUX-like and XENON-like scenarios are the same layout as ZEPLIN-III but with the gas height, electric field and pressure changed to be the same as LUX and XENON100 respectively.

	h_g (cm)	P_g (bar)	E_g (kV/cm)	γ/e^-	σ/μ (%)	Max Field on Wire (kV/cm)			Optical Transmission (%)			e ⁻ Emission Probability
						A	G	C	A	G	C	
ZEPLIN-III	0.40	1.6	7.6	314.2 ± 0.06	0.58 ± 0.01	-	-	75	-	-	-	82
LUX Run03	0.56	1.6	5.6	285.0 ± 0.1	1.16 ± 0.02	17*	46	19	77.4	98.0	-	49
XENON100	0.25	2.3	11.9	328.0 ± 0.1	1.15 ± 0.02	120*	22*	-	90.3	90.3	-	≈ 100
ZEPLIN-III LUX-like	0.56	1.6	5.6	286.2 ± 0.1	1.11 ± 0.02	-	-	-	-	-	-	-
ZEPLIN-III XENON-like	0.25	2.3	11.9	318.8 ± 0.05	0.45 ± 0.01	-	-	-	-	-	-	-

7.5.4 Comparison to LUX data

Two comparisons are made to LUX data. One is to simply validate the mean yield of Garfield++ under the normal Run 3 conditions. The single electron size in LUX is 24.88 phd with a photon detection efficiency of 9%. This implies 276 emitted photons per electron. The Garfield++ simulation is within 3% of this value.

In addition to this, there are LUX data available where the gate and anode voltages were increased until the onset of electron emission from the gate. This was studied in detail in Chapter 6, but here it serves as extra validation. Emission data were taken with the liquid level above and below the gate. In particular, one emission spot with the gate in the gaseous phase appeared to be consistent with single electron emission followed by electron multiplication in the high field region near to the wire. Single electron emitters in the liquid were also observed. The “gate-in-gas” scenario can be simulated in Garfield++. As well as the photon yield, this allows a validation of the secondary electron production and pulse shape.

The details of this simulation are summarised here, for more details refer to Chapter 6. It was found that due to the very small pitch of the LUX anode, that the shape of the field near to the anode wires had a negligible effect on the photon yield and pulse shape. Therefore, the scenario is completely simulated in Garfield++. The gate is a parallel wire grid between two planes at constant potential representing the liquid surface and the anode. The potential on the liquid surface is calculated in the same way as the LUX Run3 scenario in section 7.5.1. The anode and gate grids were at ± 3.8 kV respectively, with the bottom PMT shield at -2 kV/cm and the cathode and top PMT shield were grounded. This creates an electric field of 6.2 kV/cm between the gate and anode. The operating pressure was 1.55 bar.

The azimuthal position of the emitter on the wire surface is unknown. This does affect the drifting electrons, as the electric field strength varies around the wire. Also the pulse length is affected: if the emitter is on the top of the gate wire the path to the anode will be shorter than if it were on the side or bottom. The nominal gate-anode distance is taken to be 0.96 ± 0.03 cm. The photon detection efficiency for S2 photons in the gas is also unknown, a dedicated Monte Carlo simulation indicates it to be approximately 10%, similar to that in the liquid.

In the data individual peaks in the histogram of the number of photoelectrons emitted per electron could be seen (up to approximately 500 phe). These correspond to the number of secondary electrons produced in the high electric field near to the gate wires. Therefore, the photon detection efficiency is fixed by making sure that the first peak (for when no electron multiplication occurred) matches for both the simulation and data. The previous validations show that the photon yield (without multiplication) is reproduced within 10% by Garfield++.

Various starting angles were simulated and it was found this had a significant effect on the results. The top of the wire had the highest field and more multiplication. It was found that an angle of 45° gave the best agreement. In order to match the first peak without any secondary electrons produced, a photon detection efficiency of 9.2% was

required. This is in good agreement with the expected value. Figure 7.25 shows the number of phe per electron drift.

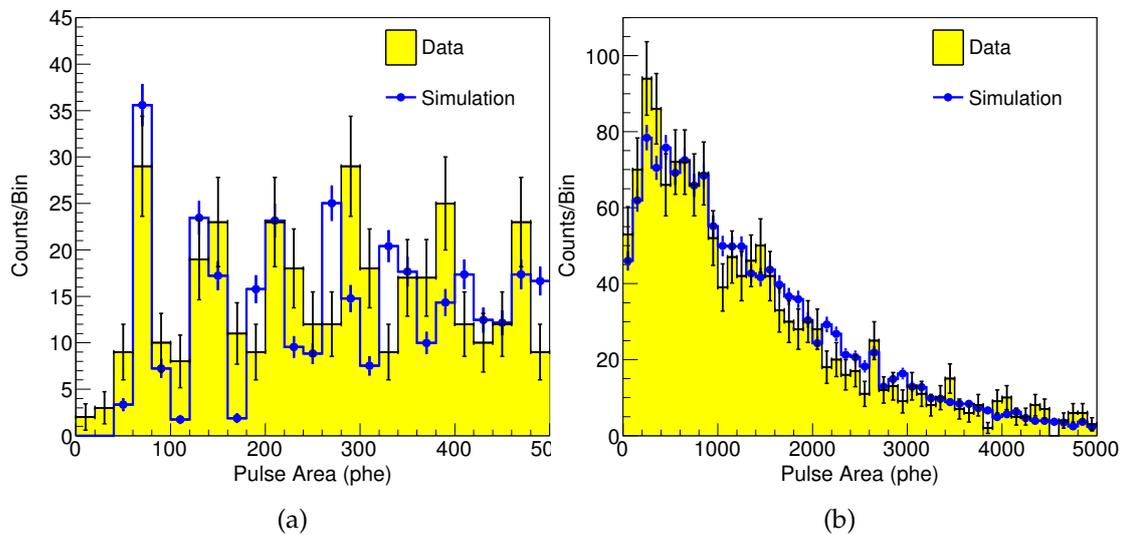


Figure 7.25: Pulse areas for the single electron emitter from the LUX gate in the gas (emitter E on Jan. 10th 2014); electrons in the simulation start at 45° from the horizontal on the wire. (a) 0 to 500 phe range. (b) Whole range up to 5000 phe.

There is good agreement in both the peaks seen at low energy and the overall shape of the spectrum. This means that the number of secondary electrons is being reproduced well, in addition to the number of excitations.

The other investigation with these data was to examine the pulse shape. In the simulation, the position and timing of excitations are known, so the waveform can be constructed. There is a time delay between the excitation and photon detection in the PMTs which is not included in the simulation, but this is added *a posteriori*. Additional time is added to the time of each excitation. This includes a component for the relaxation time of the Xe_2^* excimer (sampled from an exponential with a time constant of 60 ns [328]) and the transit time jitter in the PMTs (Gaussian with $\sigma = 1.7$ ns), all being assumed equal. The pulse shape from the $1e^-$ peak in the data (at 70.5 phe) is compared to the simulation with electron multiplication turned off. Then, the average pulse shape from all of the data is compared to that from the simulation.

These simulations are compared to average waveforms through the LUX DAQ in Fig. 7.26. For the 45° starting angle the pulse length is too short. However, the shape is reproduced well. The peak at the start is due to the high electric field near the gate wire. Then it plateaus once it reaches the constant field. The tail at the end in the simulation is due to the variance in time taken for the electrons to reach the anode. The simulation starting from -20° reproduced the pulse length better and matched very well for the $1e^-$ case. However, when looking at all of the pulses, the pulse area does not agree, indicating that this starting angle does not produce enough secondary electrons. It is possible that there is extra delay between the excitation and phe detection that is not accounted for here. Also, the gate-anode distance could be larger than the 0.96 cm used in this analysis.

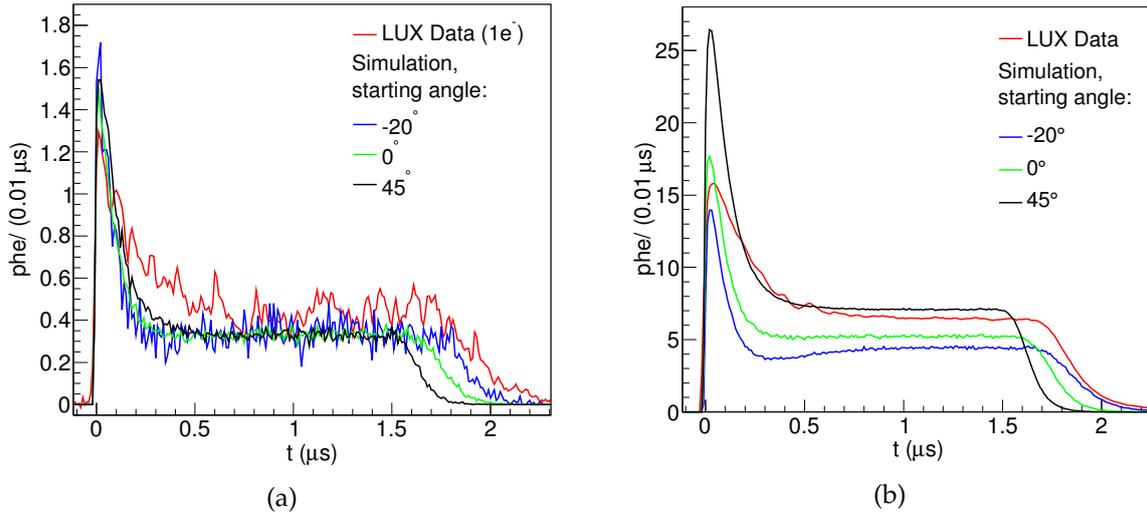


Figure 7.26: Average waveforms for data and simulation for the single electron emitter in the gas (emitter E on Jan. 10th 2014). The angle is measured from the wire plane, with the top of the wire at 90° . (a) For events with no secondary electrons. The waveform from the data is an average of all pulses in the $1e^-$ peak, simulation is only tracking the primary electron. (b) All data, and including electron multiplication in the simulation.

Taking into account the uncertainties involved, the absolute agreement in the photon yield is actually remarkably good. It is important that a process involving both excitations and ionisations can be reproduced, as this validates that the simulation is accurate when electron multiplication occurs. It also means that the variance in the EL response has been reproduced well.

7.5.5 Comparison to XENON100 data

Experimental data are available for the mean single electron response at varying fields for XENON100 [329]. This can be used to further validate the simulations for more complex geometries, and to assess whether charge multiplication may be occurring near the electrodes in this and other experiments. The data were taken at different liquid heights, then scaled to a gas height of 2.9 mm. The electric field in the bulk of the gas, E_g , was calculated by the authors using the simple two-phase, parallel plate approximation:

$$E_g = \frac{\epsilon_r V_a}{\epsilon_r h_g + d - h_g}, \quad (7.20)$$

where ϵ_r ($=1.96$) is the relative permittivity of liquid xenon, V_a is the voltage on the anode, h_g is the gas gap and d the gate-anode distance. The XENON100 single electron data is reproduced in Fig. 7.27. The uncertainties are dominated by the light collection and PMT quantum efficiency, given as $\sim 20\%$ and $\sim 25\%$, respectively. At approximately 13 kV/cm in the XENON100 data, an up-turn is observed which is attributed by the authors to electron multiplication near the anode [329].

For the simulations the gas gap was initially set to 2.9 mm so that no scaling was required. To investigate the effect of the scaling, simulations were also carried out for

a 1.25 mm gas height, with the number of emitted photons then multiplied by a factor of 2.9/1.25 to scale to the correct gas height.

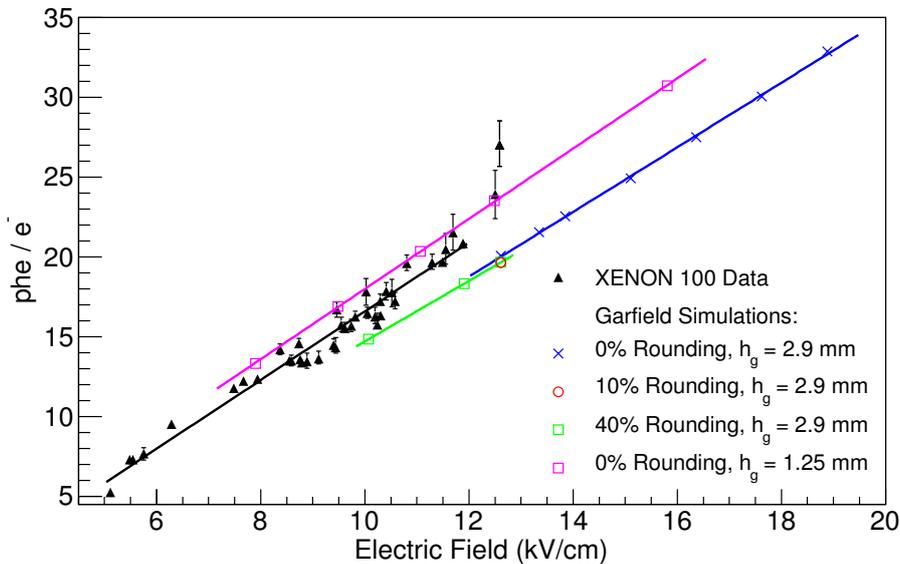


Figure 7.27: Single electron data from XENON100 [329] and simulations of that geometry from the present work. The data were taken at various gas gaps and scaled to 2.9 mm. The black line is a fit to the data. The simulations with 1.25 mm gas gap are also scaled to 2.9 mm. The point with 10% rounding at 12.5 kV/cm field gives the same result as 40% rounding. The two XENON100 data points that deviate from the linear trend are attributed to charge multiplication near the wires in [329]. For points without visible error bars, the errors are small.

This comparison to the XENON100 data shows that the overall trend is reproduced well, but there is some systematic shift in the number of phe per electron. The simulation points with a 2.9 mm gas gap points are only $\sim 10\%$ below the experimental data, so the hexagonal mesh simulation reproduces the data well. The amount of rounding on the wire was found to have very little effect. A reason for the systematic error could be the uncertainty in the photon detection efficiency and quantum efficiency. Another unknown is the gas gap for each measurement in the data, which is not reported in [329]. The most notable discrepancy is, however, the points in the experimental data where charge multiplication is claimed: the sharp upturn is not observed in the simulation. This could be related to the scaling and gas gap. If the data points at high field were taken with a gas gap smaller than 2.9 mm, the scaling would multiply the effects of the high field near to the wires and artificially increase the photon yield. The 1.25 mm gas gap simulation scaled to 2.9 mm shows this behavior.

7.5.6 Conclusions from validation studies

Validations against experimental data showed that the photon yield and drift velocity are reproduced well by Garfield++. Electroluminescence simulations of the existing experiments confirmed that the planar ZEPLIN-III anode geometry had the lowest variance in photon production, which is not surprising given the uniform field. When comparing the LUX mesh to a planar geometry with the same parameters (pressure,

temperature, electric field, gas height), the relative standard deviation was similar, so the LUX anode is indeed fine enough. On the other hand, the comparison between XENON100 and its planar version showed that the shape of the field increased the variance in photon production. However, it is important to note that the XENON100 grids have a high optical transmission and the high field improves the electron extraction efficiency.

In addition, the simulation reproduced the response of LUX to single electrons emitted from the gate very successfully, including the photon yield in a scenario where electron multiplication occurs, and the pulse shape. The agreement with XENON100 single electron size was reasonable over a range of electric fields, but disagrees with the interpretation that electron multiplication occurred at high fields.

These simulations also help to guide the LZ grid design. It is known that a 1% variance in the S2 photon production can be achieved [330]. Ideally, this would also be achieved in LZ, although it is challenging. Both a uniform field shape and stronger field were found to improve the variance, as expected.

7.6 Simulations of LZ Grids

Finally, we are well-equipped to simulate possible designs for LZ. There are three sets of gate-anode geometry considered, which are labelled as the “initial”, “CDR” and “pre-TDR” designs.¹ The initial design features 1 cm between the anode and gate, with the liquid level half way between them; this was a pre-CDR scenario, with -3.5 kV on the gate and $+4$ kV on the anode. The second had the same dimensions, but with -4 kV on the gate, giving an electric field of ~ 9 kV/cm in the EL region and was adopted for the CDR. An alternative design had a larger gate-anode distance of 1.5 cm, with the liquid level 0.5 cm above the gate. The anode-gate potential is now ± 7 kV, to deliver a similar electric field as previously — this was a pre-TDR design. The advantage of the larger gas gap is that it increases the S2 photon yield, reduces the relative variance and means the effects of grid deflection are relatively smaller. The parameters of each scenario are summarised in Fig. 7.28.

For the pre-CDR design (1 cm gate-anode distance) a range of wire pitches (1 to 3 mm) and diameters (50 to 200 μm) were investigated for both parallel and crossed-wire grids. Although the parallel wire grids are unlikely to have acceptable mechanical properties, it is useful to see the effect it has on the EL signal. In terms of the EL production, crossed grids have similar performance to woven meshes although their mechanical properties are different. For the initial design only the anode grid is simulated.

For the CDR and pre-TDR scenarios only a few different anode parameters are simulated, based on the previous simulations of the initial design and which grids can actually be manufactured. Table 7.5 shows the different combinations of anode and

¹The reader should note that often we use this terminology to describe the general grid configuration rather than the wire pitch and diameter, which can depart from the actual values used to calculate the detector performance. It should be clear when this is so.

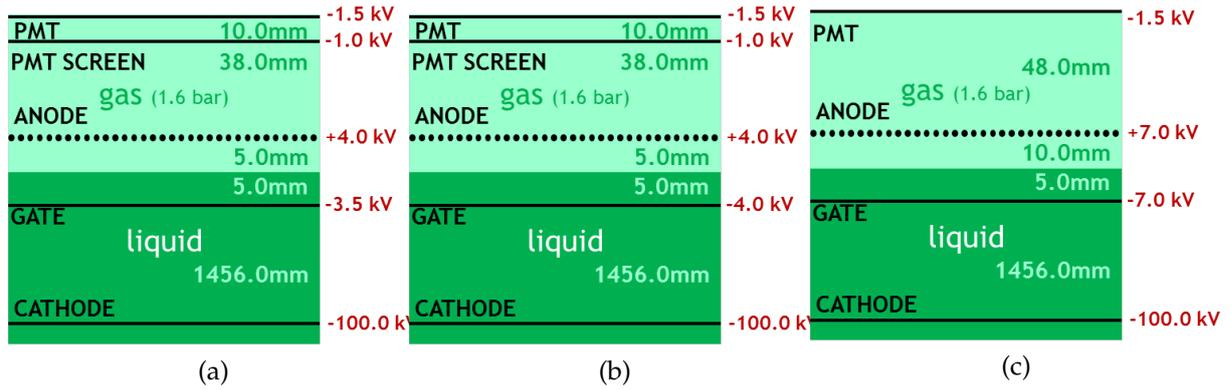


Figure 7.28: Parameters for the different possible LZ scenarios including distances between grids and voltages. Note that we adopt the cathode HV goal of 100 kV rather than the baseline of 50 kV. This difference has little effect on the field in the EL region. (a) Pre-CDR version used for preliminary simulations. (b) CDR model. (c) Pre-TDR model.

gate pitches simulated for the CDR and pre-TDR scenarios. Unlike the LUX simulations, the anode and gate have comparable pitches so electron focussing by the gate is important. Also, the larger pitch means that field leakage through grids is more significant. Therefore, the CDR and pre-TDR design simulations include both the gate and anode. The other consideration is the alignment of the grids. Depending on whether the gate and anode are aligned or not, the starting positions of the electrons will be different. This affects the paths that the electrons take to the anode wires and can change the observed mean photon yield and variance.

Table 7.5: Parameters of the woven wire meshes tested for the CDR and pre-TDR designs. Each of these scenarios can have the gate and anode aligned or offset by half the gate pitch.

Anode		Gate	
Pitch (mm)	Wire Diameter (μm)	Pitch (mm)	Wire Diameter (μm)
5	150	5	75
2.5	150	5	75
7.5	150	7.5	75

7.6.1 Creating LZ field maps

The simulations of the initial pre-CDR scenario considered only the region between the liquid surface and the anode. The potential at the liquid surface was calculated using the multiple wire grid calculation from [154]. Then, the top of the simulation domain was defined by a plane at constant potential located 0.3 cm above the anode. For the parallel wire grids, the field is calculated in Garfield++. The crossed wire grid geometry is created in GMSH, using a unit cell with two overlapping wires as shown in Fig. 7.29. The mesh density was chosen based on what was found to be acceptable for the previous LUX and XENON100 scenarios. The starting positions of the electron

drifts were sampled at random from the whole of the unit cell. Therefore, focussing from the gate is not accounted for in this instance.

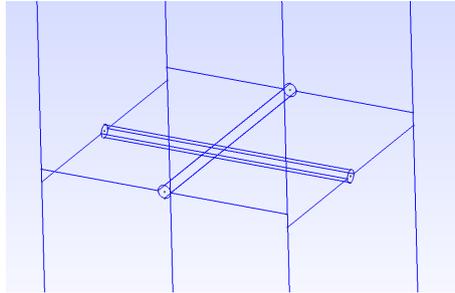


Figure 7.29: Unit cell for the crossed grid simulations.

For the CDR and pre-TDR scenarios the simulations include both the gate and the anode as woven meshes to include focussing from the gate. The woven meshes are created in GMSH, using the same method as the LUX anode in section 7.5.1. To more accurately model the field leakage, the first step is to make a field map for the cathode to the top of the simulation. For the CDR simulation the top boundary is the top PMT screen, which is taken to be a plane. For the pre-TDR, there is no top PMT shield, so the top of the simulation is a plane for the top PMT array. The cathode is also a plane, with the gate and anode as woven meshes. The initial unit cell is a very large volume, so the mesh density has to be low. This means that the field near to the wires will not be very accurate, but it is acceptable for the field away from the wires.

The next step is to shorten the simulation boundaries by finding planes of constant potential below the gate and above the anode. This simulation is done with a higher mesh density. Figure 7.30 shows the electric field for the pre-TDR scenario (pitch 5 mm, gate wire diameter $75\ \mu\text{m}$, anode wire diameter $150\ \mu\text{m}$) from the cathode-top and gate-anode simulations. The agreement here is good, which again shows the field away from the wires is not affected by the mesh density. The gate-anode field map is used to find electron drift lines. Electrons are assumed to come from a uniform plane below the gate and diffusion in the liquid as they follow the drift lines is ignored. The point at which the drift lines cross the liquid surface determines where the electron drifts can start from in the gas.

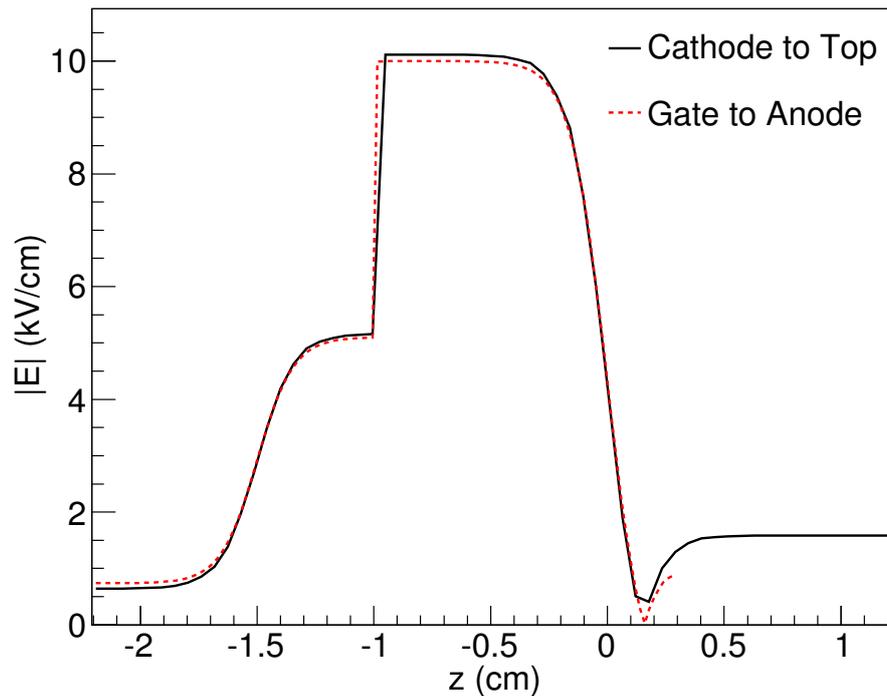


Figure 7.30: Electric field in the LZ pre-TDR scenario with 5 mm gate and anode pitch, aligned grids. Field is sampled going upwards (z), in the centre of the wires. The gate is at -1.5 cm, liquid surface at -1 cm and anode at 0 cm.

For the drift lines it is also important to check the effect of having aligned or offset grids. For aligned grids of the same pitch, the gate and anode wires are simply on top of each other. Alternatively, if they are offset the crossing point of the gate wires is in the middle of a square of the anode mesh as shown in Fig. 7.31. For the offset case the electrons are focussed onto the anode wires by the gate. The offset case requires a larger unit cell in order to have mirror symmetry at the edges for both the anode and gate. Figure 7.32 shows the field lines for the pre-TDR scenario with aligned and offset grids. The area covered by the drift lines on the liquid surface is the same in both cases and is only displaced relative to the anode wires.

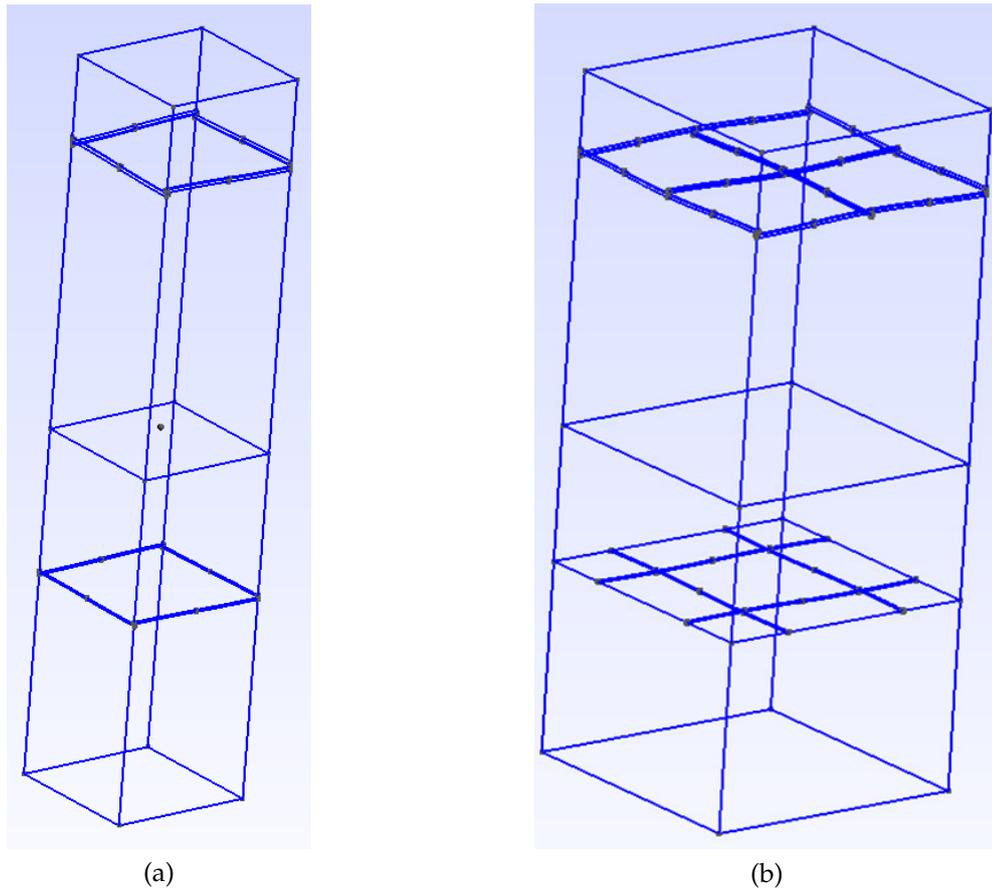


Figure 7.31: Woven mesh gate-anode unit cell for LZ with equal anode and gate pitches. (a) Aligned grids. (b) Grids offset by half of the pitch.

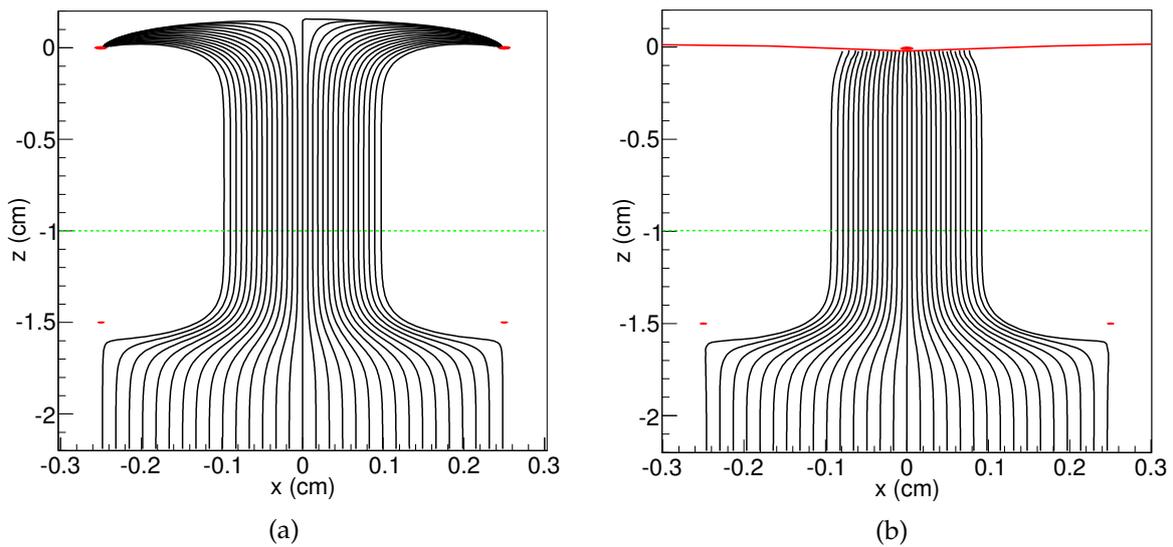


Figure 7.32: Electron drift lines for the LZ pre-TDR scenario (see Fig. 7.28c) with gate and anode aligned with a pitch of 5 mm. The green line indicates the liquid surface. The red circles and lines are the locations of the wires in the $y = 0$ plane. (a) Aligned grids. (b) Grids offset by half of the pitch.

Figure 7.33a shows the geometry for the aligned case where the anode is half of the gate pitch. It is possible to make the offset geometry with the half-anode pitch as shown in Fig. 7.33b. However a 4×4 anode grid is required to be able to have mirror symmetry at the boundary of the unit cell. This is a large volume of wires, and it is not possible to solve the field map for this geometry. Therefore, it is assumed that the drift lines for the offset case are simply displaced by half of a gate pitch, as was the case for an equal anode and gate pitch. The drift lines for the pre-TDR scenario with aligned grids, anode pitch of 2.5 mm and 5 mm gate pitch is shown in Fig. 7.34.

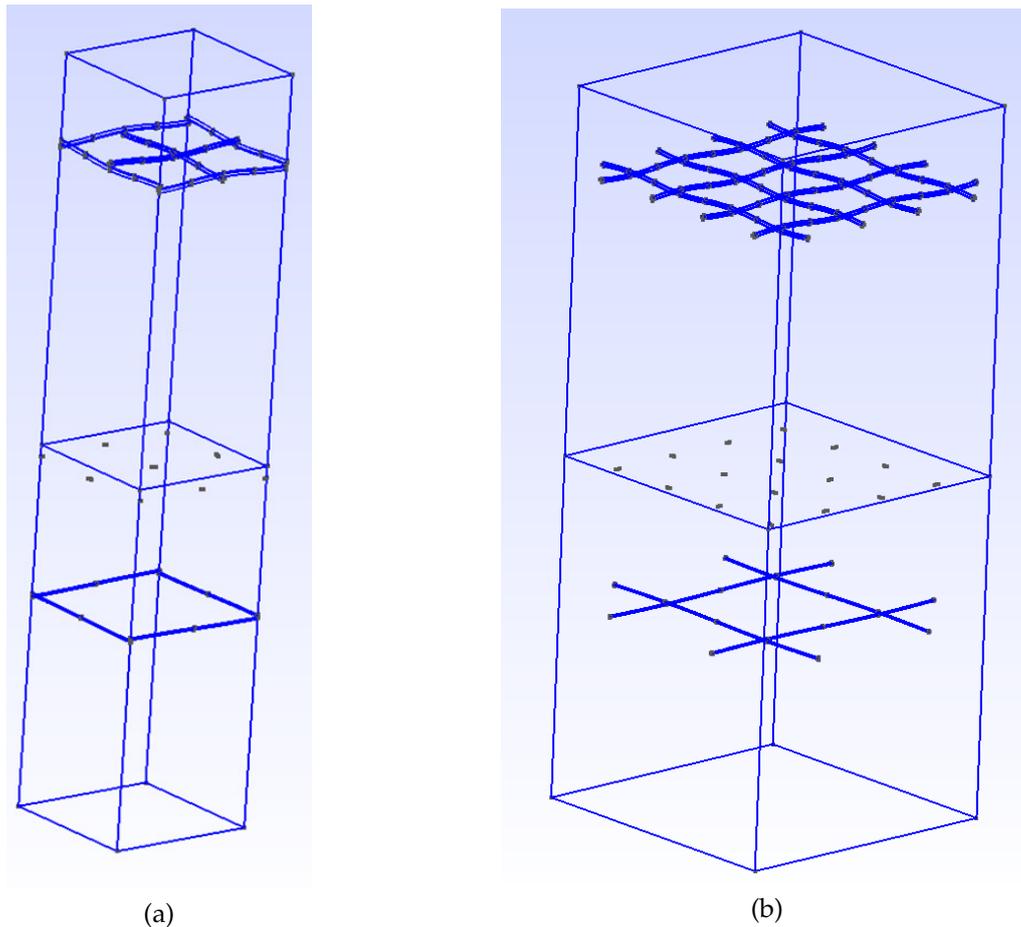


Figure 7.33: Unit cell for the woven mesh, with the anode pitch equal to half of the gate pitch. (a) Aligned grids. (b) Grids offset by half of the pitch.

The other point to note with the anode pitch being half of the gate pitch is that the aligned case focusses electrons onto the anode wires, whereas the offset case focusses them into the centre of an anode unit cell. This is shown in Fig. 7.35, which is looking up at the grids from below.

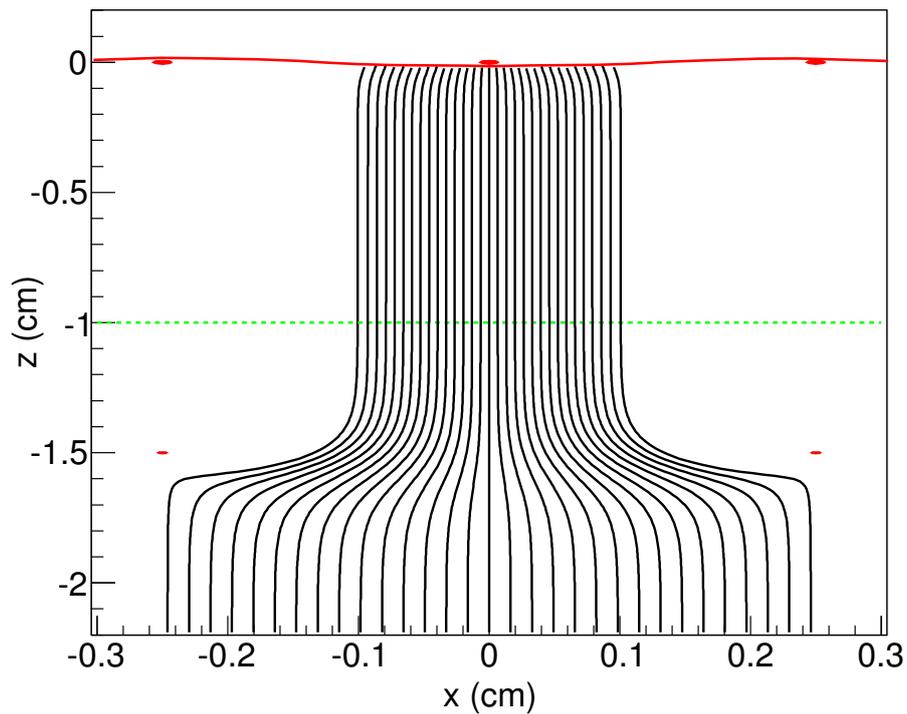


Figure 7.34: Drift lines for the LZ pre-TDR scenario, with the anode pitch equal to half the gate pitch and both grids aligned.

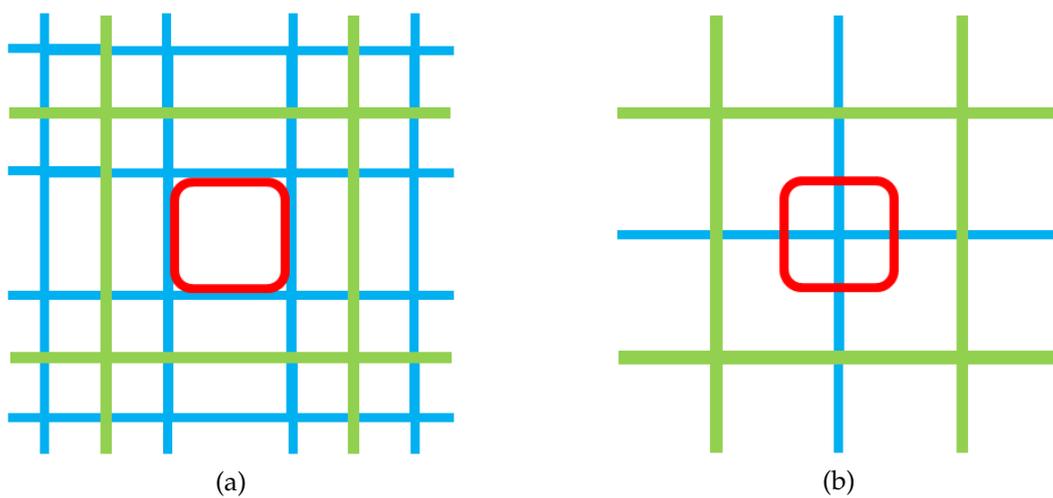


Figure 7.35: View of the anode (blue) and gate (green) meshes looking from below with the anode pitch equal to half the gate pitch. The red square shows the approximate location of electrons on the liquid surface after focusing through the gate. (a) Grids offset by half the gate pitch. (b) Aligned grid wires. Note that there are anode wires underneath all green gate wires shown.

Once the drift lines have been found, the simulation can be shortened even more. The potential at the liquid surface is obtained from the gate-anode field map, and used to create a surface-anode field map. This can use the same mesh density as the previous LUX simulation. This method allows a high quality field map to be made without vastly increasing computation time. The electrons start from points on the surface determined by the drift lines. The same field map is used for both the aligned and the offset case, but the start positions are displaced for the offset case.

7.6.2 Parallel and crossed grid results

The initial simulations of the LZ geometry (see Fig. 7.28a) used single and crossed grids. A range of both wire diameters and pitches were simulated for both. Three pitches were chosen: 1, 2 and 3 mm. Each of these was simulated with three wire diameters: 50.8, 101.6 and 203.2 μm . The results are summarised in Table 7.6 for the parallel wire grids, and Table 7.7 for the crossed grids.

Table 7.6: Results for the mean number of photons per electron (γ/e^-), relative standard deviation (σ/μ), maximum field on the anode ($\text{Max } E_a$) and optical transmission at normal incidence (T_a), for possible parallel wire grids. The maximum field on the wires is calculated using equation 7.15

Pitch \rightarrow Wire Diameter (μm)	1 mm				2 mm				3 mm			
	γ/e^-	σ/μ (%)	$\text{Max } E_a$ (kV/cm)	T_a	γ/e^-	σ/μ (%)	$\text{Max } E_a$ (kV/cm)	T_a	γ/e^-	σ/μ (%)	$\text{Max } E_a$ (kV/cm)	T_a
50.8	529.4 ± 0.2	1.13 ± 0.02	79	94.9	566.5 ± 1.7	9.44 ± 0.21	150	97.5	1028 ± 16	48.0 ± 1.3	221	98.3
101.6	529.1 ± 0.1	0.69 ± 0.01	44	89.8	528.1 ± 0.1	2.64 ± 0.06	79	94.9	554.5 ± 1.8	10.1 ± 0.2	115	96.6
203.2	529.8 ± 0.1	0.55 ± 0.01	26	79.7	525.9 ± 0.2	1.25 ± 0.03	44	89.8	523.4 ± 0.5	3.14 ± 0.07	61.6	93.2

Table 7.7: Results for the mean number of photons per electron (γ/e^-), relative standard deviation (σ/μ), maximum field on the anode ($\text{Max } E_a$) and optical transmission at normal incidence (T_a), for possible crossed wire grids. The maximum field on the wire is calculated by sampling the field near to the wire surface.

Pitch \rightarrow Wire Diameter (μm)	1 mm				2 mm				3 mm			
	γ/e^-	σ/μ (%)	$\text{Max } E_a$ (kV/cm)	T_a	γ/e^-	σ/μ (%)	$\text{Max } E_a$ (kV/cm)	T_a	γ/e^-	σ/μ (%)	$\text{Max } E_a$ (kV/cm)	T_a
50.8	539.8 ± 0.1	0.63 ± 0.02	51	90.1	535.8 ± 0.3	1.94 ± 0.04	81	95.0	539.2 ± 0.7	4.04 ± 0.09	98	96.6
101.6	540.1 ± 0.1	0.72 ± 0.02	33	80.7	536.1 ± 0.2	1.44 ± 0.04	49	90.1	533.3 ± 0.4	2.63 ± 0.06	62	93.3
203.2	540.5 ± 0.1	0.78 ± 0.02	24	63.5	537.3 ± 0.2	1.38 ± 0.04	33	80.7	534.0 ± 0.4	2.13 ± 0.06	40	86.9

The simulations of the parallel and crossed grids show that a relative width of less than 5% is achievable, so that the variance in photon production is small compared to other sources of variance. It can even be near the 1% level observed in the simulations of previous experiments. Next, some example histograms of the photons emitted per electron are shown. Firstly, Fig. 7.36 shows examples for the single wire grids.

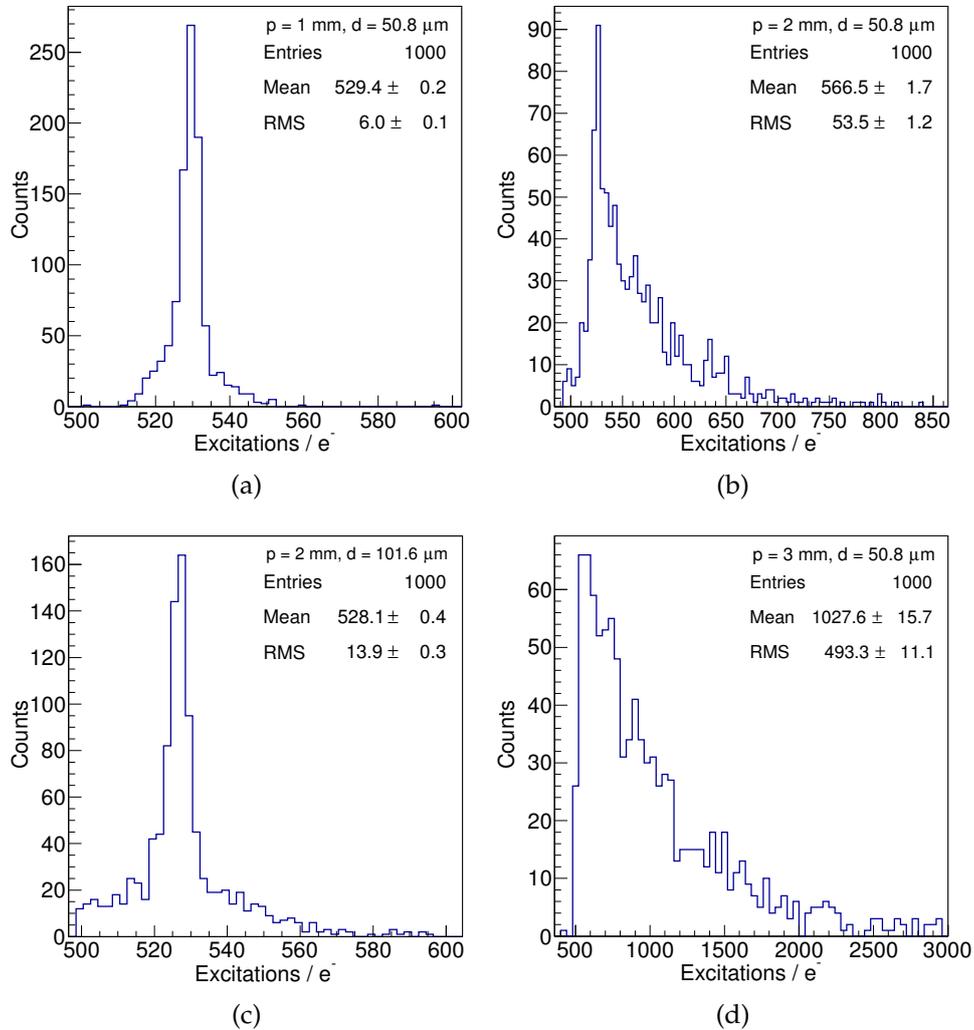


Figure 7.36: Photons emitted per electron for initial simulations with parallel wire anode grids. (a) to (d) show different anode pitches (p) and wire diameters (d).

In the parallel wire grids, even the smallest pitch and wire diameter has some multiplication occurring near the anode, which causes the tail to the right of the main peak as shown in Fig. 7.36a. Despite this, the variance in photon production is acceptable. When compared to the 2 mm pitch with 50.8 μm wires, the increase in the amount of multiplication can clearly be seen, due to the higher field near the wires. Looking at Figs. 7.36b and 7.36c for the 2 mm pitch with larger wires, the tail to the right of the mean decreases. Increasing the wire diameter reduces ionisation near the wire, at the cost of some optical transparency. When going to the largest pitch of 3 mm, both the 50.8 and 101.6 μm wires have significant electron multiplication, and for the former this greatly increases the mean photon yield.

The crossed grids have less electron multiplication. Examples of the photon and ion production are shown in Figs. 7.37 and 7.38 respectively.

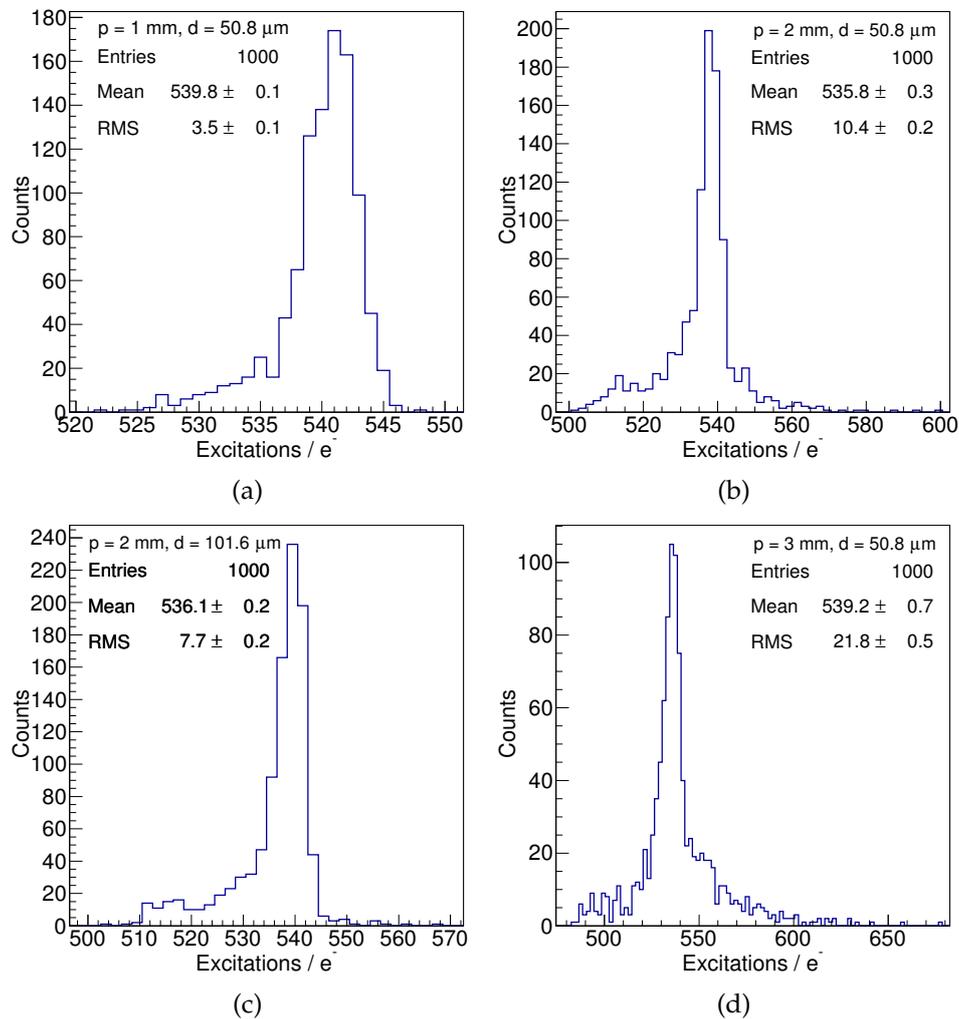


Figure 7.37: Photons emitted per electron for initial simulations with crossed wire anode grids. (a) to (d) show different anode pitches (p) and wire diameters (d).

The improvement when going from parallel to crossed grids can clearly be seen in the photon histograms. The tail at the left of the main peak from the shorter electron paths is still present in these scenarios, but they achieve acceptable variance in EL production. For the crossed grids there is little difference between the 101.6 μm and 203.2 μm diameter wires. Therefore, it is better to use the lower of the two because of the improved optical transparency. The only scenarios with a more significant contribution from electron multiplication are the 2 and 3 mm pitches with 50.8 μm wires, which can be seen in Figs. 7.38b and 7.38d.

In conclusion, both the parallel and crossed grids provide a range of pitches and wire diameters that produce acceptable performance in the EL signal. For single grids the acceptable scenarios with less than 10% relative width in the photon yield were: those with 1,2,3 mm pitch and 203.2 μm wires, although 2 mm pitch with 50.8 μm wires was only just above 10%. The others had increased variance due to electron multiplication (up to 50% relative width in the worst scenario). Despite the reason-

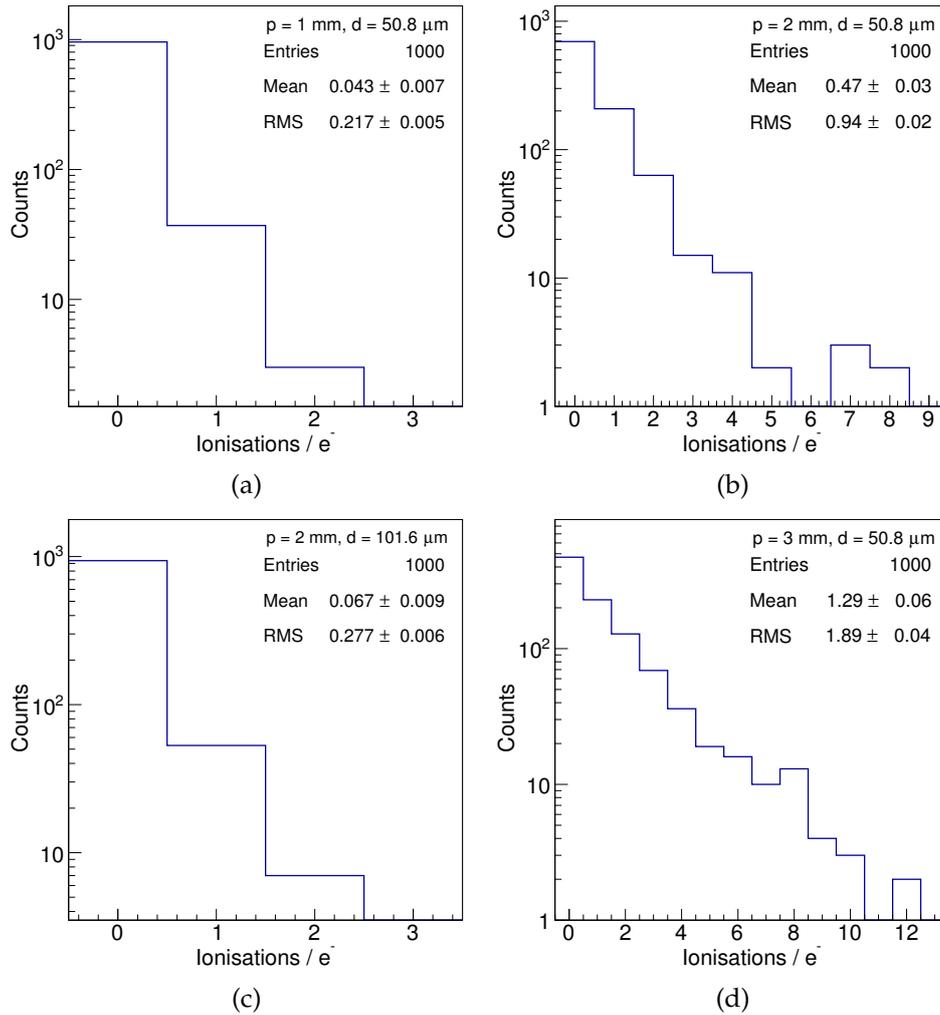


Figure 7.38: Ionisations emitted per electron for initial simulations with crossed wire anode grids. (a) to (d) show different anode pitches (p) and wire diameters (d).

able performance in the EL light production the parallel grids do not have acceptable mechanical properties. For the crossed grids, the maximum relative width in EL yield was 4.0% (with 50.8 μm wires and 3 mm pitch) which would be acceptable although there is electron multiplication occurring. The best scenario when only considering the EL variance and optical transparency is the 1 mm pitch with 50.8 μm wires. The 90.1% optical transparency is high and it achieves a sub-1% resolution. However, the smaller pitch means more wires and the diameter is small. Therefore, it may be difficult to manufacture a grid with these parameters. The other scenarios are still good, a $\sim 2\%$ relative width is possible even with a 2 or 3 mm pitch. If the pitch needs to be increased, it can be around 2 to 3 mm as long as the wire diameter is over 100 μm .

7.6.3 Woven mesh results

The previous results gave some guidance on acceptable anode pitches and wire diameters. The following simulations are now using the subsequent revisions of the LZ geometry, the one used in the CDR (Fig. 7.28b) and then later proposed for the TDR

(pre-TDR, Fig. 7.28c). LZ collaborators at SLAC are able to manufacture a woven mesh with a pitch of around 2 to 3 mm with 150 μm diameter wires although it would be more practical to weave an even larger pitch, which motivated the selection of the 5 mm pitch for the simulations. The following simulations include the focussing effect from the gate, which can improve the resolution compared to the previous simulations. Both the anode and the gate are woven wire meshes.

In addition to finding the relative variance in the S2 signal, the timing information is used to recreate the pulse shape. It was found that there were sometimes gaps in the photon emission, creating a main pulse followed by a secondary “trailing” pulse up to a few μs later. This occurs when the electrons overshoot between the anode wires, and pass through a low field (below the EL threshold) region just above the anode on the way back to that electrode. A basic pulse splitting algorithm was used to determine when the first pulse ends and the trailing pulse begins. It is intentionally conservative towards splitting pulses more. Some example pulses are shown in Fig. 7.39 for the 5 mm pitch. All of the waveforms are shown as photons per 0.01 μs , the expected sample size of the LZ DAQ.

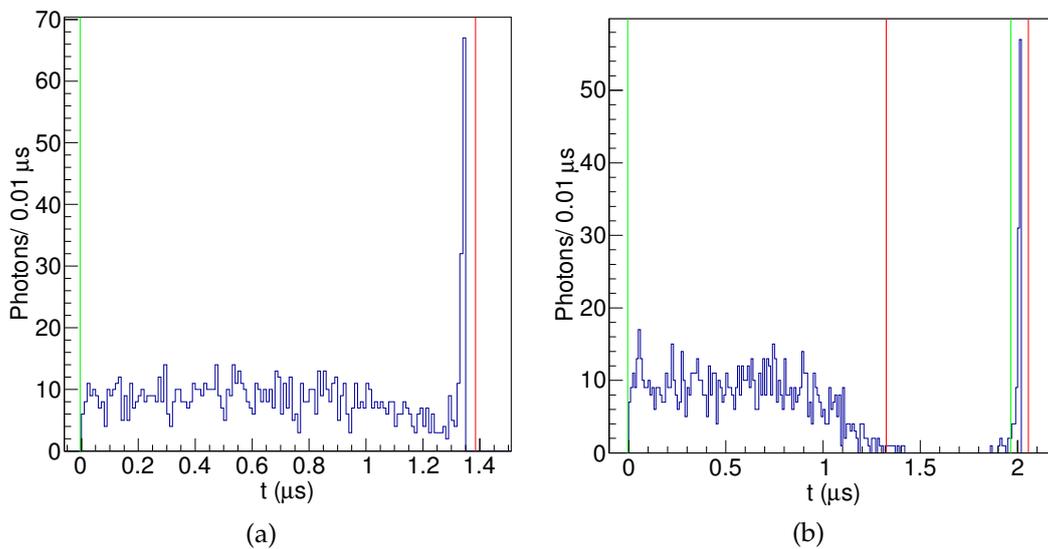


Figure 7.39: Example single electron pulses from the 5 mm pre-TDR geometry scenario with aligned grids, in emitted photons. Green lines show the start of pulses, red are the end times. (a) Pulse where the electron goes straight to the anode wire, peak at the end due to higher field near to the anode. (b) Pulse where the electron goes above the anode into a low field region where there is no emission before producing a short peak as it goes back down to the anode wire. The pulse is split into two, only the first one is used for the pulse area.

The gap in the photon emission is where the electron is above the anode, and the electric field strength is not large enough to cause electroluminescence (<2.5 kV/cm at 1.6 bar). When the electron comes back down onto an anode wire, it experiences a higher field and causes electroluminescence for a short time. The high field near the wire can also be seen in the pulses that are not split and creates a short peak at the end.

The following plots include photon histograms like those shown in the previous section. Here they are split into two colors: blue where the pulse was not split and

red when there was a trailing pulse. If the pulse is split, it reduces the area of the first pulse, as the electroluminescence from the vicinity of the wire is not included.

Firstly, Fig. 7.40 shows the scenarios with a 5 mm anode and gate pitch. Some of these had trailing pulses, the time difference between the trailing pulses and the first pulse is shown in Fig. 7.41.

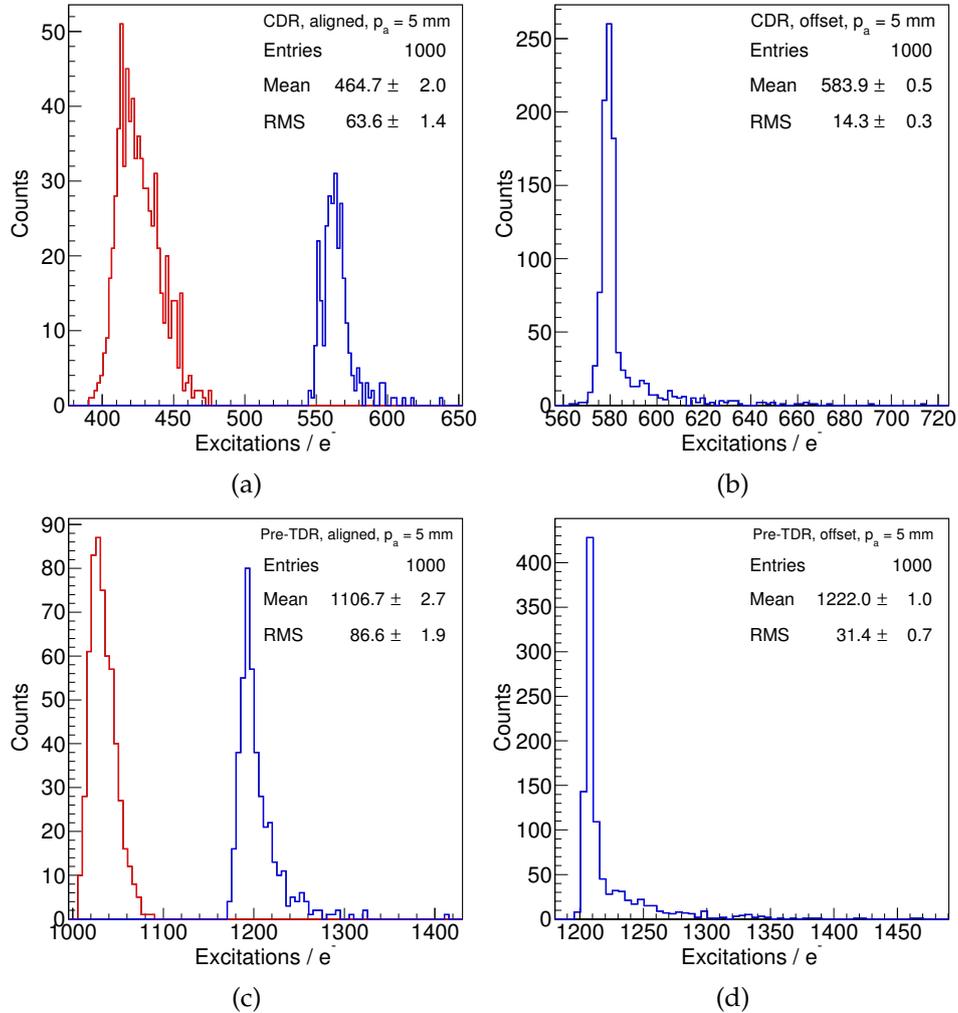


Figure 7.40: Excitations per electron for the geometries with 5 mm anode and gate pitch. (a) and (b) are for the LZ CDR model, with aligned grids, and grids offset by 2.5 mm respectively. (c) and (d) are the same for the pre-TDR model.

This study reveals a clear difference between the aligned and the offset grids. With the gate and anode aligned, the electrons are focussed between the wires and are able to go above the anode grid. The peak for events where there is a separated trailing pulse form a completely separate peak. The time to the trailing pulses is significant, but it is not unacceptable. The longest delay observed was 3.5 μ s. For larger pulses, it is likely that most of these would be included in the first pulse because the difference in arrival time of the electrons will be comparable.

Offsetting the grids removes the trailing pulses completely. Now the gate focusses the electrons towards the wires so they are not able to pass through the gaps in the

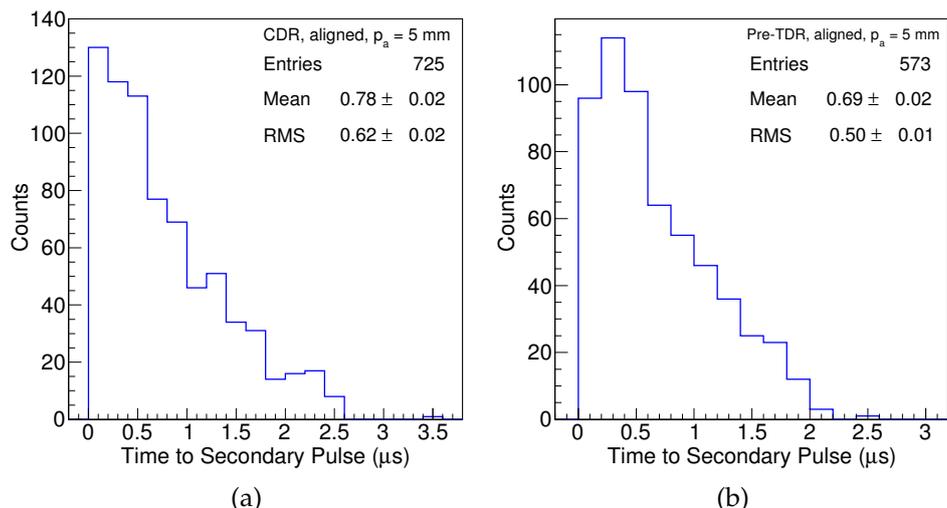


Figure 7.41: Time between first pulse and secondary pulses with 5 mm anode pitch and aligned grids for: (a) LZ CDR scenario and (b) Pre-TDR model.

anode mesh. Comparing the CDR and pre-TDR models with offset grids, the performance is very similar. There is significant ionisation in both, but the relative width is acceptable at 2.45% for the CDR and 2.57% for the pre-TDR model. This is not as good as obtained from the simulations of the other experiments, but here the pitch is much larger and this design is still likely to meet the S2 resolution requirement at high energy.

Because the 5 mm pitch with offset grids was acceptable, an even larger pitch was simulated. The focussing from the gate was a significant improvement, so it is possible that the pitch can be increased further. This would improve the optical transparency and make the grids easier to manufacture. The results from this are shown in Fig. 7.42.

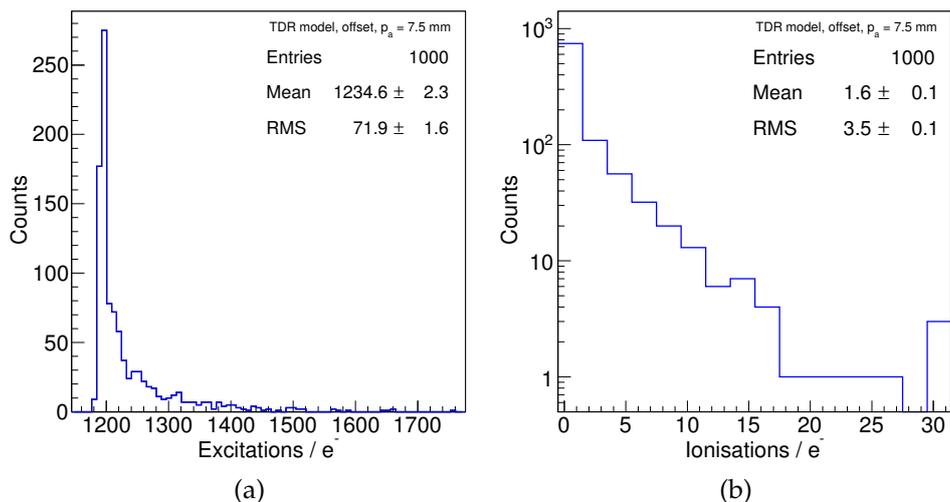


Figure 7.42: Results for the pre-TDR scenario, with a 7.5 mm pitch anode and gate, offset from each other. (a) Shows the number of ionisations per electron and (b) the number of ionisations per primary electron.

This scenario with the offset grids with a 7.5 mm pitch still did not lead to any trailing pulses. However, it can be seen that the variance in photon production has increased to 5.8%, so it is starting to become significant.

The final set of results relates to scenarios with a 2.5 mm anode pitch and a 5 mm gate pitch. In this instance the aligned grid case is the one which focusses the electrons onto the wire (see Fig. 7.35). The variance in photon production is shown in Fig. 7.43 and the time difference to trailing pulses in Fig. 7.44.

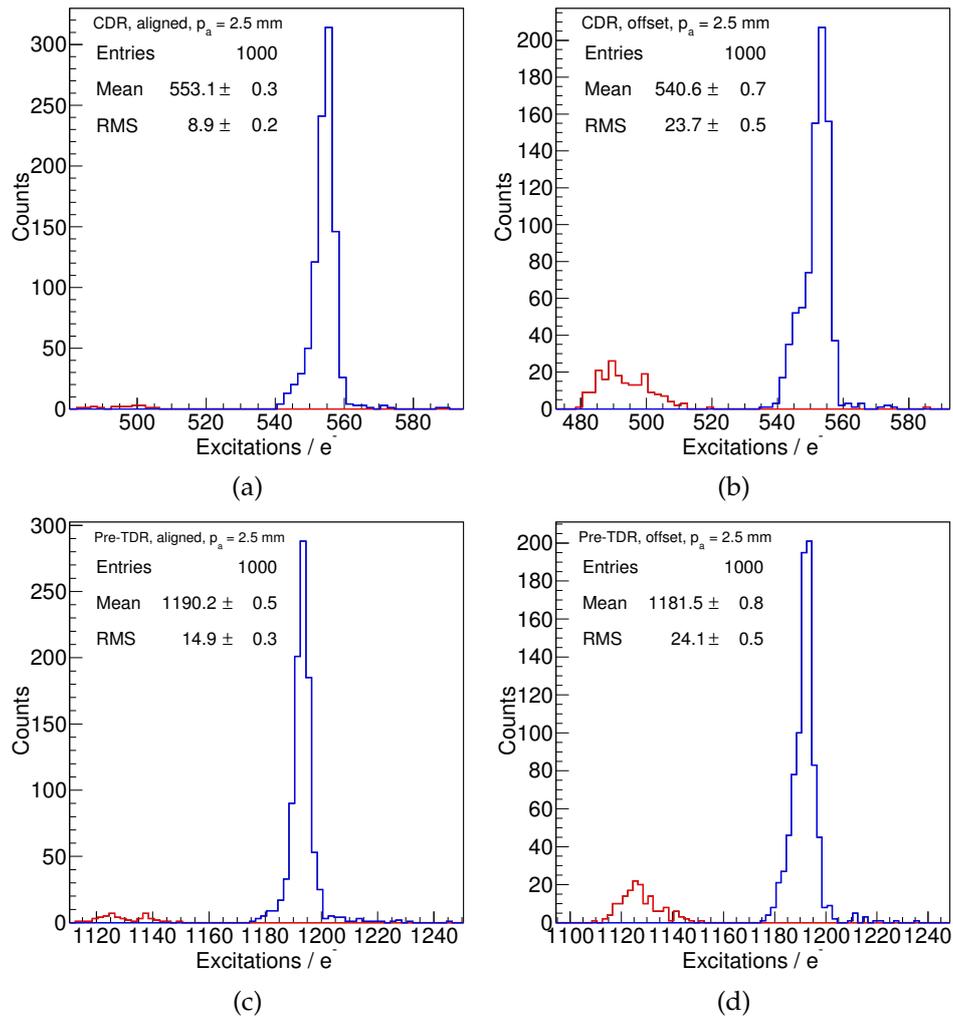


Figure 7.43: Excitations per electron for the scenarios with 2.5 mm anode pitch and 5 mm gate pitch. (a) and (b) are for the LZ CDR scenario, with aligned grids, and grids offset by 2.5 mm respectively. (c) and (d) are the same for the pre-TDR model.

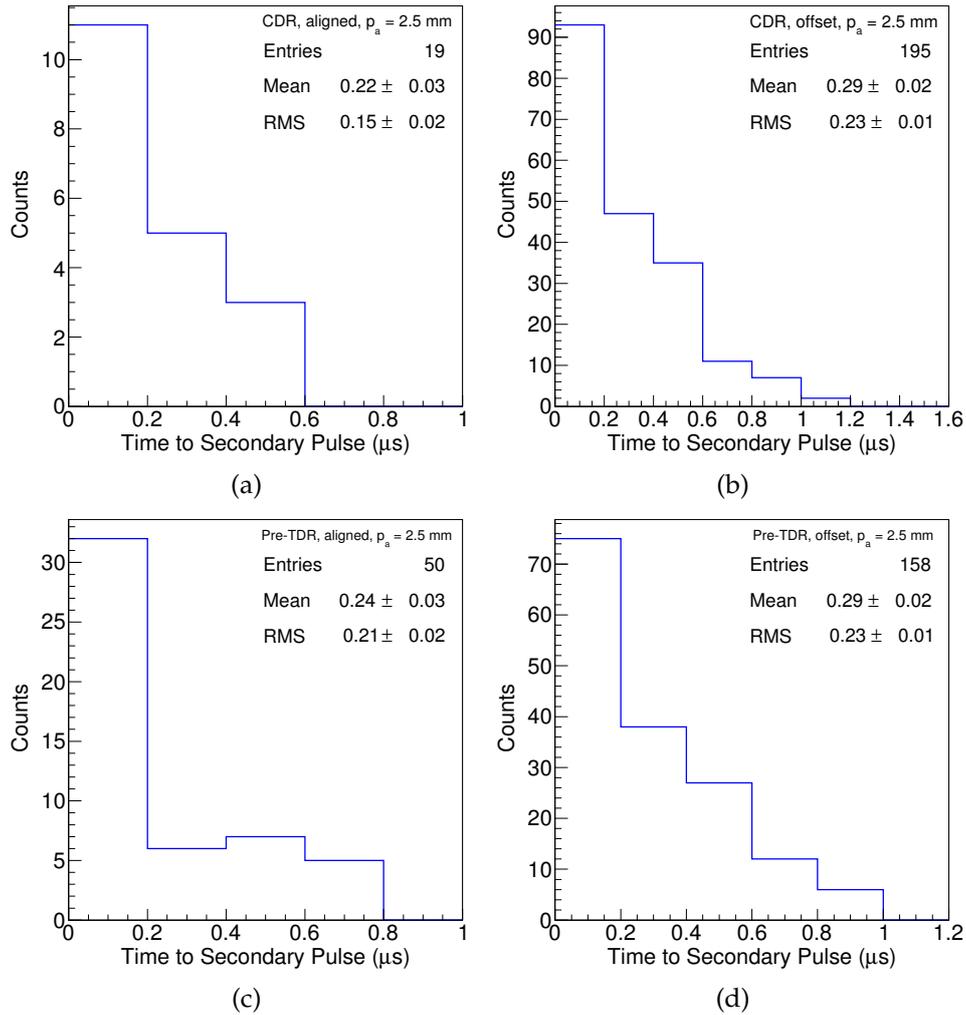


Figure 7.44: Time between first and secondary pulses for LZ scenarios with 2.5 mm anode pitch and 5 mm gate pitch. (a) and (b) are for the CDR, with aligned grids, and grids offset by 2.5 mm respectively. (c) and (d) are the same for the pre-TDR.

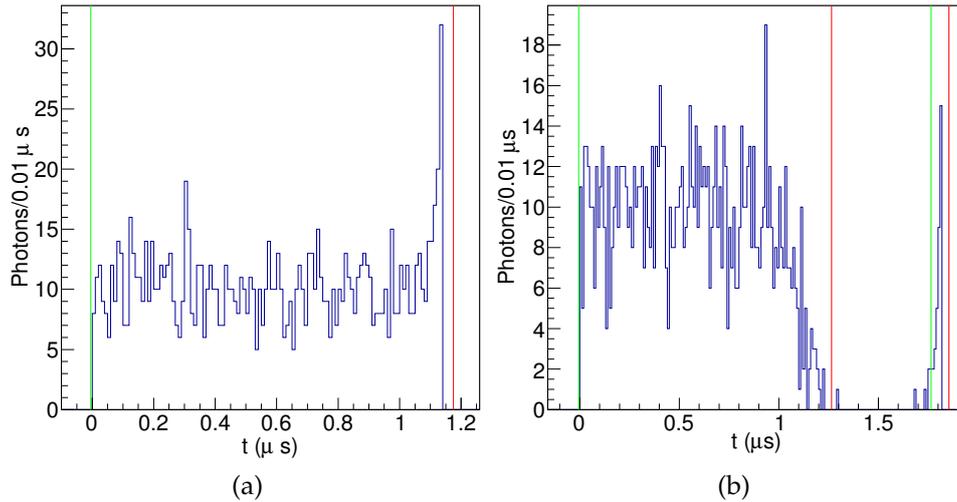


Figure 7.45: Example single electron pulses from the 2.5 mm pre-TDR scenario with aligned grids (in emitted photons). Green lines show the start of pulses, red are the end times. (a) Pulse where the electron goes straight to the anode wire. (b) Pulse where the electron goes above the anode for a short time.

The results from the CDR and pre-TDR scenarios are summarised in Table 7.8. As expected, with the anode pitch being half the gate pitch the aligned grids are better. Some trailing pulses can still be seen, causing a small peak with a lower mean. However, the overall variance is still acceptable even when including this. For the CDR aligned case the relative width is 1.61%, and for the pre-TDR 1.25%. These are comparable to the previous experiments despite the larger pitch. Also, this is a conservative estimate, because in this case the time to the secondary pulses is very small, less than $1 \mu\text{s}$. Some typical pulses are shown in Fig. 7.45. For events of more than one electron, this time is comparable to the longitudinal diffusion. For events at the detector centre the standard deviation on the arrival time at the liquid surface is $0.4 \mu\text{s}$, so these trailing pulses are likely to be clustered into the main pulse.

Table 7.8: Results for the mean number of photons per electron (γ/e^-), relative standard deviation (σ/μ), maximum field on the anode (Max E_a) and optical transmission at normal incidence (T_a), for the LZ woven mesh results. h_g is the gas height, P_G and P_A are the gate and anode pitches, respectively.

Configuration	h_g (mm)	P_G (mm)	P_A (mm)	G-A Voltage (kV)	Align- ment	γ/e^-	σ/μ (%)	Max Field on Wire (kV/cm)		T_a (%)	
								A	G	A	G
CDR	5	5	5	± 4	Aligned	464.7 ± 0.2	13.7 ± 0.31	52	20	94	97
CDR	5	5	5	± 4	Offset	583.9 ± 0.5	2.45 ± 0.05	52	20	94	97
CDR	5	5	2.5	± 4	Aligned	553.1 ± 0.3	1.61 ± 0.04	39	22	88	97
CDR	5	5	2.5	± 4	Offset	540.6 ± 0.7	4.38 ± 0.09	39	22	88	97
Pre-TDR	10	5	5	± 7	Aligned	1106 ± 2.7	7.83 ± 0.17	59	24	94	97
Pre-TDR	10	5	5	± 7	Offset	1222 ± 1.0	2.57 ± 0.06	59	24	94	97
Pre-TDR	10	5	2.5	± 7	Aligned	1190 ± 0.5	1.25 ± 0.03	43	24	88	97
Pre-TDR	10	5	2.5	± 7	Offset	1182 ± 0.5	2.04 ± 0.04	43	24	88	97
Pre-TDR	10	7.5	7.5	± 7	Offset	1235 ± 2.3	5.82 ± 0.13	64	30	96	98

7.6.4 Effects of grid deflection

For the LZ pre-TDR configuration (1.5 cm anode-gate distance with ± 7 kV), the effect of grid deformation on the electroluminescence region performance is investigated. The gate is taken to be a 5 mm pitch (75 μm diameter wires) and the anode 2.5 mm pitch (150 μm diameter wires). A maximum deflection at the centre of the grids is chosen, then the shape of the grid is assumed to be parabolic.

For each maximum displacement of the anode (downwards) and the gate (upwards) we calculate several parameters in an approximate fashion as a function of radial position in the detector, r . The grids are assumed to move by an equal amount. Electric fields are calculated for the new distances, using the method from [293]. Then the photon yield is taken from the parameterisation of Ref. [167] by Ref. [170]. The S2 width is calculated using the reduced electron drift speed in the gas v_r , which is $\approx 1.5 \text{ mm}/\mu\text{s}/(\text{V}/\text{cm}/\text{torr})$.

Three maximum deflections are chosen, up to 3 mm at the centre of each grid. The calculations of each parameter as a function of radial position are shown in Fig. 7.46.

The 1 mm displacement scenario appears acceptable. The field in the gas goes up to 12.3 kV/cm in the centre. The effect of this can be seen on the photon yield which increases from 1150 to 1215 per electron (with this field calculation). The maximum field on the gate is already over 50 kV/cm, and it only increases slightly to a maximum of 66.5 kV/cm. The S2 width is decreased at the centre, which can be a problem for S2-only analyses because it reduces the ability to determine the depth of the interaction from the pulse width. The change is about 30%, going down to just over 1 μs at the centre. When the deflection at the centre is increased further the changes are much more drastic. The field on the gate wires is approaching 100 kV/cm, and the S2 width is almost half in the centre than at the edges of the detector. It is clear that going beyond the 1 mm deflection limit quickly starts to have more severe effects.

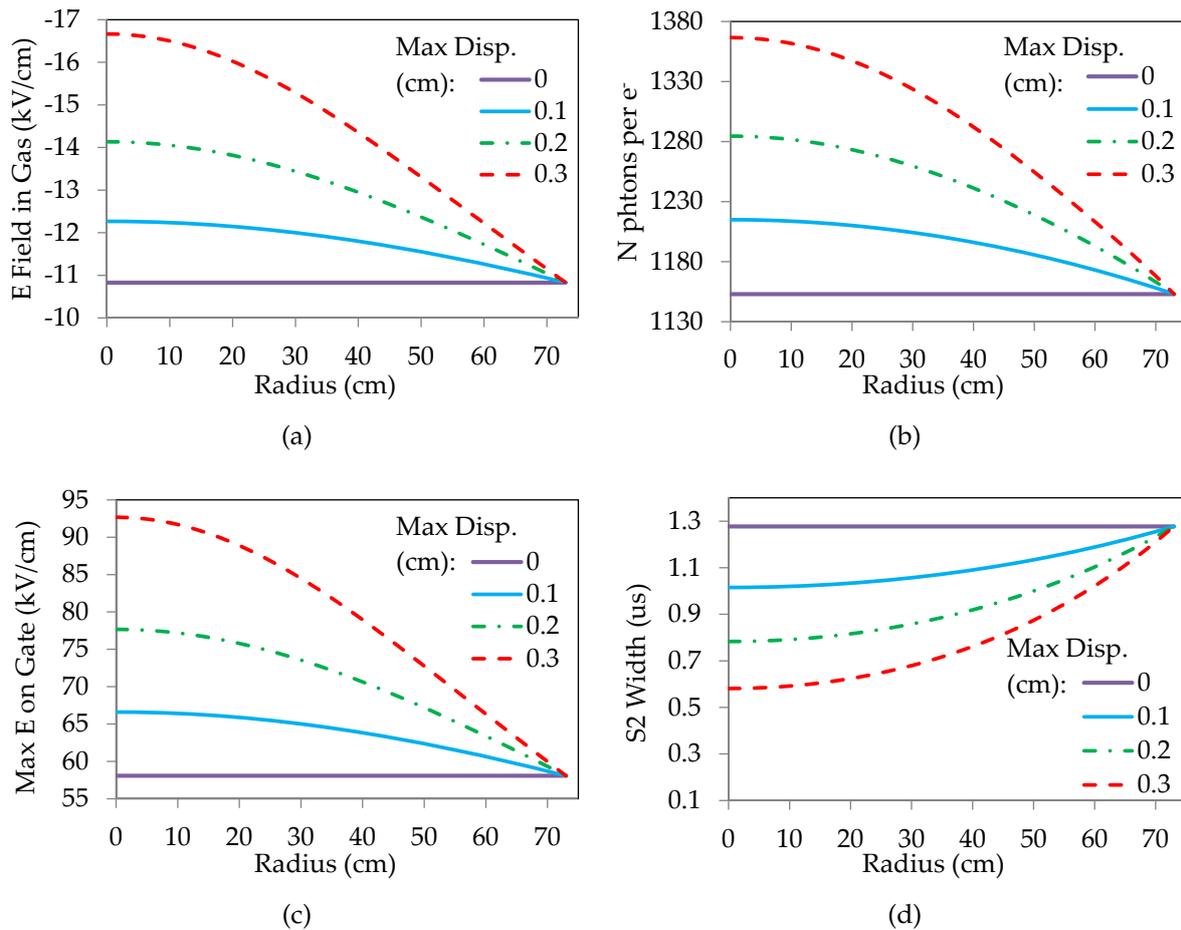


Figure 7.46: Effects of grid deflection on electroluminescence region for the LZ pre-TDR scenario (see Fig. 7.28c), $V_{Gate} = -7$ kV, $V_{Anode} = 7$ kV, gate-anode distance 1.5 cm with 1 cm gas gap. (a) E field in gas. (b) Photon yield per electron. (c) Maximum E field on gate wire surface. (d) S2 width.

7.6.5 Conclusions

An acceptable variance in the single electron response can be achieved even with the larger anode pitches required for LZ. With a 5 mm anode pitch, the variance was small when the gate and anode grids were offset by half the pitch. The CDR version had 2.45% variance and the pre-TDR model yielded a very similar 2.57%. This variance is degraded if the grids are aligned, going to 13.7% and 7.9% respectively. Therefore, if using this pitch it is particularly important that the grids are aligned properly. The reason the aligned meshes give worse results is due to gaps in the electroluminescence production, when the electrons go above the anode before dropping back down onto the wires.

The variance is improved when using an anode pitch of 2.5 mm. The best scenario here is to have aligned grids: the pre-TDR model gives 1.26% variance and the CDR configuration 1.59%. For the offset pre-TDR case it is only slightly worse, with 2.03% variance. These numbers are slightly conservative, due to the trailing pulses. In these scenarios, the maximum gap to the trailing pulses is ~ 1 μ s, so for events with multiple

electrons these are likely to be included in the main pulse. Therefore, the best scenario to use is the 2.5 mm pitch with the anode and gate aligned.

After full LZ sensitivity studies were conducted with the CDR and pre-TDR scenarios, an intermediate configuration was selected to decrease the S2 size (for linearity reasons) and the applied voltages (7 kV is somewhat aggressive). The actual model described in the TDR [130], which is the LZ baseline design, reduces the gate-anode distance to 13 mm, the gas gap to 8 mm and the electrode voltages to ± 5.5 kV. The results presented above are qualitatively unchanged.

7.7 Energy Resolution for Larger Signals

Now that scenarios which give an acceptable single electron response for uniformly distributed electrons have been found, it is important to confirm that this is acceptable for localised interactions at higher energies too. This should be the case, as higher electron statistics will improve the relative resolution, which must involve the combination of S1 and S2 signals since these are anti-correlated for electron recoils. However, EL production may still have an effect due to the position dependence. Different positions in x - y below the gate lead to slightly different EL yields. For the remainder of this section, the pre-TDR configuration was chosen.

In order to build large events, the Garfield++ simulations are combined with the FastNEST simulations. As described previously in section 7.1.3, FastNEST accounts for the production of initial quanta, liquid purity, extraction efficiency, EL light production and photon detection. The standard method of EL light production (sampling from a Gaussian) is replaced by the more comprehensive treatment described above. Also, transverse diffusion in the liquid is added and the ability to choose the x - y position of events. For the chosen scenario, simulations are conducted for a number of x - y positions in the unit cell. For each, a histogram of the number of photons per extracted electron can be created. These are all read into FastNEST. When an electron reaches just below the gate (where the calculated electron drift trajectories in the liquid start), the nearest drift line is found. Then the photon histogram corresponding to the location where that drift line reaches the liquid surface is selected and the EL yield for each electron is sampled at random from that histogram.

First, events from neutrinoless double beta decay are studied; their energy is some 1000 times higher than the region of interest for WIMP searches and so they probe the intrinsic resolution of the LZ detector. These are dual (back-to-back) ERs with $\sim 1.2 \times 10^5$ ionisation electrons drifted away from the vertex, which are sampled from a 3D Gaussian with a standard deviation of 3 mm. However, in order to see the effects of the position dependence, vertices are first generated from particular positions at the top of the liquid with no spread in their start positions. The centre and corner of the unit cell (corner of a gate grid square) are chosen and the resulting S2 (z position corrected for losses to electronegative impurities) sizes are shown in Fig. 7.47.

The relative width of both of these come out the same at 4.36%. The mean values are also very similar, the difference between them is 1.76% of the mean yield from the corner of the gate cell. The electrons coming from the corner of the gate cell (between

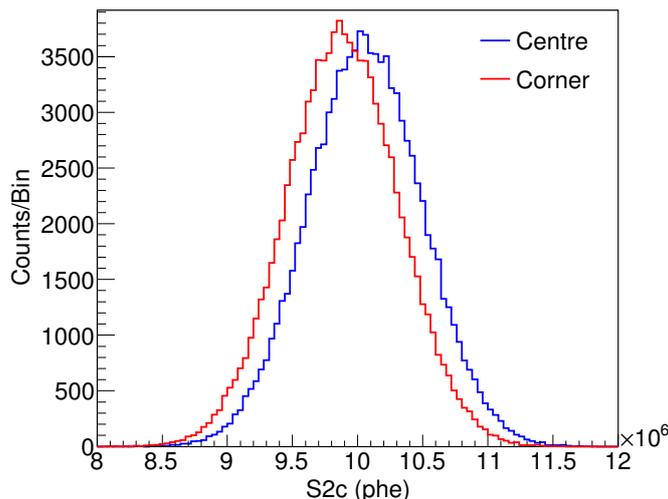


Figure 7.47: S2 signals from NDBD events coming from the top of the detector, at different x - y positions (z position corrected). Blue is from the centre of the gate unit cell, or where two anode wires cross and red from the corner of the gate unit cell, where electrons are focused into the region between anode wires (see red square on Fig. 7.35b).

anode wires) give a smaller yield due to some of the pulses being split; electrons from this position are more likely to go above the anode. The emission after the gap is not included in the pulse area here. Conversely, the electrons that go straight to the anode wire do not have a gap in the EL emission so have a larger pulse area. Nevertheless, the difference between the yields from the two starting positions is small, even for these two extreme cases. Once the spread in position (from the initial positions in the NDBD event and transverse diffusion in the liquid) is included in the simulation, many points on the grids will be sampled at once so this effect will be even smaller.

Now the NDBD events are simulated with the relevant effects taken into account. The start position of the event is sampled at random, both in x - y and z between 55 and 818 μ s to cut out regions at the top and bottom expected to be outside of the (WIMP-search) fiducial volume. The position corrected S1 and S2 yields are shown in Fig. 7.48.

The relative S2 distribution width is 4.37% (S1 distribution width 8.6%). The combined (S1+S2) energy resolution can be obtained from the event-by-event S1 and S2 responses using the standard formula which accounts for their anti-correlation:

$$E = W \left(\frac{S1_c}{g_1} + \frac{S2_c}{g_2} \right) . \quad (7.21)$$

The required quantities are taken directly from FastNEST — so it is assuming each is perfectly known. The S1 photon detection efficiency $g_1 = 0.075$ phd/photon. The S2 yield g_2 , in phd/electron, is the product of the extraction efficiency (95%), mean EL yield (1190 photons per extracted electron) and photon detection efficiency in the gas (also assumed to be 7.5%). The reconstructed energy for NDBD events is shown in Fig. 7.49.

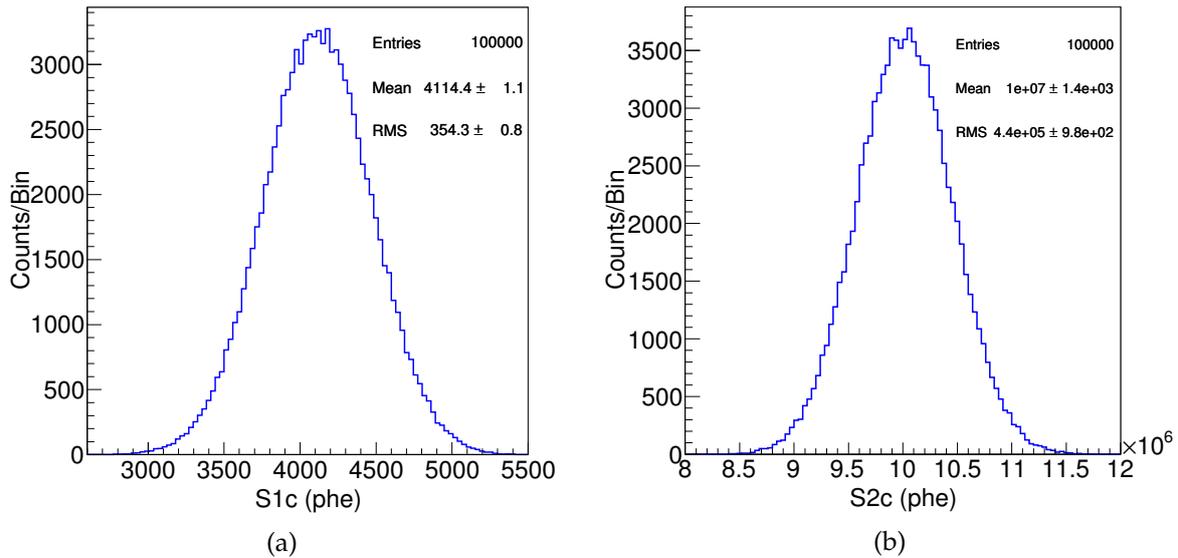


Figure 7.48: Simulated S1 and S2 yields from NDBD events in LZ, using FastNEST plus EL response modelling. The start position is randomly sampled in the fiducial volume. (a) S1 yield, z position corrected. (b) S2 yield, z position corrected.

The overall energy resolution at the ^{136}Xe Q-value is 0.64%. There is some bias from the expected mean value, due to the simple version of the energy reconstruction formula used. As expected, the resolution remains small even with the position dependence included. Repeating the simulation with the EL variance fixed to zero had a negligible effect. Therefore, the grid design does not degrade the resolution at high energy. This means that the lower estimate of the LZ NDBD sensitivity of 4×10^{26} years should be achievable with enriched xenon [197], as long as the radiopurity of detector materials and spatial resolution of events are also acceptable.

Now that the resolution at the NDBD energy has been investigated, the same simulation method is used for smaller events, as the combined S1 + S2 resolution was not previously analysed. Lower energy ER events at 5 and 10 keV are simulated. These have 11.7 and 26.5 S1 phe respectively, within the region of interest for standard WIMP searches. All electrons within an event are assumed to be from the same position in the liquid. Again, the vertex position is taken at random from within the fiducial volume.

The energy resolution for ERs at 5 and 10 keV is 14.5% and 10.9%, respectively. The simulation was repeated with a high electron lifetime of 5 ms and the energy resolutions improved slightly to 14.1% and 10.7%, respectively. With the EL variance removed ($\sigma_{EL} = 0$) the energy resolution was 14.4% at 5 keV and remained at 10.9% for 10 keV. Therefore, the variance in the EL yield is not a significant contribution to the overall energy resolution.

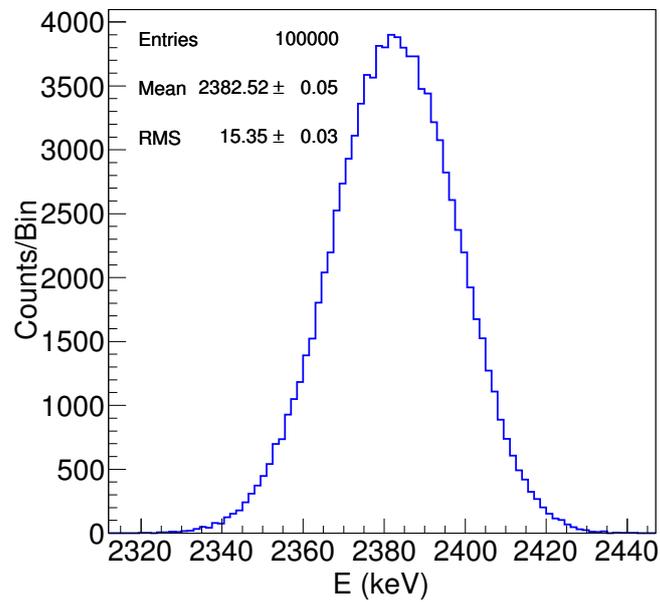


Figure 7.49: Reconstructed energy of NDBD events in LZ from FastNEST plus EL response modelling.

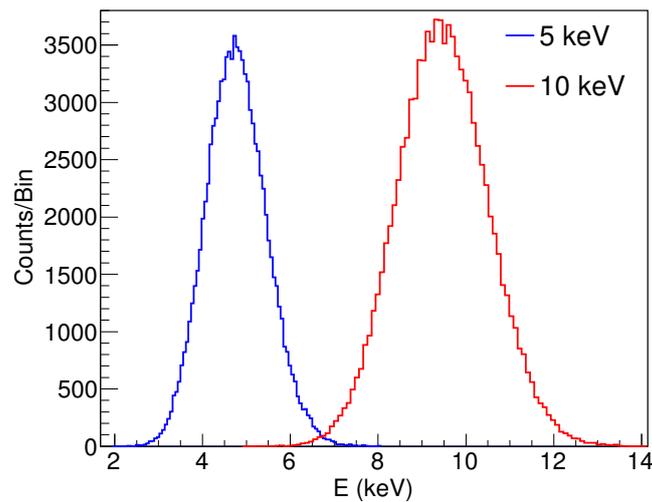


Figure 7.50: Reconstructed energy for ER events in LZ from FastNEST simulations plus EL response modelling.

7.8 Summary

In this chapter the design of the EL region and was investigated, to determine the impact on the EL light production and resolution. This was done by simulating single electron drifts from the liquid surface to the anode in various EL region geometries. The simulation was validated by comparing the EL response to experimental data; the agreement between the two was very good. In particular, the simulation reproduced LUX data of single electron emission originating from a gate wire (with the gate in the gas phase). In this scenario multiplication occurred, which means that the photon yield, ionisation yield, and their variances are accurate.

A FastNEST simulation was used to assess the impact of the relative width of the single electron distribution for LZ on the discrimination. NR and ER events were simulated with a Gaussian EL yield, with various widths. No effect on the discrimination was observed until the relative width in the EL photon production reached 10%.

Single electron drift simulations in geometries of the ZEPLIN-III, LUX, and XENON-100 experiments were carried out, which showed that a relative width of 1% in the EL photon yield was achieved. The LUX anode, a very fine woven mesh was found to have no measurable effect on the EL light production variance.

Two LZ EL region designs were considered, with a gas gap of either 5 or 10 mm as used in the CDR [197] and pre-TDR [130] respectively. For each, various pitches were chosen based on practical considerations. The single electron response was investigated for each scenario. For a 5 mm gate and anode pitch, the EL photon production resolution was acceptable (2.5%) as long as the grids were offset relative to each other (with the crossing point of the gate wires in centre of an anode square). For a 5 mm gate pitch with a 2.5 mm anode pitch the aligned grids gave the best resolution of 2% for the pre-TDR scenario, but the resolution for offset grids was only slightly larger.

For the pre-TDR scenario with a 2.5 mm anode pitch higher energy events events containing many electrons were simulated by combining FastNEST with the detailed EL simulations conducted in this study. At the Q-value of the ^{136}Xe NDBD the combined (S1 + S2) energy resolution was 0.64%, well below the 2.0% requirement. It was also confirmed that the variance in EL light production did not degrade the energy resolution for lower energy ERs, and so it will not affect discrimination.

After this study was conducted, the configuration actually adopted for the TDR (LZ baseline) changed: the gas gap was reduced to 8 mm with a gate-anode distance of 13 mm. The 2.5 mm anode pitch with a 5 mm gate pitch was adopted with the grids aligned. This study shows that the 5 and 10 mm gas gaps were acceptable with these grid parameters, so the intermediate value will be too.

Chapter 8

Conclusions and Outlook

Although LUX has not detected a WIMP signal, it has been very successful. It is the most sensitive direct dark matter experiment to date, which set world-leading limits and explored new regions of WIMP parameter space.

The first WIMP search result from LUX in 2013 was discussed in Chapter 4. Data were taken over 85.3 live days with a 118.3 kg fiducial mass. Extensive calibration was carried out, including: ^{83m}Kr , tritiated methane, AmBe, and ^{252}Cf . The analysis was non-blind, and therefore only used a minimal set of analysis cuts to avoid bias. Conservative assumptions were also made for the detector response to low energy nuclear recoils. This provided a robust result, that still achieved a relatively low energy threshold. An upper limit was set on the spin-independent WIMP-nucleon cross section with a minimum of $7.6 \times 10^{-46} \text{ cm}^2$ at a WIMP mass of 33 GeV (90% CL) [185]. This result excluded regions favoured by possible light WIMP signals found in other direct detection experiments.

Various improvements were made after the first analysis, which motivated a reanalysis of the 2013 data. New low energy NR calibration from a DD neutron generator source allowed the detector response to be evaluated down to 1.1 keV, down from 3 keV used previously. Other improvements included improvements in: background modelling, position reconstruction algorithms, and event selection. The modelling of backgrounds originating on the detector walls allowed the fiducial mass to be increased to 145.4 kg. Extra data were added, taking the live time up to 95 days. The reanalysis improved the spin-independent exclusion limit, with a minimum of $5.6 \times 10^{-46} \text{ cm}^2$ at 33 GeV (90% CL) [148]. It also motivated the first spin-dependent analysis from LUX, which was discussed in Chapter 5. This was a world-leading limit on the spin-dependent WIMP-neutron scattering cross section, excluding down to $9.4 \times 10^{-41} \text{ cm}^2$ at 33 GeV WIMP mass (90% CL) [231]. Using the latest calculations of the spin structure functions, there is also significant sensitivity to the spin-dependent WIMP-proton cross section in xenon. With this result possible light WIMP signals were excluded in both the spin-independent and the spin-dependent channels by LUX.

In Chapter 6 a study of spurious electron emission processes from the cathodic gate wires was undertaken, using special LUX data with increased voltages on the gate and the anode. The localised electron emission is the limiting factor for the maximum voltage in dual-phase TPCs. If the process is understood and mitigated it leads to improvements in both the drift field and the extraction field. The emission was found to be a few electrons and photons from the gate at once with a similar number of each.

The electrons were found to follow a Polya function, which suggests an electron multiplication process is taking place. These effects could be consistent with an enhanced electric field at the emitter location. However, the emission rate was found to increase for a short time before dropping again as the field on the gate wires was increased. With a locally enhanced field the emission should rise sharply, and so this is instead some kind of resonant process. This could be caused by an impurity on the wire surface, such as an insulator or oxide layer. The exact process is still unclear but it should produce resonant effects with electron multiplication and scintillation. It could possibly be a combination of an impurity with a small field enhancement.

For LZ the electroluminescence region design was investigated in Chapter 7. Single electron drifts from the liquid surface to the anode were simulated in various possible designs. The impact of the design on detector resolution was assessed at low energy for WIMP searches, and at high energy for spectroscopy of background γ -rays and neutrinoless double beta decay searches. The simulation software was validated against experimental data. It was found that despite LZ having sparser grids than LUX, the grids did not degrade the overall energy resolution.

The last WIMP search run of LUX was carried out in 2016, with an exposure of 33 500 kg-day from 332 live days. This sets the most stringent limit on the spin-independent WIMP-nucleon scattering cross section, $2.2 \times 10^{-46} \text{ cm}^2$ for a 50 GeV WIMP (90% CL) [1]. New exclusion limits on the spin-dependent cross sections are expected soon. Even though dark matter still evades detection, there are many prospects in direct detection, indirect detection, and collider experiments. LZ is one of the experiments, due to start commissioning in 2020, with an expected sensitivity of $2.5 \times 10^{-48} \text{ cm}^2$ at 40 GeV [130]. There is a real possibility that upcoming experiments could detect a WIMP. This would be a major discovery, and a direct detection would prove the existence of particle dark matter in the universe.

Bibliography

- [1] D. S. Akerib *et al.*, “Results from a search for dark matter in LUX with 332 live days of exposure,” *arXiv:1608.07648*, 2016.
- [2] A. Tan *et al.*, “Dark Matter Results from First 98.7-day Data of PandaX-II Experiment,” *arXiv:1607.07400*, 2016.
- [3] K. G. Begeman, A. H. Broeils, and R. H. Sanders, “Extended rotation curves of spiral galaxies: dark haloes and modified dynamics,” *Mon. Not. R Astron. Soc.*, **249** 3, 523–537, 1991. DOI: 10.1093/mnras/249.3.523.
- [4] M. Persic, P. Salucci, and F. Stel, “The universal rotation curve of spiral galaxies – I. The dark matter connection,” *Mon. Not. R Astron. Soc.*, **281** 1, 27–47, 1996. DOI: 10.1093/mnras/278.1.27.
- [5] K. Freese, “Review of observational evidence for dark matter in the universe and in upcoming searches for dark stars,” *EAS Publications Series*, **36**, 113–126, 2009. DOI: 10.1051/eas/0936016.
- [6] F Zwicky, “Die Rotverschiebung von extragalaktischen Nebeln,” *Helvetica Physica Acta*, **6**, 110–127, 1933.
- [7] A. B. Newman, T. Treu, R. S. Ellis, *et al.*, “The Density Profiles of Massive, Relaxed Galaxy Clusters. I. the Total Density Over Three Decades in Radius,” *Astrophys. J.*, **765** 1, 24, 2013. DOI: 10.1088/0004-637X/765/1/24.
- [8] N. A. Bahcall and X. Fan, “A lightweight universe?” *Proc. Natl. Acad. Sci. USA*, **95** 11, 5956–5959, 1998. eprint: <http://www.pnas.org/content/95/11/5956.full.pdf>.
- [9] P. Gorenstein and W. Tucker, “Astronomical Signatures of Dark Matter,” *Advan. High Energy Phys.*, **2014** 878203, 2014. DOI: <http://doi:10.1155/2014/878203>.
- [10] A. Cavaliere and R. Fusco-Femiano, “X-rays from hot plasma in clusters of galaxies,” *Astron. & Astrophys.*, **49**, 137–144, 1976.
- [11] P. Gorenstein, D. Fabricant, K. Topka, and F. R. Harnden Jr, “Soft X-ray structure of the Perseus cluster of galaxies,” *Astrophys. J.*, **224**, 718–723, 1978. DOI: <http://dx.doi.org/10.1086/156421>.

- [12] A. Vikhlinin, A. Kravtsov, W. Forman, *et al.*, "Chandra sample of nearby relaxed galaxy clusters: mass, gas fraction, and mass-temperature relation," *Astrophys. J.*, **640** 2, 691, 2006. DOI: <http://dx.doi.org/10.1086/500288>.
- [13] L. P. David, C. Jones, and W. Forman, "The Universal Gas Mass Fraction in Clusters of Galaxies," *Astrophys. J.*, **748**, 120, 2012. DOI: [10.1088/0004-637X/748/2/120](https://doi.org/10.1088/0004-637X/748/2/120).
- [14] R. B. Metcalf and L. A. Moustakas, "Spectroscopic gravitational lensing and limits on the dark matter substructure in Q2237+ 0305," *Astrophys. J.*, **607**, 43–59, 2004.
- [15] L. A. Moustakas and R. B. Metcalf, "Detecting dark matter substructure spectroscopically in strong gravitational lenses," *Mon. Not. R Astron. Soc.*, **339** 3, 607–615, 2003. DOI: [10.1046/j.1365-8711.2003.06055.x](https://doi.org/10.1046/j.1365-8711.2003.06055.x).
- [16] J. A. Tyson, G. P. Kochanski, and I. P. Dell'Antonio, "Detailed mass map of CL 0024+ 1654 from strong lensing," *The Astrophysical Journal*, **498**, 107–110, 1998. DOI: <http://dx.doi.org/10.1086/311314>.
- [17] H. Hoekstra, H. K. C. Yee, and M. D. Gladders, "Current status of weak gravitational lensing," *New Astron. Rev.*, **46** 12, 767–781, 2002. DOI: [http://dx.doi.org/10.1016/S1387-6473\(02\)00245-2](http://dx.doi.org/10.1016/S1387-6473(02)00245-2).
- [18] M. Bartelmann and P. Schneider, "Weak gravitational lensing," *Phys. Rep.*, **340**, 291–472, 2001. DOI: [http://dx.doi.org/10.1016/S0370-1573\(00\)00082-X](http://dx.doi.org/10.1016/S0370-1573(00)00082-X).
- [19] R. Gavazzi, T. Treu, J. D. Rhodes, *et al.*, "The sloan lens ACS survey. IV. The mass density profile of early-type galaxies out to 100 effective radii," *Astrophys. J.*, **667**, 176–190, 2007. DOI: <http://dx.doi.org/10.1086/519237>.
- [20] R. Massey, T. Kitching, and J. Richard, "The dark matter of gravitational lensing," *Rep. Prog. Phys.*, **73**, 086901, 2010. DOI: [10.1088/0034-4885/73/8/086901](https://doi.org/10.1088/0034-4885/73/8/086901).
- [21] D. Clowe, M. Bradač, A. H. Gonzalez, *et al.*, "A direct empirical proof of the existence of dark matter," *Astrophys. J.*, **648** 2, L109, 2006. DOI: <http://dx.doi.org/10.1086/508162>.
- [22] M. J. Jee, H. C. Ford, G. D. Illingworth, *et al.*, "Discovery of a Ringlike Dark Matter Structure in the Core of the Galaxy Cluster Cl 0024+17," *The Astrophysical Journal*, **661** 2, 728, 2007. DOI: <http://dx.doi.org/10.1086/517498>.
- [23] M. Bradač, S. W. Allen, T. Treu, *et al.*, "Revealing the properties of dark matter in the merging cluster MACS J0025.4–1222," *Astrophys. J.*, **687**, 959–967, 2008. DOI: <http://dx.doi.org/10.1086/591246>.

- [24] J. Merten, D. Coe, R. Dupke, *et al.*, "Creation of cosmic structure in the complex galaxy cluster merger Abell 2744," *Mon. Not. R Astron. Soc.*, **417** 1, 333–347, 2011. DOI: 10.1111/j.1365-2966.2011.19266.x.
- [25] M. Markevitch, "Chandra observation of the most interesting cluster in the Universe," *arXiv:astro-ph/0511345*, 2005.
- [26] Y. Sofue, M. Honma, and T. Omodaka, "Unified rotation curve of the Galaxy Decomposition into de Vaucouleurs bulge, disk, dark halo, and the 9-kpc rotation dip," *Pubs. Astron. Soc. Japan*, **61** 2, 227–236, 2009. DOI: 10.1093/pasj/61.2.227.
- [27] M. Honma, T. Bushimata, Y. K. Choi, *et al.*, "Astrometry of galactic star-forming region sharpless 269 with vera: parallax measurements and constraint on outer rotation curve," *Publ. Astron. Soc. Japan*, **59** 5, 889–895, 2007. DOI: 10.1093/pasj/59.5.889.
- [28] X. X. Xue, H. W. Rix, G. Zhao, *et al.*, "The Milky Way's Circular Velocity Curve to 60 kpc and an Estimate of the Dark Matter Halo Mass from the Kinematics of ~ 2400 SDSS Blue Horizontal-Branch Stars," *Astrophys. J.*, 2003, 1143–1158, 2008. DOI: <http://dx.doi.org/10.1086/589500>.
- [29] G. Battaglia, a. Helmi, H. Morrison, *et al.*, "The radial velocity dispersion profile of the Galactic halo: constraining the density profile of the dark halo of the Milky Way," *Mon. Not. R Astron. Soc.*, **364** 2, 433–442, 2005. DOI: 10.1111/j.1365-2966.2005.09367.x.
- [30] M. I. Wilkinson and N. W. Evans, "The present and future mass of the Milky Way halo," *Mon. Not. R Astron. Soc.*, **310** 3, 645–662, 1999. DOI: 10.1046/j.1365-8711.1999.02964.x.
- [31] P. J. McMillan, "Mass models of the milky way," *Mon. Not. R Astron. Soc.*, **414** 3, 2446–2457, 2011. DOI: 10.1111/j.1365-2966.2011.18564.x.
- [32] M. Weber and W. De Boer, "Determination of the local dark matter density in our galaxy," *Astron. & Astrophys.*, **509** A25, 2010. DOI: <http://dx.doi.org/10.1051/0004-6361/200913381>.
- [33] R. Catena and P. Ullio, "A novel determination of the local dark matter density," *J. Cosmol. Astropart. Phys.*, **08** 004, 2010. DOI: 10.1088/1475-7516/2010/08/004.
- [34] L. E. Strigari and R. Trotta, "Reconstructing WIMP properties in direct detection experiments including galactic dark matter distribution uncertainties," *J. Cosmol. Astropart. Phys.*, **11** 019, 2009. DOI: 10.1088/1475-7516/2009/11/019.
- [35] J. I. Read, "The local dark matter density," *J. Phys. G: Nuc. Phys.*, **41** 6, 063101, 2014. DOI: 10.1088/0954-3899/41/6/063101.

- [36] Planck Collaboration, "Planck 2015 results. XIII. Cosmological parameters," *Astron. & Astrophys., Forthcoming Article*, 2016. DOI: <http://dx.doi.org/10.1051/0004-6361/201525830>.
- [37] Planck Collaboration, "Planck 2015 results. XI. CMB power spectra, likelihoods, and robustness of parameters," *Astron. & Astrophys., Forthcoming Article*, 2016. DOI: <http://dx.doi.org/10.1051/0004-6361/201526926>.
- [38] J. S. Bagla and T Padmanabhan, "Cosmological N-body simulations," *Pramana*, **49** 2, 161–192, 1997. DOI: 10.1007/BF02845853.
- [39] J. Schaye, R. a. Crain, R. G. Bower, *et al.*, "The EAGLE project: simulating the evolution and assembly of galaxies and their environments," *Mon. Not. R Astron. Soc.*, **446**, 521–554, 2014. DOI: 10.1093/mnras/stu2058.
- [40] V. Springel, S. D. M. White, A. Jenkins, *et al.*, "Simulations of the formation, evolution and clustering of galaxies and quasars.," *Nature*, **435**, 629–636, 2005. DOI: 10.1038/nature03597.
- [41] M. Vogelsberger, S. Genel, V. Springel, *et al.*, "Introducing the Illustris Project: simulating the coevolution of dark and visible matter in the Universe," *Mon. Not. R Astron. Soc.*, **444**, 1518–1547, 2014. DOI: 10.1093/mnras/stu1536.
- [42] A. G. Sánchez, M. Crocce, A. Cabré, *et al.*, "Cosmological parameter constraints from SDSS luminous red galaxies: a new treatment of large-scale clustering," *Mon. Not. R Astron. Soc.*, **400**, 1643–1664, 2009. DOI: 10.1111/j.1365-2966.2009.15572.x.
- [43] L. Anderson, E. Aubourg, S. Bailey, *et al.*, "The clustering of galaxies in the SDSS-III Baryon Oscillation Spectroscopic Survey: baryon acoustic oscillations in the Data Releases 10 and 11 Galaxy samples," *Mon. Not. R Astron. Soc.*, **441**, 24–62, 2014. DOI: 10.1093/mnras/stu523.
- [44] T Lasserre, C Afonse, J. Albert, *et al.*, "Not enough stellar mass Machos in the Galactic halo," *Astron. & Astrophys.*, **355**, L39–L42, 2000.
- [45] C. Alcock, R. A. Allsman, D. R. Alves, *et al.*, "The MACHO project: microlensing results from 5.7 years of large magellanic cloud observations," *Astrophys. J.*, **542**, 281–307, 2000. DOI: <http://dx.doi.org/10.1086/309512>.
- [46] C. Alcock, R. A. Allsman, D. R. Alves, *et al.*, "MACHO Project Limits on Black Hole Dark Matter in the 1-30 M_{\odot} Range," *Astrophys. J.*, **550**, L169–L172, 2001. DOI: <http://dx.doi.org/10.1086/319636>.

- [47] R. H. Cyburt, "Primordial nucleosynthesis for the new cosmology: Determining uncertainties and examining concordance," *Phys. Rev. D*, **70** 2, 023505, 2004. DOI: 10.1103/PhysRevD.70.023505.
- [48] J. R. Bond, G. Efstathiou, and J. Silk, "Massive neutrinos and the large-scale structure of the universe," *Phys. Rev. Lett.*, **45** 24, 1980–1984, 1980. DOI: <http://dx.doi.org/10.1103/PhysRevLett.45.1980>.
- [49] J. R. Bond and A. S. Szalay, "The collisionless damping of density fluctuations in an expanding universe," *Astrophys. J.*, **274**, 443–468, 1983. DOI: <http://dx.doi.org/10.1086/161460>.
- [50] K. Abazajian, "Linear cosmological structure limits on warm dark matter," *Phys. Rev. D*, **73**, 063513, 2006. DOI: 10.1103/PhysRevD.73.063513.
- [51] G. Jungman, M. Kamionkowski, and K. Griest, "Supersymmetric dark matter," *Physics Reports*, **267**, 195–373, 1996. DOI: [http://dx.doi.org/10.1016/0370-1573\(95\)00058-5](http://dx.doi.org/10.1016/0370-1573(95)00058-5).
- [52] J. L. Feng, "Naturalness and the Status of Supersymmetry," *Annual Review of Nuclear and Particle Science*, **63** 1, 351–382, 2013. DOI: 10.1146/annurev-nucl-102010-130447.
- [53] G. Aad *et al.*, "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC," *Phys. Lett. B*, **716** 1, 1–29, 2012. DOI: <http://dx.doi.org/10.1016/j.physletb.2012.08.020>.
- [54] S. Chatrchyan *et al.*, "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC," *Phys. Lett. B*, **716** 1, 30–61, 2012. DOI: <http://dx.doi.org/10.1016/j.physletb.2012.08.021>.
- [55] P. Nath, "Supersymmetry after the higgs," *Annalen der Physik*, **528** 1-2, 167–178, 2016. DOI: 10.1002/andp.201500005.
- [56] M. Schmaltz and D. Tucker-Smith, "Little Higgs Review," *Annu. Rev. Nucl. Part. Sci.*, **55** 3, 1–44, 2005. DOI: 10.1146/annurev.nucl.55.090704.151502.
- [57] M. Perelstein, "Little Higgs Models and Their Phenomenology," *Prog. Part. Nucl. Phys.*, 247–291, 1 2008. DOI: <http://dx.doi.org/10.1016/j.pnpnp.2006.04.001>.
- [58] J. Hubisz and P. Meade, "Phenomenology of the littlest Higgs model with T-parity," *Phys. Rev. D*, **71** 3, 035016, 2005. DOI: 10.1103/PhysRevD.71.035016.
- [59] D. Hooper and S. Profumo, "Dark matter and collider phenomenology of universal extra dimensions," *Phys. Rep.*, **453** 2-4, 29–115, 2007. DOI: 10.1016/j.physrep.2007.09.003.

- [60] G. Servant and T. M. P. Tait, "Elastic scattering and direct detection of Kaluza-Klein dark matter," *New J. Phys.*, **4**, 99.1–99.19, 2002. DOI: <http://dx.doi.org/10.1088/1367-2630/4/1/399>.
- [61] E. W. Kolb and M. S. Turner, "The Early Universe," *Addison-Wesley*, 1990.
- [62] S. Dodelson and L. M. Widrow, "Sterile neutrinos as dark matter," *Phys. Rev. Lett.*, **72** 1, 17–20, 1994. DOI: [10.1103/PhysRevLett.72.17](https://doi.org/10.1103/PhysRevLett.72.17).
- [63] K. Abazajian, G. M. Fuller, and M. Patel, "Sterile neutrino hot, warm, and cold dark matter," *Phys. Rev. D*, **64** 2, 023501, 2001. DOI: [10.1103/PhysRevD.64.023501](https://doi.org/10.1103/PhysRevD.64.023501).
- [64] X. Shi and G. M. Fuller, "New dark matter candidate: nonthermal sterile neutrinos," *Phys. Rev. Lett.*, **82** 14, 2832–2835, 1999. DOI: <http://dx.doi.org/10.1103/PhysRevLett.82.2832>.
- [65] M. H. Chan, "Can decaying sterile neutrinos account for all dark matter?" *Astrophys. Space Sci.*, **361**, 116, 2016. DOI: [10.1007/s10509-016-2703-6](https://doi.org/10.1007/s10509-016-2703-6).
- [66] R. Adhikari, M. Agostini, N. Anh Ky, *et al.*, "A white paper on keV sterile neutrino dark matter," *arXiv:1602.04816*, 2016.
- [67] R. D. Peccei and H. R. Quinn, "CP conservation in the presence of pseudoparticles," *Phys. Rev. Lett.*, **38** 25, 1440–1443, 1977. DOI: <http://dx.doi.org/10.1103/PhysRevLett.38.1440>.
- [68] J. E. Kim, "Light pseudoscalars, particle physics and cosmology," *Phys. Rep.*, **150** 1 and 2, 1–177, 1987. DOI: [http://dx.doi.org/10.1016/0370-1573\(87\)90017-2](http://dx.doi.org/10.1016/0370-1573(87)90017-2).
- [69] M. Dine, W. Fischler, and M. Srednicki, "A simple solution to the strong CP problem with a harmless axion," *Phys. Lett. B*, **104** 3, 199–202, 1981. DOI: [http://dx.doi.org/10.1016/0370-2693\(81\)90590-6](http://dx.doi.org/10.1016/0370-2693(81)90590-6).
- [70] J. E. Kim, "Weak-interaction singlet and strong CP invariance," *Phys. Rev. Lett.*, **43** 2, 103–107, 1979. DOI: <http://dx.doi.org/10.1103/PhysRevLett.43.103>.
- [71] L. Visinelli and P. Gondolo, "Dark matter axions revisited," *Phys. Rev. D*, **80**, 035024, 2009. DOI: [10.1103/PhysRevD.80.035024](https://doi.org/10.1103/PhysRevD.80.035024).
- [72] A. Derevianko, V. A. Dzuba, V. V. Flambaum, and M. Pospelov, "Axio-electric effect," *Phys. Rev. D*, **82** 6, 065006, 2010. DOI: [10.1103/PhysRevD.82.065006](https://doi.org/10.1103/PhysRevD.82.065006).
- [73] G. Raffelt and L. Stodolsky, "Mixing of the photon with low-mass particles," *Phys. Rev. D*, **37** 5, 1237–1249, 1988. DOI: <http://dx.doi.org/10.1103/PhysRevD.37.1237>.

- [74] L. J. Rosenberg, "Dark-matter QCD-axion searches.," *PNAS*, **112** 40, 12278–12281, 2015. DOI: 10.1073/pnas.1308788112.
- [75] V. A. Mitsou, "Overview of searches for dark matter at the LHC," *J. of Phys. Conf. Ser.*, **651**, 012023, 2015. DOI: 10.1088/1742-6596/651/1/012023.
- [76] J. Goodman, M. Ibe, A. Rajaraman, *et al.*, "Constraints on dark matter from colliders," *Phys. Rev. D*, **82** 11, 116010, 2010. DOI: 10.1103/PhysRevD.82.116010.
- [77] S. A. Malik *et al.*, "Interplay and characterization of dark matter searches at colliders and in direct detection experiments," *Phys. Dark Universe*, **9–10**, 51–58, 2015. DOI: <http://dx.doi.org/10.1016/j.dark.2015.03.003>.
- [78] S. Y. Hoh, J. R. Komaragiri, and W. a. T. Wan Abdullah, "Dark Matter searches at the Large Hadron Collider," *AIP Conf. Proc.*, **1704**, 020005, 2016. DOI: 10.1063/1.4940063.
- [79] G Aad *et al.*, "Search for new phenomena in final states with an energetic jet and large missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector.," *Eur. Phys. J. C*, **75**, 299, 2015. DOI: 10.1140/epjc/s10052-015-3517-3.
- [80] CMS Collaboration, "Search for dark matter, extra dimensions, and unparticles in monojet events in proton-proton collisions at $\sqrt{s} = 8$ TeV," *Eur. Phys. J. C*, **75** 5, 235, 2015. DOI: 10.1140/epjc/s10052-015-3451-4.
- [81] G Aad *et al.*, "Search for dark matter in events with missing transverse momentum and a Higgs boson decaying to two photons in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector," *Phys. Rev. Lett.*, **115**, 131801, 2015. DOI: 10.1103/PhysRevLett.115.131801.
- [82] CMS Collaboration, "Search for new phenomena in monophoton final states in protonproton collisions at $\sqrt{s} = 8$ TeV," *Phys. Lett. B*, **755**, 102–124, 2016. DOI: 10.1016/j.physletb.2016.01.057.
- [83] G. Aad *et al.*, "Search for invisible decays of a Higgs boson using vector-boson fusion in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector," *JHEP*, **01**, 172, 2016. DOI: 10.1007/JHEP01(2016)172.
- [84] CMS Collaboration, "Searches for invisible Higgs boson decays with the CMS detector," *CMS PAS*, **16**, 016, 2016.
- [85] M. Klasen, M. Pohl, and G. Sigl, "Indirect and direct search for dark matter," *Prog. Part. Nucl. Phys.*, **85**, 1–32, 2015. DOI: 10.1016/j.ppnp.2015.07.001.

- [86] T. Bringmann and C. Weniger, "Gamma ray signals from dark matter: Concepts, status and prospects," *Phys. Dark Universe*, **1**, 194–217, 2012. DOI: 10.1016/j.dark.2012.10.005.
- [87] C. Rott, "Review of Indirect WIMP Search Experiments," *Nucl. Phys. B: Proc. Suppl.*, **235-236**, 413–420, 2013. DOI: 10.1016/j.nuclphysbps.2013.04.040.
- [88] M. Cirelli, "Dark Matter indirect detection: Some anomalies and many constraints," *EPJ Web of Conferences*, **121**, 06002, 2016.
- [89] O. Adriani, G. C. Barbarino, G. A. Bazilevskaya, *et al.*, "An anomalous positron abundance in cosmic rays with energies 1.5–100 GeV," *Nature*, **458**, 607–609, 2009. DOI: 10.1038/nature07942.
- [90] M. Ackermann, A. M. A. A., *et al.*, "Measurement of separate cosmic-ray electron and positron spectra with the Fermi Large Area Telescope," *Phys. Rev. Lett.*, **108** 1, 011103, 2012. DOI: 10.1103/PhysRevLett.108.011103.
- [91] M. Aguilar, G. Alberti, B. Alpat, *et al.*, "First result from the alpha magnetic spectrometer on the international space station: precision measurement of the positron fraction in primary cosmic rays of 0.5–350 GeV," *Phys. Rev. Lett.*, **110**, 141102, 14 2013. DOI: 10.1103/PhysRevLett.110.141102.
- [92] A. A. Abdo, M. Ackermann, M. Ajello, *et al.*, "Measurement of the cosmic ray $e^+ + e^-$ spectrum from 20 GeV to 1 TeV with the Fermi Large Area Telescope," *Phys. Rev. Lett.*, **102**, 181101, 2009. DOI: 10.1103/PhysRevLett.102.181101.
- [93] M. Ackermann, M. Ajello, W. B. Atwood, *et al.*, "Fermi LAT observations of cosmic-ray electrons from 7 GeV to 1 TeV," *Phys. Rev. D*, **82**, 092004, 2010. DOI: 10.1103/PhysRevD.82.092004.
- [94] F. Aharonian, A. G. Akhperjanian, U. Barres de Almeida, *et al.*, "Energy spectrum of cosmic-ray electrons at TeV energies," *Phys. Rev. Lett.*, **101**, 261104, 2008. DOI: 10.1103/PhysRevLett.101.261104.
- [95] F. Aharonian, A. G. Akhperjanian, G. Anton, *et al.*, "Probing the ATIC peak in the cosmic-ray electron spectrum with HESS," *Astron. & Astrophys.*, **508**, 561–564, 2009. DOI: <http://dx.doi.org/10.1051/0004-6361/200913323>.
- [96] O. Adriani *et al.*, "PAMELA results on the cosmic-ray antiproton flux from 60 MeV to 180 GeV in kinetic energy," *Phys. Rev. Lett.*, **105** 12, 121101, 2010. DOI: 10.1103/PhysRevLett.105.121101.

- [97] M. Ackermann, M. Ajello, A. Albert, *et al.*, “Constraining dark matter models from a combined analysis of Milky Way satellites with the Fermi Large Area Telescope,” *Phys. Rev. Lett.*, **107**, 241302, 2011. DOI: 10.1103/PhysRevLett.107.241302.
- [98] M. Papucci and A. Strumia, “Robust implications on dark matter from the first FERMI sky γ map,” *J. Cosmol. Astropart. Phys.*, **03**, 014, 2010. DOI: 10.1088/1475-7516/2010/03/014.
- [99] M. Cirelli, P. Panci, and P. D. Serpico, “Diffuse gamma ray constraints on annihilating or decaying Dark Matter after Fermi,” *Nucl. Phys. B*, **840**, 284–303, 2010. DOI: 10.1016/j.nuclphysb.2010.07.010.
- [100] S. Galli, F. Iocco, G. Bertone, and A. Melchiorri, “CMB constraints on dark matter models with large annihilation cross section,” *Phys. Rev. D*, **80**, 023505, 2009. DOI: 10.1103/PhysRevD.80.023505.
- [101] M. Di Mauro, F. Donato, N. Fornengo, *et al.*, “Interpretation of AMS-02 electrons and positrons data,” *J. Cosmol. Astropart. Phys.*, **04**, 006, 2014. DOI: 10.1088/1475-7516/2014/04/006.
- [102] K. Kohri, K. Ioka, Y. Fujita, and R. Yamazaki, “Can we explain AMS-02 antiproton and positron excesses simultaneously by nearby supernovae without pulsars or dark matter?” *Prog. Theor. Exp. Phys.*, **2016** 2, 021E01, 2016. DOI: 10.1093/ptep/ptv193.
- [103] T. Bringmann, X. Huang, A. Ibarra, *et al.*, “Fermi LAT search for internal bremsstrahlung signatures from dark matter annihilation,” *J. Cosmol. Astropart. Phys.*, **07**, 054, 2012. DOI: 10.1088/1475-7516/2012/07/054.
- [104] C. Weniger, “A tentative gamma-ray line from Dark Matter annihilation at the Fermi Large Area Telescope,” *J. Cosmol. Astropart. Phys.*, **08**, 007, 2012. DOI: 10.1088/1475-7516/2012/08/007.
- [105] A. Hektor, M. Raidal, and E. Tempel, “Evidence for Indirect Detection of Dark Matter From Galaxy Clusters in Fermi Γ -Ray Data,” *Astrophys. J.*, **762**, L22, 2013. DOI: 10.1088/2041-8205/762/2/L22.
- [106] A. Boyarsky, D. Malyshev, and O. Ruchayskiy, “Spectral and spatial variations of the diffuse γ -ray background in the vicinity of the Galactic plane and possible nature of the feature at 130 GeV,” *Phys. Dark Univ.*, **2** 2, 90–96, 2013.
- [107] N. Mirabal, “The dark knight falters,” *Mon. Not. R Astron. Soc.*, **429**, L109–L113, 2012. DOI: 10.1093/mnrasl/sls034.

- [108] A. Hektori, M. Raidal, and E. Tempel, "Double gamma-ray lines from unassociated Fermi-LAT sources revisited," *arXiv:1208.1996*, 2012.
- [109] L. Goodenough and D. Hooper, "Possible evidence for dark matter annihilation in the inner Milky Way from the Fermi Gamma Ray Space Telescope," *arXiv:0910.2998*, 2009.
- [110] T. Daylan, D. P. Finkbeiner, D. Hooper, *et al.*, "The characterization of the gamma-ray signal from the central Milky Way: A case for annihilating dark matter," *Phys. Dark Univ.*, **12**, 1–23, 2016. DOI: 10.1016/j.dark.2015.12.005.
- [111] K. N. Abazajian, "The consistency of Fermi-LAT observations of the galactic center with a millisecond pulsar population in the central stellar cluster," *J. Cosmol. Astropart. Phys.*, **03**, 2011. DOI: 10.1088/1475-7516/2011/03/010.
- [112] Q. Yuan and B. Zhang, "Millisecond pulsar interpretation of the Galactic center gamma-ray excess," *J. High Energy Astrophys.*, **3-4**, 1–8, 2014. DOI: 10.1016/j.jheap.2014.06.001.
- [113] A. Boyarsky, D. Malyshev, and O. Ruchayskiy, "A comment on the emission from the Galactic Center as seen by the Fermi telescope," *Phys. Lett. B*, **705**, 165–169, 2011. DOI: 10.1016/j.physletb.2011.10.014.
- [114] E. Carlson and S. Profumo, "Cosmic ray protons in the inner Galaxy and the Galactic Center gamma-ray excess," *Phys. Rev. D*, **90**, 023015, 2014. DOI: 10.1103/PhysRevD.90.023015.
- [115] J. Petrović, P. D. Serpico, and G. Zaharijaš, "Galactic Center gamma-ray "excess" from an active past of the Galactic Centre?" *J. Cosmol. Astropart. Phys.*, **10**, 052, 2014. DOI: 10.1088/1475-7516/2014/10/052.
- [116] R. Bartels, S. Krishnamurthy, and C. Weniger, "Strong Support for the Millisecond Pulsar Origin of the Galactic Center GeV Excess," *Phys. Rev. Lett.*, **116**, 051102, 2016. DOI: 10.1103/PhysRevLett.116.051102.
- [117] S. K. Lee, M. Lisanti, B. R. Safdi, *et al.*, "Evidence for Unresolved γ -Ray Point Sources in the Inner Galaxy," *Phys. Rev. Lett.*, **116**, 051103, 2016. DOI: 10.1103/PhysRevLett.116.051103.
- [118] A. Boyarsky, O. Ruchayskiy, D. Iakubovskiy, and J. Franse, "Unidentified line in x-ray spectra of the Andromeda galaxy and Perseus galaxy cluster," *Phys. Rev. Lett.*, **113**, 251301, 2014. DOI: 10.1103/PhysRevLett.113.251301.
- [119] E. Bulbul, M. Markevitch, A. Foster, *et al.*, "Detection of an Unidentified Emission Line in the Stacked X-Ray Spectrum of Galaxy Clusters," *Astrophys. J.*, **789**, 13, 2014. DOI: 10.1088/0004-637X/789/1/13.

- [120] T. Jeltema and S. Profumo, "Discovery of a 3.5 keV line in the Galactic Centre and a critical look at the origin of the line across astronomical targets," *Mon. Not. R Astron. Soc.*, **450**, 2143–2152, 2015. DOI: 10.1093/mnras/stv768.
- [121] M. Aartsen, K. Abraham, M. Ackermann, *et al.*, "Improved limits on dark matter annihilation in the sun with the 79-string icecube detector and implications for supersymmetry," *J. Cosmol. Astropart. Phys.*, **2016** 04, 022, 2016. DOI: <http://dx.doi.org/10.1088/1475-7516/2016/04/022>.
- [122] J. D. Lewin and R. F. Smith, "Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil," *Astropart. Phys.*, **6** 96, 87–112, 1996. DOI: 10.1016/S0927-6505(96)00047-3.
- [123] A. Drukier, K. Freese, and D. Spergel, "Detecting cold dark-matter candidates," *Phys. Rev. D*, **33** 12, 1986. DOI: 10.1103/PhysRevD.33.3495.
- [124] D.-M. Mei and A. Hime, "Muon-induced background study for underground laboratories," *Phys. Rev. D*, **73** 5, 053004, 2006. DOI: 10.1103/PhysRevD.73.053004.
- [125] D.-M. Mei, C. Zhang, and A. Hime, "Evaluation of (α , n) induced neutrons as a background for dark matter experiments," *Nucl. Instrum. Meth. A*, **606**, 651–660, 2009. DOI: 10.1016/j.nima.2009.04.032.
- [126] D. S. Akerib *et al.*, "Radiogenic and Muon-Induced Backgrounds in the LUX Dark Matter Detector," *Astropart. Phys.*, **62**, 33–46, 2015. DOI: 10.1016/j.astropartphys.2014.07.009.
- [127] A. Gütlein, C. Ciemniak, F. von Feilitzsch, *et al.*, "Solar and atmospheric neutrinos: Background sources for the direct dark matter searches," *Astropart. Phys.*, **34**, 90–96, 2010. DOI: 10.1016/j.astropartphys.2010.06.002.
- [128] J. Billard, L. E. Strigari, and E. Figueroa-Feliciano, "Solar neutrino physics with low-threshold dark matter detectors," *Phys. Rev. D*, **91**, 095023, 2015. DOI: 10.1103/PhysRevD.91.095023.
- [129] F. Ruppin, J. Billard, E. Figueroa-Feliciano, and L. Strigari, "Complementarity of dark matter detectors in light of the neutrino background," *Phys. Rev. D*, **90**, 083510, 8 2014. DOI: 10.1103/PhysRevD.90.083510.
- [130] LZ collaboration, "LZ Technical Design Report," in preparation.
- [131] T Marrodán Undagoitia and L Rauch, "Dark matter direct-detection experiments," *J. Phys. G: Nuc. Phys.*, **43**, 013001, 2016. DOI: <http://dx.doi.org/10.1088/0954-3899/43/1/013001>.

- [132] R. Bernabei, P. Belli, F. Cappella, *et al.*, “First results from DAMA/LIBRA and the combined results with DAMA/NaI,” *Eur. Phys. J. C*, **56** 3, 333–355, 2008. DOI: 10.1140/epjc/s10052-008-0662-y.
- [133] R. Bernabei, P. Belli, F. Cappella, *et al.*, “Final model independent result of DAMA/LIBRA phase 1,” *Eur. Phys. J. C*, **73** 12, 2648, 2013. DOI: 10.1140/epjc/s10052-013-2648-7.
- [134] C. Savage, G. Gelmini, P. Gondolo, and K. Freese, “Compatibility of DAMA/LIBRA dark matter detection with other searches,” *J. Cosmol. Astropart. Phys.*, **04**, 010, 2009. DOI: 10.1088/1475-7516/2009/04/010.
- [135] 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** 18, 181301, 2012. DOI: 10.1103/PhysRevLett.108.181301.
- [136] E. Shields, J. Xu, and F. Calaprice, “SABRE: A New NaI(Tl) Dark Matter Direct Detection Experiment,” *Phys. Procedia*, **61**, 169–178, 2015. DOI: 10.1016/j.phpro.2014.12.028.
- [137] J. Amaré, S. Cebrián, C. Cuesta, *et al.*, “From ANAIS-25 towards ANAIS-250,” *Phys. Procedia*, **61**, 157–162, 2015. DOI: 10.1016/j.phpro.2014.12.026.
- [138] J. Cherwinka, R. Co, D. Cowen, *et al.*, “A search for the dark matter annual modulation in South Pole ice,” *Astroparticle Physics*, **35** 11, 749–754, 2012. DOI: 10.1016/j.astropartphys.2012.03.003.
- [139] R. Agnese, a. J. Anderson, D. Balakishiyeva, *et al.*, “Demonstration of surface electron rejection with interleaved germanium detectors for dark matter searches,” *App. Phys. Lett.*, **103**, 164105, 2013. DOI: 10.1063/1.4826093.
- [140] P. Brink, B. Cabrera, J. Castle, *et al.*, “First test runs of a dark-matter detector with interleaved ionization electrodes and phonon sensors for surface-event rejection,” *Nucl. Instrum. Meth. A*, **559**, 414–416, 2006. DOI: 10.1016/j.nima.2005.12.026.
- [141] M. Pyle, B. Serfass, P. L. Brink, *et al.*, “Surface electron rejection from ge detector with interleaved charge and phonon channels,” *AIP Conf. Proc.*, **1185** 1, 223–226, 2009. DOI: <http://dx.doi.org/10.1063/1.3292319>.
- [142] R. Agnese, A. J. Anderson, T. Aramaki, *et al.*, “New Results from the Search for Low-Mass Weakly Interacting Massive Particles with the CDMS Low Ionization Threshold Experiment.,” *Physical review letters*, **116**, 071301, 2016. DOI: 10.1103/PhysRevLett.116.071301.

- [143] C. Amole *et al.*, “Improved Dark Matter Search Results from PICO-2L Run-2,” *Phys. Rev. D*, **93**, 061101, 2016. DOI: 10.1103/PhysRevD.93.061101.
- [144] C. Amole *et al.*, “Dark Matter Search Results from the PICO-60 CF₃I Bubble Chamber,” *Phys. Rev. D*, **93**, 052014, 2015. DOI: 10.1103/PhysRevD.93.052014.
- [145] M. Kuniak, P.-A. Amaudruz, M. Batygov, *et al.*, “Deap-3600 dark matter search,” *Nucl. Part. Phys. Proc.*, **273–275**, 340–346, 2016. DOI: <http://dx.doi.org/10.1016/j.nuclphysbps.2015.09.048>.
- [146] M. Boulay and A. Hime, “Technique for direct detection of weakly interacting massive particles using scintillation time discrimination in liquid argon,” *Astropart. Phys.*, **25**, 179–182, 2006. DOI: 10.1016/j.astropartphys.2005.12.009.
- [147] E. Aprile, J. Aalbers, F. Agostini, *et al.*, “Physics reach of the XENON1T dark matter experiment,” *J. Cosmol. Astropart. Phys.*, **04**, 027, 2016. DOI: <http://dx.doi.org/10.1088/1475-7516/2016/04/027>.
- [148] D. S. Akerib *et al.*, “Improved Limits on Scattering of Weakly Interacting Massive Particles from Reanalysis of 2013 LUX Data,” *Phys. Rev. Lett.*, **116**, 161301, 2016. DOI: <http://dx.doi.org/10.1103/PhysRevLett.116.161301>.
- [149] P. Agnes, T. Alexander, A. Alton, *et al.*, “First results from the DarkSide-50 dark matter experiment at Laboratori Nazionali del Gran Sasso,” *Phys. Lett. B*, **743**, 456–466, 2015. DOI: 10.1016/j.physletb.2015.03.012.
- [150] X. Xiao, X. Chen, A. Tan, *et al.*, “Low-mass dark matter search results from full exposure of the PandaX-I experiment,” *Phys. Rev. D*, **92**, 052004, 2015. DOI: 10.1103/PhysRevD.92.052004.
- [151] E. Aprile, M. Alfonsi, K. Arisaka, *et al.*, “Dark Matter Results from 225 Live Days of XENON100 Data,” *Phys. Rev. Lett.*, **109** 18, 181301, 2012. DOI: 10.1103/PhysRevLett.109.181301.
- [152] E. A. Bagnaschi, O. Buchmueller, R. Cavanaugh, *et al.*, “Supersymmetric dark matter after LHC run 1,” *Eur. Phys. J. C*, **75** 10, 500, 2015. DOI: 10.1140/epjc/s10052-015-3718-9.
- [153] M. Pato, L. Baudis, G. Bertone, *et al.*, “Complementarity of dark matter direct detection targets,” *Phys. Rev. D*, **83**, 083505, 2011. DOI: 10.1103/PhysRevD.83.083505.
- [154] C. E. Dahl, “The physics of background discrimination in liquid xenon, and first results from XENON 10 in the hunt for WIMP dark matter,” PhD thesis, Princeton University, 2009.

- [155] D. S. Akerib *et al.*, "The Large Underground Xenon (LUX) experiment," *Nucl. Instrum. Meth. A*, **704**, 111–126, 2013. DOI: 10.1016/j.nima.2012.11.135.
- [156] T. Doke and K. Masuda, "Present status of liquid rare gas scintillation detectors and their new application to gamma-ray calorimeters," *Nucl. Instrum. Meth.*, **420**, 62–80, 1999. DOI: doi:10.1016/S0168-9002(98)00933-4.
- [157] J. Jortner, L. Meyer, S. A. Rice, and E. G. Wilson, "Localized Excitations in Condensed Ne, Ar, Kr, and Xe," *J. Chem. Phys.*, **42** 12, 4250, 1965. DOI: 10.1063/1.1695927.
- [158] S. Kubota, M. Hishida, and J. Raun, "Evidence for a triplet state of the self-trapped exciton states in liquid argon, krypton and xenon," *J. Phys. C: Solid State Phys.*, **2645**, 1978. DOI: <http://dx.doi.org/10.1088/0022-3719/11/12/024>.
- [159] A. Hitachi, T. Takahashi, and N. Funayama, "Effect of ionization density on the time dependence of luminescence from liquid argon and xenon," *Phys. Rev. B*, **27** 9, 5279–5285, 1983. DOI: 10.1103/PhysRevB.27.5279.
- [160] A. Baldini, C. Bemporad, F. Cei, *et al.*, "Absorption of scintillation light in a 100 l liquid xenon γ -ray detector and expected detector performance," *Nucl. Instrum. Meth. A*, **545** 3, 753–764, 2005. DOI: <http://dx.doi.org/10.1016/j.nima.2005.02.029>.
- [161] L. S. Miller, S. Howe, and W. E. Spear, "Charge transport in solid and liquid Ar, Kr, and Xe," *Physical Review*, **166**, 871–878, 3 1968. DOI: 10.1103/PhysRev.166.871.
- [162] G. Bakale, U. Sowada, and W. F. Schmidt, "Effect of an Electric Field on Electron Attachment to SF₆, N₂O, and O₂ in Liquid Argon and Xenon," **80** 23, 2556–2559, 1976. DOI: 10.1021/j100564a006.
- [163] W. T. Chen, H. Carduner, J. P. Cussonneau, *et al.*, "Measurement of the Transverse Diffusion Coefficient of Charge in Liquid Xenon," *Defect and Diffusion Forum*, 567–572, 2012. DOI: 10.4028/www.scientific.net/DDF.326-328.567.
- [164] P. Sorensen, "Anisotropic diffusion of electrons in liquid xenon with application to improving the sensitivity of direct dark matter searches," *Nucl. Instrum. Meth. A*, **635** 1, 41–43, 2011. DOI: 10.1016/j.nima.2011.01.089.
- [165] A. I. Bolozdynya, "Two-phase emission detectors and their applications," *Nucl. Instrum. Meth. A*, **422**, 314–320, 1999. DOI: 10.1016/S0168-9002(98)00965-6.

- [166] T. Takahashi, S. Himi, and M. Suzuki, "Emission spectra from Ar–Xe, Ar–Kr, Ar–N₂, Ar–CH₄, Ar–CO₂ and Xe–N₂ gas scintillation proportional counters," *Nuclear Instruments and ...*, **205**, 591–596, 3 1983. DOI: [http://dx.doi.org/10.1016/0167-5087\(83\)90028-5](http://dx.doi.org/10.1016/0167-5087(83)90028-5).
- [167] A. C. Fonseca, R. Meleiro, V. Chepel, *et al.*, "Study of secondary scintillation in xenon vapour," *IEEE Nuc. Sci. Symp. Conf.*, **1**, 572–576, 2004. DOI: 10.1109/NSSMIC.2004.1462260.
- [168] C. M. B. Monteiro, L. M. P. Fernandes, J. A. M. Lopes, *et al.*, "Secondary scintillation yield in pure xenon," *JINST*, **05001**, 2007. DOI: <http://dx.doi.org/10.1088/1748-0221/2/05/P05001>.
- [169] F. P. Santos, T. H. V. T. Dias, A. D. Stauffer, and C. A. Conde, "Three-dimensional Monte Carlo calculation of the VUV electroluminescence and other electron transport parameters in xenon," *J. Phys. D: Appl. Phys.*, **27** 1, 42, 1994. DOI: <http://dx.doi.org/10.1088/0022-3727/27/1/007>.
- [170] V. Chepel and H. Araújo, "Liquid noble gas detectors for low energy particle physics," *JINST*, **8**, R04001, 2013. DOI: 10.1088/1748-0221/8/04/R04001.
- [171] T. Hashimoto and Y. Nakamura, *Papers of Gas Discharge Technical Committee ED-90.61 (Japan: IEE)*, 1990.
- [172] T. Doke, A. Hitachi, J. Kikuchi, *et al.*, "Absolute Scintillation Yields in Liquid Argon and Xenon for Various Particles," *Japanese Journal of Applied Physics*, **41** Part 1, No. 3A, 1538–1545, 2002. DOI: 10.1143/JJAP.41.1538.
- [173] E. Aprile, K. L. Giboni, P. Majewski, *et al.*, "Observation of anticorrelation between scintillation and ionization for MeV gamma rays in liquid xenon," *Phys. Rev. B*, **76** 1, 014115, 2007. DOI: 10.1103/PhysRevB.76.014115.
- [174] Q. Lin, J. Fei, F. Gao, *et al.*, "Scintillation and ionization responses of liquid xenon to low energy electronic and nuclear recoils at drift fields from 236 V/cm to 3.93 kV/cm," *Phys. Rev. D*, **92** 3, 032005, 2015. DOI: 10.1103/PhysRevD.92.032005.
- [175] 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** 6, 063501, 2011. DOI: 10.1103/PhysRevD.83.063501.
- [176] S. Kubota, A. Nakamoto, and T. Takahashi, "Recombination luminescence in liquid argon and in liquid xenon," *Phys. Rev. B*, **17**, 6–9, 1978. DOI: <http://dx.doi.org/10.1103/PhysRevB.17.2762>.

- [177] V. N. Lebedenko, H. M. Araújo, E. J. Barnes, *et al.*, “Results from the first science run of the ZEPLIN-III dark matter search experiment,” *Phys. Rev. D*, **80**, 052010, 2009. DOI: 10.1103/PhysRevD.80.052010.
- [178] D. S. Akerib *et al.*, “Tritium calibration of the LUX dark matter experiment,” *Phys. Rev. D*, **93**, 072009, 2016. DOI: <http://dx.doi.org/10.1103/PhysRevD.93.072009>.
- [179] M Szydagis *et al.*, “NEST: a comprehensive model for scintillation yield in liquid xenon,” *Journal of Instrumentation*, **6** 10, P10002, 2011. DOI: 10.1088/1748-0221/6/10/P10002.
- [180] K. T. Lesko, S. Acheson, J. Alonso, P. Bauer, *et al.*, “Deep Underground Science and Engineering Laboratory-Preliminary Design Report,” *arXiv:1108.0959*, 1, 2011.
- [181] K. T. Lesko, “The Sanford Underground Research Facility at Homestake,” *The European Physical Journal Plus*, **127** 9, 1–11, 2012. DOI: 10.1140/epjp/i2012-12107-x.
- [182] C. H. Faham, “Prototype , Surface Commissioning and Photomultiplier Tube Characterization for the Large Underground Xenon (LUX) Direct Dark Matter Search Experiment,” PhD thesis, Brown University.
- [183] D. S. Akerib *et al.*, “The LUX prototype detector: Heat exchanger development,” *Nucl. Instrum. Meth. A*, **709**, 29–36, 2013. DOI: 10.1016/j.nima.2013.01.036.
- [184] D. S. Akerib *et al.*, “Technical results from the surface run of the LUX dark matter experiment,” *Astropart. Phys.*, **45**, 34–43, 2013. DOI: <http://dx.doi.org/10.1016/j.astropartphys.2013.02.001>.
- [185] D. S. Akerib *et al.*, “First results from the LUX dark matter experiment at the Sanford Underground Research Facility,” *Phys. Rev. Lett.*, **112**, 091303, 2014. DOI: 10.1103/PhysRevLett.112.091303.
- [186] D. S. Akerib *et al.*, “Radio-assay of Titanium samples for the LUX Experiment,” *arXiv:1112.1376*, 2011.
- [187] D. S. Akerib *et al.*, “An ultra-low background PMT for liquid xenon detectors,” *Nucl. Instrum. Meth. A*, **703**, 1–6, 2013. DOI: 10.1016/j.nima.2012.11.020.
- [188] M. Attisha, L. De Viveiros, R. Gaitksell, and J. Thompson, “Soudan Low Background Counting Facility (SOLO),” *AIP Conf. Proc.*, **785** 1, 75–78, 2005. DOI: <http://dx.doi.org/10.1063/1.2060455>.

- [189] A. Dobi, D. S. Leonard, C. Hall, *et al.*, "Study of a zirconium getter for purification of xenon gas," *Nucl. Instrum. Meth. A*, **620** 2-3, 594–598, 2010. DOI: 10.1016/j.nima.2010.03.151.
- [190] D. S. Leonard, A. Dobi, C. Hall, *et al.*, "A simple high-sensitivity technique for purity analysis of xenon gas," *Nucl. Instrum. Meth. A*, **621** 1-3, 678–684, 2010. DOI: 10.1016/j.nima.2010.04.152.
- [191] A. Dobi, C. Davis, C. Hall, *et al.*, "Detection of krypton in xenon for dark matter applications," *Nucl. Instrum. Meth. A*, **665**, 1–6, 2011. DOI: 10.1016/j.nima.2011.11.043.
- [192] A. Bolozdynya, A. Bradley, S. Bryan, *et al.*, "Cryogenics for the LUX Detector," *IEEE Trans. Nuc. Sci.*, **56** 4, 2309–2312, 2009. DOI: 10.1109/TNS.2009.2023443.
- [193] D. S. Akerib *et al.*, "Data acquisition and readout system for the LUX dark matter experiment," *Nucl. Instrum. Meth. A*, **668**, 1–8, 2012. DOI: <http://dx.doi.org/10.1016/j.nima.2011.11.063>.
- [194] D. S. Akerib *et al.*, "FPGA-based Trigger System for the LUX Dark Matter Experiment," *Nucl. Instrum. Meth. A*, **818**, 57–67, 2016. DOI: [doi:10.1016/j.nima.2016.02.017](https://doi.org/10.1016/j.nima.2016.02.017). eprint: 1511.03541.
- [195] L. W. Kastens, S. B. Cahn, A. Manzur, and D. N. McKinsey, "Calibration of a liquid xenon detector with ^{83}Kr ," *Phys. Rev. C*, **80** 4, 045809, 2009. DOI: 10.1103/PhysRevC.80.045809.
- [196] A. Manalaysay, T. Marrodán Undagoitia, A. Askin, *et al.*, "Spatially uniform calibration of a liquid xenon detector at low energies using $(^{83}\text{m})\text{Kr}$," *The Review of Scientific Instruments*, **81** 7, 073303, 2010. DOI: 10.1063/1.3436636.
- [197] LZ Collaboration, "LUX-ZEPLIN (LZ) Conceptual Design Report," 2015. arXiv: 1509.02910.
- [198] M. B. Chadwick *et al.*, "Nuclear data sheets," *special issue on ENDF/B-VII.1 library*, **112**, 2887, 2011.
- [199] M. J. Berger *et al.*, "Xcom: photon cross section database (version 1.5)," *National Institute of Standards and Technology, 100 Bureau Drive, Stop 1070, Gaithersburg, MD 20899-1070.*, 2010.
- [200] A. I. Bolozdynya, P. P. Brusov, T. Shutt, *et al.*, "A chromatographic system for removal of radioactive ^{85}Kr from xenon," *Nucl. Instrum. Meth. A*, **579** 1, 50–53, 2007. DOI: 10.1016/j.nima.2007.04.011.

- [201] J. B. Albert *et al.*, “Improved measurement of the $2\nu\beta\beta$ half-life of ^{136}Xe with the EXO-200 detector,” *Phys. Rev. C*, **89** 1, 015502, 2014. DOI: 10.1103/PhysRevC.89.015502.
- [202] X. Guo *et al.*, “A Precision measurement of the neutrino mixing angle θ_{13} using reactor antineutrinos at Daya-Bay (proposal),” *arXiv:hep-ex/0701029*, 2007.
- [203] F. P. An, Q. An, J. Z. Bai, *et al.*, “A side-by-side comparison of Daya Bay antineutrino detectors,” *Nucl. Instrum. Meth. A*, **685**, 78–97, 2012. DOI: 10.1016/j.nima.2012.05.030.
- [204] D. S. Akerib *et al.*, “Chromatographic separation of radioactive noble gases from xenon,” *arXiv:1605.03844*, 2016.
- [205] V. N. Solovov, V. A. Belov, D. Y. Akimov, *et al.*, “Position reconstruction in a dual phase xenon scintillation detector,” *IEEE Trans. Nuc. Sci.*, **59** 6, 3286–3293, 2012. DOI: 10.1109/TNS.2012.2221742.
- [206] A. Chan, H. B. Crawley, D. M. Edsall, *et al.*, “Performance of the HPC calorimeter in DELPHI,” *IEEE*, **42** 4, 491–498, 1995. DOI: 10.1109/23.467923.
- [207] A. Dobi, “Measurement of the Electron Recoil Band of the LUX Dark Matter Detector with a Tritium Calibration Source,” PhD thesis, University of Maryland.
- [208] G. Drexlin, V. Hannen, S. Mertens, and C. Weinheimer, “Current direct neutrino mass experiments,” *Advan. High Energy Phys.*, **2013**, 2013. DOI: <http://dx.doi.org/10.1155/2013/293986>.
- [209] S. Agostinelli *et al.*, “Geant4 – a simulation toolkit,” *Nucl. Instrum. Meth. A*, **506** 3, 250–303, 2003. DOI: [http://dx.doi.org/10.1016/S0168-9002\(03\)01368-8](http://dx.doi.org/10.1016/S0168-9002(03)01368-8).
- [210] D. S. Akerib *et al.*, “LUXSim: A component-centric approach to low-background simulations,” *Nucl. Instrum. Meth. A*, **675**, 63–77, 2012. DOI: 10.1016/j.nima.2012.02.010.
- [211] P. Collon, T. Antaya, and B. Davids, “Measurement of ^{81}Kr in the atmosphere,” *Nucl. Instrum. Meth. B*, **123**, 122–127, 1–4 1997. DOI: doi : 10.1016/S0168-583X(96)00674-X.
- [212] L. Wilhelmová, M. Tomávek, and K. Stukheil, “Monitoring of krypton-85 activity in the atmosphere around Prague,” *Environmental monitoring and assessment*, **34** 2, 145–9, 1995. DOI: 10.1007/BF00546028.
- [213] H. Loosli, B. E. Lehmann, and W. M. Smethie, “Noble gas radio-isotopes: ^{37}Ar , ^{85}Kr , ^{39}Ar , ^{81}Kr ,” *Pages in Environmental Tracers in Subsurface Hydrology*, Springer US, 379–397, 2000. DOI: 10.1007/978-1-4615-4557-6_12.

- [214] H. Araújo, D. Akimov, E. Barnes, *et al.*, “Radioactivity backgrounds in ZEPLINIII,” *Astropart. Phys.*, **35**, 495–502, 2012. DOI: 10.1016/j.astropartphys.2011.11.001.
- [215] H. Sievers, NDS 62,271, National Nuclear Data Center, NuDat 2 database, <http://www.nndc.bnl.gov/nudat2>.
- [216] D. C. Malling, “Dark Matter Search Backgrounds from Primordial Radionuclide Chain Disequilibrium,” PhD thesis, Brown University, 2013.
- [217] A. Manzur, A. Curioni, L. Kastens, *et al.*, “Scintillation efficiency and ionization yield of liquid xenon for monoenergetic nuclear recoils down to 4 keV,” *Phys. Rev. C*, **81** 2, 025808, 2010. DOI: 10.1103/PhysRevC.81.025808.
- [218] G. Plante, E. Aprile, R. Budnik, *et al.*, “New measurement of the scintillation efficiency of low-energy nuclear recoils in liquid xenon,” *Phys. Rev. C*, **84** 4, 045805, 2011. DOI: 10.1103/PhysRevC.84.045805.
- [219] C. Savage, K. Freese, and P. Gondolo, “Annual modulation of dark matter in the presence of streams,” *Phys. Rev. D*, **74** 4, 043531, 2006. DOI: 10.1103/PhysRevD.74.043531.
- [220] R. H. Helm, “Inelastic and elastic scattering of 187-MeV electrons from selected even-even nuclei,” *Phys. Rep.*, **104** 5, 1466–1475, 1956. DOI: 10.1103/PhysRev.104.1466.
- [221] E. Armengaud *et al.*, “Final results of the EDELWEISS-II WIMP search using a 4-kg array of cryogenic germanium detectors with interleaved electrodes,” *Phys. Lett. B*, **702** 5, 329–335, 2011. DOI: 10.1016/j.physletb.2011.07.034.
- [222] CDMS II Collaboration, “Dark matter search results from the CDMS II experiment,” *Science*, **327** 5973, 1619–1621, 2010. DOI: 10.1126/science.1186112.
- [223] D. Yu. Akimov *et al.*, “WIMP-nucleon cross-section results from the second science run of ZEPLIN-III,” *Phys. Lett. B*, **709**, 14, 2012. DOI: 10.1016/j.physletb.2012.01.064.
- [224] J. Angle *et al.*, “Search for light dark matter in XENON10 data,” *Phys. Rev. Lett.*, **107** 5, 051301, 2011. DOI: 10.1103/PhysRevLett.107.051301.
- [225] C. E. Aalseth, P. S. Barbeau, J. Colaresi, *et al.*, “CoGeNT: A search for low-mass dark matter using p-type point contact germanium detectors,” *Phys. Rev. D*, **88** 1, 012002, 2013. DOI: 10.1103/PhysRevD.88.012002.

- [226] Z. Ahmed *et al.*, “Results from a Low-Energy Analysis of the CDMS II Germanium Data,” *Phys. Rev. Lett.*, **106** 13, 131302, 2011. DOI: 10.1103/PhysRevLett.106.131302.
- [227] R Agnese *et al.*, “Silicon detector dark matter results from the final exposure of CDMS II,” *Phys. Rev. Lett.*, **111** 25, 251301, 2013. DOI: 10.1103/PhysRevLett.111.251301.
- [228] G. Angloher, M. Bauer, I. Bavykina, *et al.*, “Results from 730 kg·days of the CRESST-II Dark Matter search,” *Eur. Phys. J. C*, **72** 4, 1971, 2012. DOI: 10.1140/epjc/s10052-012-1971-8.
- [229] J. D. Sloane, M. R. Buckley, A. M. Brooks, and F. Governato, “Assessing Astrophysical Uncertainties in Direct Detection with Galaxy Simulations,” *arXiv:1601.05402*, 2016.
- [230] M. T. Frandsen, F. Kahlhoefer, C. McCabe, *et al.*, “Resolving astrophysical uncertainties in dark matter direct detection,” *J. Cosmol. Astropart. Phys.*, **01**, 024, 2012. DOI: 10.1088/1475-7516/2012/01/024.
- [231] D. S. Akerib *et al.*, “Results on the Spin-Dependent Scattering of Weakly Interacting Massive Particles on Nucleons from the Run 3 Data of the LUX Experiment,” *Phys. Rev. Lett.*, **116**, 161302, 16 2016. DOI: 10.1103/PhysRevLett.116.161302.
- [232] J. Engel, S. Pittel, and P. Vogel, “Nuclear physics of dark matter detection,” *Int. J. Mod. Phys. E*, **1**, 1–37, 1992. DOI: <http://dx.doi.org/10.1142/S0218301392000023>.
- [233] V. A. Bednyakov, H. V. Klapdor-Kleingrothaus, and S. G. Kovalenko, “Direct detection of supersymmetric dark matter and the role of the target nucleus spin,” *Phys. Rev. D*, **50** 12, 1994. DOI: <http://dx.doi.org/10.1103/PhysRevD.50.7128>.
- [234] A. Birkedal, A. Noble, M. Perelstein, and A. Spray, “Little Higgs dark matter,” *Phys. Rev. D*, **74**, 035002, 3 2006. DOI: <http://dx.doi.org/10.1103/PhysRevD.74.035002>.
- [235] H. C. Cheng, J. L. Feng, and K. Matchev, “Kaluza-Klein dark matter,” *Phys. Rev. Lett.*, **89**, 211301, 2002. DOI: <http://dx.doi.org/10.1103/PhysRevLett.89.211301>.
- [236] P. Klos, J. Menéndez, D. Gazit, and A. Schwenk, “Large-scale nuclear structure calculations for spin-dependent WIMP scattering with chiral effective field theory currents,” *Phys. Rev. D*, **88** 8, 083516, 2013. DOI: 10.1103/PhysRevD.88.083516.

- [237] E. Caurier, G. Martinez-Pinedo, and F. Nowacki, "The shell model as a unified view of nuclear structure," *Rev. Mod. Phys.*, **77**, 427–488, 2005. DOI: 10.1103/RevModPhys.77.427.
- [238] M. T. Ressel and D. J. Dean, "Spin-dependent neutralino-nucleus scattering for $A \sim 127$ nuclei," *Phys. Rev. C*, **56** 1, 535, 1997. DOI: 10.1103/PhysRevC.56.535.
- [239] R. Machleidt, K. Holinde, and C. Elster, "The bonn meson-exchange model for the nucleon–nucleon interaction," *Phys. Rep.*, **149**, 1–89, 1987. DOI: 10.1016/S0370-1573(87)80002-9.
- [240] V. G. J. Stoks, R. A. M. Klomp, C. P. F. Terheggen, and J. J. de Swart, "Construction of high-quality NN potential models," *Phys. Rev. C*, **49** 6, 2950–2963, 1994. DOI: <http://dx.doi.org/10.1103/PhysRevC.49.2950>.
- [241] P. Toivanen, M. Kortelainen, J. Suhonen, and J. Toivanen, "Large-scale shell-model calculations of elastic and inelastic scattering rates of lightest supersymmetric particles (LSP) on ^{127}I , ^{129}Xe , ^{131}Xe , and ^{133}Cs nuclei," *Phys. Rev. C*, **79** 4, 044302, 2009. DOI: 10.1103/PhysRevC.79.044302.
- [242] M. Hjorth-Jensen, T. T. S. Kuo, and E. Osnes, "Realistic effective interactions for nuclear systems," *Phys. Rep.*, **261**, 125–270, 1995. DOI: doi : 10.1016/0370-1573(95)00012-6.
- [243] A. L. Fitzpatrick, W. Haxton, E. Katz, *et al.*, "The effective field theory of dark matter direct detection," *J. Cosmol. Astropart. Phys.*, **2**, 1–55, 2013. DOI: <http://dx.doi.org/10.1088/1475-7516/2013/02/004>.
- [244] J. Menéndez, D. Gazit, and A. Schwenk, "Spin-dependent WIMP scattering off nuclei," *Phys. Rev. D*, **86** 10, 103511, 2012. DOI: 10.1103/PhysRevD.86.103511.
- [245] E. Caurier, J. Menéndez, F. Nowacki, and A. Poves, "Influence of pairing on the nuclear matrix elements of the neutrinoless $\beta\beta$ decays," *Phys. Rev. Lett.*, **100** 5, 052503, 2008. DOI: 10.1103/PhysRevLett.100.052503.
- [246] J. Menéndez, A. Poves, E. Caurier, and F. Nowacki, "Disassembling the nuclear matrix elements of the neutrinoless $\beta\beta$ decay," *Nucl. Phys. A*, **818**, 139–151, 2009. DOI: 10.1016/j.nuclphysa.2008.12.005.
- [247] L. Vietze, P. Klos, J. Menéndez, *et al.*, "Nuclear structure aspects of spin-independent WIMP scattering off xenon," *Phys. Rev. D*, **91** 4, 043520, 2015. DOI: 10.1103/PhysRevD.91.043520.

- [248] G Dūda, A. Kemper, and P. Gondolo, "Model-independent form factors for spin-independent neutral nucleon scattering from elastic electron scattering data," *J. Cosmol. Astropart. Phys.*, **3**, 2007. DOI: <http://dx.doi.org/10.1088/1475-7516/2007/04/012>.
- [249] M. C. Smith, G. R. Ruchti, A. Helmi, *et al.*, "The RAVE survey: constraining the local Galactic escape speed," *Mon. Not. R Astron. Soc.*, **379**, 755–772, 2007. DOI: [10.1111/j.1365-2966.2007.11964.x](https://doi.org/10.1111/j.1365-2966.2007.11964.x).
- [250] C McCabe, "Astrophysical uncertainties of dark matter direct detection experiments," *Phys. Rev. D*, **82** 2, 023530, 2010. DOI: [10.1103/PhysRevD.82.023530](https://doi.org/10.1103/PhysRevD.82.023530).
- [251] R. Barlow, "Extended maximum likelihood," *Nucl. Instrum. Meth. A*, **297**, 496–506, 1990. DOI: [10.1016/0168-9002\(90\)91334-8](https://doi.org/10.1016/0168-9002(90)91334-8).
- [252] L Moneta, K Belasco, and K Cranmer, "The RooStats project," *Proceedings of Science*, **57**, 2010.
- [253] W Verkerke and D Kirkby, "The RooFit toolkit for data modeling," *Tech. Rep. SLAC, Stanford*, 2003.
- [254] C. Lee, "Mitigation of backgrounds for the Large Underground Xenon dark matter experiment," PhD thesis, Case Western Reserve University, 2015.
- [255] D. S. Akerib *et al.*, "Calibration, event reconstruction, data analysis and limits calculation for the LUX dark matter experiment", (in preparation).
- [256] D. S. Akerib *et al.*, "Low-energy (0.7–74 keV) nuclear recoil calibration of the LUX dark matter experiment using D-D neutron scattering kinematics," *arXiv:1608.05381*, 2016.
- [257] J. Thomas and D. A. Imel, "Recombination of electron-ion pairs in liquid argon and liquid xenon," *Phys. Rev. A*, **36**, 614–616, 2 1987. DOI: [10.1103/PhysRevA.36.614](https://doi.org/10.1103/PhysRevA.36.614).
- [258] A. Hitachi, "Properties of liquid xenon scintillation for dark matter searches," *Astropart. Phys.*, **24** 3, 247–256, 2005. DOI: [10.1016/j.astropartphys.2005.07.002](https://doi.org/10.1016/j.astropartphys.2005.07.002).
- [259] M. S. J. Lindhard V. Nielsen and P. V. Thomsen, "Integral equations governing radiation effect," *Mat. Fis. Medd. Dan. Vid. Selsk.*, **33** 10, 1963.
- [260] F. Bezrukov, F. Kahlhoefer, and M. Lindner, "Interplay between scintillation and ionization in liquid xenon Dark Matter searches," *Astroparticle Physics*, **35** 3, 119–127, 2011. DOI: [10.1016/j.astropartphys.2011.06.008](https://doi.org/10.1016/j.astropartphys.2011.06.008).

- [261] D.-M. Mei, Z. B. Yin, L. Stonehill, and A. Hime, "A model of nuclear recoil scintillation efficiency in noble liquids," *Astroparticle Physics*, **30** 1, 12–17, 2008. DOI: 10.1016/j.astropartphys.2008.06.001.
- [262] C. H. Faham, V. M. Gehman, A. Currie, *et al.*, "Measurements of wavelength-dependent double photoelectron emission from single photons in vuv-sensitive photomultiplier tubes," *JINST*, **10** 09, P09010, 2015. DOI: <http://dx.doi.org/10.1088/1748-0221/10/09/P09010>.
- [263] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, "Power-constrained limits," *arXiv:1105.3166*, **2**, 1–11, 2011.
- [264] T. Shutt, C. Dahl, J. Kwong, *et al.*, "Performance and Fundamental Processes at Low Energy in a Two-Phase Liquid Xenon Dark Matter Detector," *Nucl. Phys. B: Proc. Suppl.*, **173**, 160–163, 2007. DOI: 10.1016/j.nuclphysbps.2007.08.140.
- [265] H. Lee *et al.*, "Limits on Interactions between Weakly Interacting Massive Particles and Nucleons Obtained with CsI(Tl) Crystal Detectors," *Phys. Rev. Lett.*, **99** 9, 091301, 2007. DOI: 10.1103/PhysRevLett.99.091301.
- [266] S. Archambault *et al.*, "Constraints on low-mass WIMP interactions on ^{19}F from PICASSO," *Phys. Lett. B*, **711** 2, 153–161, 2012. DOI: 10.1016/j.physletb.2012.03.078.
- [267] J. Angle, E. Aprile, F. Arneodo, *et al.*, "Limits on Spin-Dependent WIMP-Nucleon Cross Sections from the XENON10 Experiment," *Phys. Rev. Lett.*, **101** 9, 091301, 2008. DOI: 10.1103/PhysRevLett.101.091301.
- [268] E. Aprile *et al.*, "Limits on Spin-Dependent WIMP-Nucleon Cross Sections from 225 Live Days of XENON100 Data," *Phys. Rev. Lett.*, **111** 2, 021301, 2013. DOI: 10.1103/PhysRevLett.111.021301.
- [269] V. Lebedenko *et al.*, "Limits on the Spin-Dependent WIMP-Nucleon Cross Sections from the First Science Run of the ZEPLIN-III Experiment," *Phys. Rev. Lett.*, **103** 15, 151302, 2009. DOI: 10.1103/PhysRevLett.103.151302.
- [270] 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.*, **114** April, 141301, 2015. DOI: 10.1103/PhysRevLett.114.141301.
- [271] D. R. Tovey, R. J. Gaitskell, P. Gondolo, *et al.*, "A new model-independent method for extracting spin-dependent cross section limits from dark matter searches," *Phys. Lett. B*, **488** May, 17, 2000. DOI: [doi:10.1016/S0370-2693\(00\)00846-7](https://doi.org/10.1016/S0370-2693(00)00846-7).

- [272] F. Giuliani and T. A. Girard, "Model-independent limits from spin-dependent WIMP dark matter experiments," *Phys. Rev. D*, **71** 12, 123503, 2005. DOI: 10.1103/PhysRevD.71.123503.
- [273] 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** 1, 011301, 2009. DOI: 10.1103/PhysRevLett.102.011301.
- [274] G Wikström and J Edsjö, "Limits on the WIMP-nucleon scattering cross-section from neutrino telescopes," *Journal of Cosmology and Astroparticle Physics*, **2009** 04, 009–009, 2009. DOI: 10.1088/1475-7516/2009/04/009.
- [275] R. H. Fowler and L. Nordheim, "Electron emission in intense electric fields," *Proc. R Soc. A*, **119** 781, 173–181, 1928. DOI: 10.1098/rspa.1928.0091.
- [276] L. W. Nordheim, "The effect of the image force on the emission and reflexion of electrons by metals," *Proc. R Soc. London A*, **121** 788, 626–639, 1928.
- [277] R. J. Noer, "Electron field emission from broad-area electrodes," *Appl. Phys. A Solids and Surfaces*, **28** 1, 1–24, 1982. DOI: 10.1007/BF00617778.
- [278] D. W. Williams and W. T. Williams, "Prebreakdown and breakdown characteristics of stainless steel electrodes in vacuum," *J. Phys. D: Appl. Phys.*, **7** 8, 1173, 1974. DOI: <http://dx.doi.org/10.1088/0022-3727/7/8/315>.
- [279] B. M. Cox and W. T. Williams, "Field-emission sites on unpolished stainless steel," *J. Phys. D: Appl. Phys.*, **10**, L5, 1977. DOI: <http://dx.doi.org/10.1088/0022-3727/10/3/001>.
- [280] R. P. Little and W. T. Whitney, "Electron Emission Preceding Electrical Breakdown in Vacuum," *J. Appl. Phys.*, **34** 8, 2430, 1963. DOI: 10.1063/1.1702760.
- [281] R. P. Little and S. T. Smith, "Field Enhancing Projections Produced by the Application of an Electric Field," *J. Appl. Phys.*, **36** 4, 1502, 1965. DOI: 10.1063/1.1714350.
- [282] H. E. Tomaschke, "Role of Submicroscopic Projections in Electrical Breakdown," *J. Vac. Sci. Tech.*, **4** 4, 192, 1967. DOI: 10.1116/1.1492545.
- [283] R. Hackam, "Structure of stainless steel electrodes subjected to very high positive and negative voltages," *Mat. Res. Bull.*, **7** 6, 655–661, 1972. DOI: doi:10.1016/0025-5408(72)90053-0.

- [284] M. Jimenez, R. J. Noer, G. Jouve, *et al.*, "Electron field emission from selectively contaminated cathodes," *J. Phys. D: Appl. Phys.*, **26**, 1503–1509, 1993. DOI: <http://dx.doi.org/10.1088/0022-3727/26/9/026>.
- [285] M. Jimenez, R. J. Noer, G. Jouve, *et al.*, "Electron field emission from large-area cathodes: evidence for the projection model," *J. Phys. D: Appl. Phys.*, **1038**, 1994. DOI: <http://dx.doi.org/10.1088/0022-3727/27/5/023>.
- [286] G. A. Farrall, "Techniques for the Study of Breakdown between Large-Area Electrodes in Vacuum," *J. Appl. Phys.*, **43** 3, 938, 1972. DOI: 10.1063/1.1661310.
- [287] N. K. Allen, R. V. Latham, and B. M. Cox, "The source of high-beta electron emission sites on broad-area high-voltage alloy electrodes," *J. Phys. D: Appl. Phys.*, **12** 6, 969, 1979. DOI: <http://dx.doi.org/10.1088/0022-3727/12/6/020>.
- [288] R. Hurley, "Electrical phenomena occurring at the surface of electrically stressed metal cathodes. II. Identification of electroluminescent (k-spot) radiation with electron emission on," *J. Phys. D: Appl. Phys.*, **2247**, 1979. DOI: <http://dx.doi.org/10.1088/0022-3727/12/12/027>.
- [289] R. V. Latham, "The origin of prebreakdown electron emission from vacuum-insulated high voltage electrodes," *Vacuum*, **32**, 3–6, 1982. DOI: 10.1016/0042-207X(82)80043-2.
- [290] C. B. Duke, "Field Emission through Atoms Adsorbed on a Metal Surface," *J. Chem. Phys.*, **46** 3, 923, 1967. DOI: 10.1063/1.1840828.
- [291] B. M. Cox, "The nature of field emission sites," *J. Phys. D: Appl. Phys.*, **2065**, 1975. DOI: <http://dx.doi.org/10.1088/0022-3727/8/17/008>.
- [292] L. W. Swanson and L. C. Crouser, "Total-energy distribution of field-emitted electrons and single-plane work functions for tungsten," *Phys. Rep.*, **163**, 622–641, 3 1967. DOI: 10.1103/PhysRev.163.622.
- [293] K. McDonald. (2003). Notes on electrostatic wire grids, [Online]. Available: <http://www.hep.princeton.edu/~mcdonald/examples/grids.pdf>.
- [294] R. Veenhof *et al.*, *Garfield++*, <http://garfieldpp.web.cern.ch/>.
- [295] G. D. Alkhasov, "Statistics of electron avalanches and ultimate resolution of proportional counters," *Nucl. Instrum. Meth.*, **89**, 155–165, 1970. DOI: 10.1016/0029-554X(70)90818-9.
- [296] W. Blum, W. Riegler, and L. Rolandi, "Particle detection with drift chambers," *Springer*, 2008.

- [297] T. Zerguerras, B. Genolini, F. Kuger, *et al.*, "Understanding avalanches in a Micro-megas from single-electron response measurement," *Nucl. Instrum. Meth. A*, **772**, 76–82, 2015. DOI: 10.1016/j.nima.2014.11.014.
- [298] W. P. Dyke and J. K. Trolan, "Field emission: large current densities, space charge, and the vacuum arc," *Phys. Rev. Lett.*, **80**, 799–808, 1953. DOI: 10.1103/PhysRev.89.799.
- [299] C. D. Ehrlich and E. W. Plummer, "Measurement of the absolute tunneling current density in field emission from tungsten(110)," *Phys. Rev. B*, **18**, 8 1978. DOI: 10.1103/PhysRevB.18.3767.
- [300] W. Tauchert, "Photoelectric determination of V_0 values and electron ranges in some cryogenic liquids," *Canadian J. Chem.*, **55** 11, 1860–1866, 1977. DOI: 10.1139/v77-260.
- [301] D. Alpert, "Initiation of Electrical Breakdown in Ultrahigh Vacuum," *J. Vacuum Sci. Tech.*, **1** 2, 35, 1964. DOI: 10.1116/1.1491722.
- [302] P. A. Chatterton, "A theoretical study of field emission initiated vacuum breakdown," *Proc. Phys. Soc.*, **88** 231, 1966. DOI: <http://dx.doi.org/10.1088/0370-1328/88/1/326>.
- [303] R. Hackam, "Formation of microprotrusions on a stainless steel cathode in high vacuum," *Mat. Res. Bull.*, **8** 7, 868–870, 1973. DOI: 10.1016/0025-5408(73)90196-7.
- [304] R. Hackam, "Determination of the electric field enhancement factor and crater dimensions in aluminum from scanning electron micrographs," *J. Appl. Phys.*, **45** 1, 114, 1974. DOI: 10.1063/1.1662944.
- [305] H. Biederman, "Metal-insulator-metal sandwich structures with anomalous properties," *Vacuum*, **26** 12, 513–523, 1976. DOI: 10.1016/S0042-207X(76)81130-X.
- [306] G. Dearnaley, "Electronic conduction through thin unsaturated oxide layers," *Phys. Lett. A*, **25A** 10, 3–4, 1967. DOI: [http://dx.doi.org/10.1016/0375-9601\(67\)90982-6](http://dx.doi.org/10.1016/0375-9601(67)90982-6).
- [307] R. E. Hurley and P. J. Dooley, "Electroluminescence produced by high electric fields at the surface of copper cathodes," *J. Phys. D: Appl. Phys.*, **10**, L195, 1977. DOI: <http://dx.doi.org/10.1088/0022-3727/10/15/002>.
- [308] E. Aprile, H. Contreras, L. W. Goetzke, *et al.*, "Measurements of proportional scintillation and electron multiplication in liquid xenon using thin wires," *JINST*, **9** 11, P11012, 2014. DOI: <http://dx.doi.org/10.1088/1748-0221/9/11/P11012>.

- [309] E. Aprile, K. Arisaka, F. Arneodo, *et al.*, “The XENON100 dark matter experiment,” *Astropart. Phys.*, **35** 9, 573–590, 2012. DOI: 10.1016/j.astropartphys.2012.01.003.
- [310] D. Yu. Akimov *et al.*, “The ZEPLIN-III dark matter detector: Instrument design, manufacture and commissioning,” *Astropart. Phys.*, **27**, 46–60, 2007. DOI: 10.1016/j.astropartphys.2006.09.005.
- [311] G. Alner *et al.*, “First limits on WIMP nuclear recoil signals in ZEPLIN-II: A two-phase xenon detector for dark matter detection,” *Astroparticle Physics*, **28** 3, 287–302, 2007. DOI: 10.1016/j.astropartphys.2007.06.002.
- [312] E. M. Gushchin, A. A. Kruglov, and I. M. Obodovskii, “Electron dynamics in condensed argon and xenon,” *Zh, Eksp. Teor. Fiz*, **58** 4, 1114–1125, 1982.
- [313] S. F. Biagi, “Monte Carlo simulation of electron drift and diffusion in counting gases under the influence of electric and magnetic fields,” *Nucl. Instrum. Meth. A*, **421**, 234–240, 1–2 1999. DOI: 10.1016/S0168-9002(98)01233-9.
- [314] C. A. B. Oliveira, M. Sorel, J. Martin-Albo, *et al.*, “Energy resolution studies for NEXT,” *JINST*, **6** 05, P05007, 2011. DOI: <http://dx.doi.org/10.1088/1748-0221/6/05/P05007>.
- [315] J. E. Renner, “High Pressure Xenon Detectors for Rare Physics Searches,” PhD thesis, University of California Berkeley, 2014.
- [316] T. Doke, A. Hitachi, S. Kubota, *et al.*, “Estimation of Fano factors in liquid argon, krypton, xenon and xenon-doped liquid argon,” *Nucl. Instrum. Meth.*, **134**, 353–357, 2 1976. DOI: 10.1016/0029-554X(76)90292-5.
- [317] M. Redshaw, E. Wingfield, J. McDaniel, and E. Myers, “Mass and Double-Beta-Decay Q-Value of ^{136}Xe ,” *Phys. Rev. Lett.*, **98** 5, 053003, 2007. DOI: 10.1103/PhysRevLett.98.053003.
- [318] R. Veenhof, *Magboltz*, <http://consult.cern.ch/writeup/magboltz/>.
- [319] IT Center for Science, *ELMER*, <http://www.csc.fi/english/pages/elmer>.
- [320] C. Geuzaine and J. Remacle, *GMSH*, <http://geuz.org/gmsh/>.
- [321] M. M. Aggarwal, S. K. Badyal, V. S. Bhatia, *et al.*, “A honeycomb proportional counter for photon multiplicity measurement in the ALICE experiment,” *Nucl. Instrum. Meth. A*, **488**, 131–143, 2002. DOI: [http://dx.doi.org/10.1016/S0168-9002\(02\)00482-5](http://dx.doi.org/10.1016/S0168-9002(02)00482-5).

- [322] A. Rubin, L. Arazi, S. Bressler, *et al.*, “Optical readout: a tool for studying gas-avalanche processes,” *JINST*, **8** 08, P08001, 2013. DOI: <http://dx.doi.org/10.1088/1748-0221/8/08/P08001>.
- [323] C. A. B. Oliveira, P. M. M. Correia, H. Schindler, *et al.*, “Simulation of VUV electroluminescence in micropattern gaseous detectors: the case of GEM and MHSP,” *JINST*, **7** 09, P09006, 2012. DOI: <http://dx.doi.org/10.1088/1748-0221/7/09/P09006>.
- [324] C. A. B. Oliveira, H. Schindler, R. J. Veenhof, *et al.*, “A simulation toolkit for electroluminescence assessment in rare event experiments,” *Phys. Lett. B*, **703**, 217–222, 2011. DOI: [10.1016/j.physletb.2011.07.081](https://doi.org/10.1016/j.physletb.2011.07.081).
- [325] S. R. Hunter, J. G. Carter, and L. G. Christophorou, “Low-energy electron drift and scattering in krypton and xenon,” *Phys. Rev. A*, **38** 11, 5539–5551, 1988. DOI: <http://dx.doi.org/10.1103/PhysRevA.38.5539>.
- [326] M. Suzuki and T. Taniguchi, “Momentum transfer cross section of xenon deduced from electron drift velocity data,” *J. Phys. D: Appl. Phys.*, **25**, 50–56, 1992. DOI: <http://dx.doi.org/10.1088/0022-3727/25/1/007>.
- [327] Y. Mei, “Direct Dark Matter Search with the XENON100 Experiment,” PhD thesis, Rice University, 2011.
- [328] N. E. Kuz'menko, “Electronic transition probabilities and lifetimes of electronically excited states of diatomic molecules,” *Sov. Phys. Usp.*, **22** 3, 160–173, 1979. DOI: <http://dx.doi.org/10.1070/PU1979v022n03ABEH005425>.
- [329] E. Aprile, M. Alfonsi, K. Arisaka, *et al.*, “Observation and applications of single-electron charge signals in the XENON100 experiment,” *J. Phys. G: Nuc. Phys.*, **41** 3, 035201, 2014. DOI: [10.1088/0954-3899/41/3/035201](https://doi.org/10.1088/0954-3899/41/3/035201).
- [330] D. S. Akerib *et al.*, “LUX Run 3 ER yields and recombination,” in preparation.

Appendix A

Figure Copyright

Please see the Tables on the following pages for the copyright information for Figures reproduced from other sources in this thesis.

Page No.	Name of work	Source of work	Copyright holder and contact	Permission requested on	Permission obtained?	Permission note
25	Fig. 2.1, the rotation curve of NGC 6503...	K. Freese, Review of observational evidence for dark matter in the universe and in upcoming searches for dark stars, EAS Publications Series, 36, 113126, 2009 DOI: http://10.1051/eas/0936016 .	EDP Sciences, permission@edpsciences.org	26/07/2016	Yes	Written permission from publisher.
27	Fig. 2.2, composite image of the Buller Cluster collision ...	NASA/CXC/SAO http://chandra.harvard.edu/photo/2006/1e0657/index.html	NASA/CXC/SAO cxcedu@cfa.harvard.edu	26/07/2016	Not yet	Request sent to http://chandra.harvard.edu/photo/image_use.html
29	Fig. 2.3, the observed CMB power spectrum from Planck ...	Planck Collaboration, Planck 2015 results. XIII. Cosmological parameters, Astron. & Astrophys., Forthcoming Article, 2016. DOI: http://dx.doi.org/10.1051/0004-6361/201525830	EDP Sciences, Astronomy & Astrophysics, aanda.paris@obspm.fr	09/08/2016	Yes	Written permission from publisher.
46, 51, 52	Figs. 3.2, 3.4, and 3.5, LUX detector diagrams.	D. S. Akerib <i>et al.</i> , The Large Underground Xenon (LUX) experiment, Nucl. Instrum. Meth. A, 704, 111–126, 2013. DOI: http://10.1016/j.nima.2012.11.135 .	Elsevier, via http://www.copyright.com/about/contact/	26/07/2016	Yes	License number: 3916401213054

Page No.	Name of work	Source of work	Copyright holder and contact	Permission requested on	Permission obtained?	Permission note
57	Fig. 3.7, data flow through the DAQ of the LUX experiment ...	D. S. Akerib <i>et al.</i> , "Data acquisition and readout system for the LUX dark matter experiment, Nucl. Instrum. Meth. A, 668, 1–8, 2012. DOI: http://dx.doi.org/10.1016/j.nima.2011.11.063 .	Elsevier, via http://www.copyright.com/about/contact/	09/08/2016	Yes	License number: 3924750886043
82	Fig. 4.19, the measured tritium energy spectrum from LUX ...	D. S. Akerib <i>et al.</i> , Tritium calibration of the LUX dark matter experiment, Phys. Rev. D, 93, 072009, 2016. DOI: http://dx.doi.org/10.1103/PhysRevD.93.072009 .	American Physical Society, via http://www.copyright.com/about/contact/	14/08/2016	Yes	License number: 3927731494179
82, 103, 104, 106, 107	Figs. 4.20, 4.42, 4.43, 4.44, and 4.45, first results from LUX.	D. S. Akerib <i>et al.</i> , First results from the LUX dark matter experiment at the Sanford Underground Research Facility, Phys. Rev. Lett., 112, 091303, 2014. DOI: http://10.1103/PhysRevLett.112.091303 .	American Physical Society, via http://www.copyright.com/about/contact/	09/08/2016	Yes	License number: 3924751165039
140, 141	Figs. 5.12 and 5.13, spin-dependent analysis results from LUX	D. S. Akerib <i>et al.</i> , Results on the Spin-Dependent Scattering of Weakly Interacting Massive Particles on Nucleons from the Run 3 Data of the LUX Experiment, Phys. Rev. Lett., 116, 161302, 16 2016. DOI: http://10.1103/PhysRevLett.116.161302 .	American Physical Society, via http://www.copyright.com/about/contact/	09/08/2016	Yes	License number: 3924750566592